



HAL
open science

Analysis of Stochastic Algorithms for Sampling and Riemannian Approximation

Pablo Jiménez

► **To cite this version:**

Pablo Jiménez. Analysis of Stochastic Algorithms for Sampling and Riemannian Approximation. Machine Learning [stat.ML]. Institut Polytechnique de Paris, 2023. English. NNT : 2023IPPAX051 . tel-04214337v2

HAL Id: tel-04214337

<https://hal.science/tel-04214337v2>

Submitted on 21 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2023IPPAX051

Thèse de doctorat



Analysis of Stochastic Algorithms for Sampling and Riemannian Approximation

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à École polytechnique

École doctorale n°574 Ecole Doctorale de Mathématiques Hadamard (EDMH)
Spécialité de doctorat : Mathématiques appliquées

Thèse présentée et soutenue à Palaiseau, le 30 juin 2023, par

PABLO JIMÉNEZ

Composition du Jury :

Anne Auger Directrice de recherche INRIA, École polytechnique (UMR 7641)	Présidente du jury
Marc Arnaudon Professeur, Université de Bordeaux (UMR 5251)	Rapporteur
Konstantinos Zygalakis Reader, University of Edinburgh	Rapporteur
Stéphanie Allasonnière Professeure, Université Paris Cité (UMR S1138)	Examinatrice
Aymeric Dieuleveut Assistant Professor, École Polytechnique (UMR 7641)	Examineur
Éric Moulines Professeur, École polytechnique (UMR 7641)	Directeur de thèse
Alain Durmus Professeur, École polytechnique (UMR 7641)	Invité, Co-directeur
Salem Said Chargé de recherche CNRS, Université Grenoble Alpes (UMR 5224)	Invité

PABLO

*La del alba sería
por La Cañada.
Salvador y Sofía
se desvelaban
esperando al lucero
de la mañana.*

*Con lápiz de colores
Martín pintaba
la cara del hermano
con quien soñaba.*

*Cuando la primavera
luce sus galas,
llegó la vida nueva
que se anunciaba.*

*Ramos de rosas rojas
tiene la casa
y guirnaldas de flores
en las ventanas.*

*En la cuna hay un ángel
lleno de gracia:
es un niño moreno,
Pablo se llama.
— Salvador Jiménez*

Remerciements

Tout d'abord, je voudrais remercier mes directeurs de thèse Alain et Eric. Quand vous m'avez proposé de faire une thèse en lien avec la géométrie riemannienne, je sentais qu'on partait dans une grande aventure. Le sujet a été plus technique que ce que je pensais au début, mais on y est arrivés tout de même. Merci pour avoir cru en moi et m'avoir poussé quand il le fallait pour finir la thèse.

Alain, je te suis très reconnaissant pour tout ce que tu m'as appris, dès le stage en MCMC, j'ai bien vu que je pouvais compter sur toi pour avoir une intuition ou sortir un vieux livre soviétique qui contient le théorème dont on a besoin. Tu as su être patient avec moi quand je cherchais mes repères et ça m'a permis d'être à l'aise par la suite et de travailler dans une bonne ambiance. Je garde le souvenir d'une grande satisfaction lorsqu'on a revisité la grassmannienne et qu'on a compris tous les calculs qui semblaient si compliqués un an avant, tout comme en finissant la preuve du cut locus. Merci pour tout le temps passé à lire et refaire les preuves du livre de Lee.

Eric, merci pour ton énergie incessante, tu m'as montré ce qu'est un chercheur dévoué qui veut toujours comprendre une nouvelle branche de son domaine, se lancer dans un projet de plus. Tu m'as donné la possibilité de rencontrer beaucoup de gens et d'endroits pendant ma thèse, malgré les limitations du covid, et je t'en remercie. Ta franchise, bien qu'elle ait été un peu déstabilisante au début, m'a permis de trouver ma place dans cette codirection avec Alain, tout en sachant que vous formez une équipe très soudée.

I also want to thank the jury: Stéphanie, Aymeric, Anne and Salem, with a special thanks to Marc and Kostas for giving me honest and positive feedback about my thesis.

Merci à Salem pour m'avoir tant appris sur la Géométrie riemannienne, toujours accompagnée de ta patience et bienveillance. Je n'y serai pas arrivé sans ton aide. Thank you Francesca for your patience and your relentless work on our project on MCMC even though I know I was not always the best coauthor.

Je voudrais remercier ma famille, ma maman pour sa tendresse, sa capacité à me faire sourire, sa sincérité. Gracias por ser mi fan incondicional, se que te mueres de ganas de gritar "¡Ése es mi hijo!" o incluso "¡Ése es mi *hijito!*". Nunca dejaré de ser tu hijo pequeño. Mon papa, pour son humour et sa tendresse, ainsi que pour les recettes de cuisine. Quand tu m'emmenais dans ton bureau à l'université dès mon adolescence, je pensais déjà que j'étais à l'endroit le plus cool qui puisse exister : un bureau avec un tableau à craie. Gracias por haber compartido mi ilusión por las mates y haberme llevado tantos sábados a ESTALMAT. Mon frère pour être un si bon grand frère et avoir toujours cru en moi.

Merci beaucoup Iris pour toute la joie, la tendresse et l'amour que tu m'apportes. Ton soutien m'a permis de finir la thèse sereinement et j'ai hâte de partager mon quotidien avec toi.

Je remercie aussi mes amis, ceux de longue date qui trouvaient un moment pour qu'on se voie comme ceux qui sont apparus pendant ces années. Para mis amigos de Villanueva y alrededores, gracias a Carlos, María, Belén, Litos, Alitzel, Michel, Elvira, Alejandro y Juan por seguir siendo la pandilla que tantas ganas tengo de ver cada vez que vengo por Madrid. Merci à Romain pour avoir gardé contact après tant d'années.

Je remercie mes amis d'Orsay, Quentin, Christèle, Alexia et Florian. Merci Paul pour ton soutien, tu es toujours là pour moi et ça compte beaucoup. Merci Fabien pour les sorties, les memes et les conseils de lecture, merci Thomas pour m'avoir initié aux échecs et pour avoir partagé tant de centres d'intérêt. Je suis désolé d'avoir été méchant avec toi, je te vois comme mon égal et je suis très mauvais perdant. Courage à tous les deux pour vos thèses !

Merci à Aude, Charles, Edgar et Jude pour m'avoir accueilli avec le cœur grand ouvert.

Je remercie aussi Yuliia et Arnaud, Vitalina et Martina. Votre compagnie et amitié m'a apporté beaucoup.

Muchas gracias a Juan Diego, tu compañía ha sido especial, tengo ganas de ver cual será la próxima aventura !

Un très grand merci à Aliaume, tu sais que je t'adore, tu m'as beaucoup soutenu et compris, je suis vraiment chanceux de t'avoir comme ami. Merci aussi à Hélène et Pierre-Etienne, pour ce qu'on a partagé pendant la CIG et ce qui a suivi, la musique et les bars.

Merci à tous mes amis de la danse en particulier mon Ophélie, Basile et Wissem pour le bon temps passé sur les pistes et ailleurs.

Gracias Anna por haberme ayudado a entender tantas cosas sobre mí. Eres un cielo y me has ayudado a hacer las paces conmigo mismo y a quererme más.

I want also to thank all my hanab friends for the laughs, the bombs, the wins and everything in between. In particular Animex, Anti, Paula, Amy, Catula, Jack, Sucubis, Nitrate, Melwen, Wump, Xmil, Timehoodie and timotree. I found this community when I most needed it, and you have been very welcoming. I always felt at home playing with you.

Je voudrais aussi remercier tous mes professeurs de mathématiques, du collège à la prépa et après à l'université, pour m'avoir contaminé leur passion pour les mathématiques, chacun à sa façon. Je ne pouvais pas oublier de remercier Christophe Bonnefoy, qui dès la sixième me faisait rêver de maths avec des énigmes, puis tout le lycée m'a appris l'importance de la constance avec des interros de cours tous les mardis. Tu m'avais dit une fois que tu n'avais pas eu d'impact sur ma note de maths au BAC, mais – de la même façon que les maths ne se résument pas aux nombres – ton influence est allée bien au delà d'une note. Sache que ton attitude décontractée, ton charme – et ta rigueur – ont beaucoup aidé à ce que les maths soient quelque chose de *branché*. Je remercie aussi mes professeurs de prépa Roger Mansuy et Bernard Randé pour m'avoir si bien formé. Vous savez rendre la salle de mathématiques un endroit spécial où l'on a envie de retourner. Je

suivais vos cours comme si c'était la meilleure série Netflix.

Je veux maintenant remercier mes collègues pour m'avoir aidé à tenir jusqu'ici. En ordre chronologique, merci à Valentin, Tina, Xavier, Pierre, Roger, Mariano, Valéry, Jérémy et tous les gens du Centre Borelli pour avoir animé mes journées pendant mon stage et le début de ma thèse. Merci à Achille, Apolline, Baptiste, Benjamin, Claire, Clément, Clément, Constantin, Corentin, Guillaume, Louis, Margaux, Solange et tous les autres doctorants du CMAP avec qui j'ai pu discuter, jouer à des jeux de société, rire et manger. Je n'ai pas été très assidu aux bureaux de l'X mais vous m'avez toujours accueilli dans la joie et la bonne humeur. Merci à Gabriel, Mehdi, Louis, Matthew, Maxence, Mehdi, Pierre, Thomas, Tom, Valentin, Vincent et Yazid, ainsi qu'à l'équipe du centre Lagrange pour avoir ouvert ses portes pendant la dure période du covid afin de partager autant de choses entre doctorants. Le séjour au Maroc n'aurait pas été le même sans vous. Merci aussi aux gens du LPSM, Alexis, Antonio, Ariane, Iqraa, Miguel, Camila, Grâce, Ludovic, Paul, Pierre, Francesco, Lucas, Kimia, Nathan, Ferdinand et Nicolai, ainsi qu'aux permanents Arnaud, Anna, Claire, Hugues, Maxime, Stéphane, Sylvain, Hugues, Etienne et tous ceux que j'ai oublié, mais qui ont fait partie du quotidien dans les locaux du LPSM. Vous m'avez apporté une bouffée d'air frais, et je vous en remercie.

Contents

1	Introduction	1
1.1	Introduction to stochastic approximation	1
1.1.1	Motivation	1
1.1.2	Noise Assumptions	4
1.1.3	Convergence techniques	5
1.1.4	SGD assumptions	5
1.1.5	Projected scheme	6
1.2	Stochastic approximation on Riemannian manifolds	7
1.2.1	Geodesic versus retraction schemes	7
1.2.2	First asymptotic convergence result	9
1.2.3	Non-asymptotic analysis	11
1.2.4	Momentum and Nesterov's Accelerated Gradient in the deterministic case	11
1.3	Introduction to Markov chain Monte Carlo methods	13
1.3.1	The problem of sampling from a target distribution	13
1.3.2	Monte Carlo algorithms and their Markovian counterparts	14
1.3.3	The Metropolis-Hastings procedure for sampling	15
1.4	Asymptotic analysis of Metropolis-Hastings algorithms: optimal scaling	17
1.4.1	Introduction to optimal scaling: the symmetric random walk Metropolis algorithm	17
1.4.2	Optimal scaling of MALA	20
1.5	Setting of the Thesis and Contributions	21
1.5.1	Riemannian schemes	21
1.5.2	Noise assumptions	22
1.5.3	Convergence Techniques	23
1.5.4	Results of Chapters 2 and 3	23
1.5.5	Results of Chapter 4	26
2	Convergence Analysis of Riemannian Stochastic Approximation Schemes	29
2.1	Introduction	31
2.2	Geodesic Stochastic Approximation Schemes	33
2.3	General Retraction SA Schemes	38

2.3.1	Retraction and Quantitative Estimates of the Retraction Error . . .	38
2.3.2	Analysis of Retraction SA Schemes	40
2.4	Applications	41
2.4.1	Principal Component Analysis	42
2.4.2	Robust Barycenter in a Hadamard Manifold	43
2.4.3	Numerical Examples on the PCA Problem	45
2.5	Proofs for Section 2.2	46
2.5.1	Proof of Lemma 2.1	46
2.5.2	A technical lemma: proof of Lemma 2.3	46
2.5.3	Proof of Theorem 2.4	47
2.5.4	Proof of Proposition 2.6	48
2.5.5	Proof of Theorem 2.5	51
2.6	Proofs for Section 2.3.1	54
2.6.1	Proof of Lemma 2.7	54
2.6.2	Projective Retraction on the Sphere	56
2.6.3	Proof of Proposition 2.10	58
2.7	Complementary proofs for Section 2.3.1	60
2.7.1	Proof of Proposition 2.18	60
2.7.2	Details for the Grassmann manifold	61
2.8	Proofs for Section 2.3.2	81
2.8.1	A technical lemma	81
2.8.2	Proof of Theorem 2.11	83
2.8.3	Proof of Theorem 2.12	85
2.9	Proof of Section 2.4	92
2.9.1	Proofs of Section 2.4.1	92
2.9.2	Proofs of Section 2.4.2	96
2.9.3	Details and specification for Section 2.4.3	100
2.10	Useful results from matrix analysis	101
2.11	Preliminaries on Riemannian Geometry	103
2.11.1	Metric Tensor and Distance	104
2.11.2	Levi-Civita Connection	104
2.11.3	Geodesic Equation	105
2.11.4	Parallel Transport and Parallel Frames	105
2.11.5	Riemannian Exponential Map and Cut Locus	106
2.11.6	Injectivity Domain	107
2.11.7	Isometries and Homogeneous Spaces	107
2.11.8	Riemannian Gradient, Hessian	108
2.11.9	Taylor formula along a geodesic	109
2.11.10	Bounded Hessian Implies Lipschitz Gradient	111

3	On Riemannian Stochastic Approximation Schemes with Fixed Step-Size	113
3.1	Introduction	114
3.2	Constant stepsize analysis for a constrained scheme	117
3.2.1	Main Results	117
3.2.2	Two Examples of Lyapunov Functions	120
3.3	Asymptotic expansion and law in the unconstrained case	122
3.3.1	Asymptotic Expansion as $\eta \rightarrow 0$	122
3.3.2	A Central Limit Theorem on $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$	123
3.4	Application to SGD	124
3.5	Numerical Experiments	127
3.6	Assumptions	130
3.7	Supplementary notation	131
3.8	Proofs of Section 3.2	132
3.8.1	Proof of Theorem 3.1	132
3.8.2	An alternative to Theorem 3.1-(b)	133
3.8.3	Proof of Theorem 3.2	134
3.8.4	Proof of Theorem 3.3	136
3.8.5	Proof of Proposition 3.4	136
3.8.6	Proof of Proposition 3.5	137
3.9	Proofs of Section 3.3	138
3.9.1	Proof of Theorem 3.6	140
3.9.2	Proof of Theorem 3.7	141
3.10	Proofs for Section 3.4	149
3.10.1	Proof of Lemma 3.8	149
3.10.2	Proof of Lemma 3.10	150
3.10.3	Proof of Proposition 3.11	151
3.10.4	Proof of Proposition 3.12	151
3.10.5	Proof of Theorem 3.13	152
3.11	Background on Markov chain theory and Riemannian geometry	156
3.11.1	Markov chain notions	157
3.11.2	Useful results from Riemannian geometry	157
4	Optimal Scaling Results for a Wide Class of Proximal MALA Algorithms	167
4.1	Introduction	169
4.2	Proximal MALA Algorithms	171
4.3	Optimal scaling of Proximal MALA	174
4.3.1	Differentiable targets	175
4.3.2	Laplace target	178
4.4	Practical Implications and Numerical Simulations	181
4.4.1	Numerical Experiments	182
4.5	Discussion	184

4.6	Proof of the Result for the Laplace distribution	185
4.6.1	Proof of Theorem 4.4	186
4.6.2	Proof of Proposition 4.5	188
4.6.3	Proof of Proposition 4.6	195
4.6.4	Proof of Theorem 4.7	198
4.7	Proof of Theorem 4.1	206
4.7.1	Auxiliary Results for the Proof of Case (a)	208
4.7.2	Auxiliary Results for the Proof of Case (b)	216
4.7.3	Auxiliary Results for the Proof of Case (c)	219
4.7.4	Proof of Theorem 4.1	222
4.8	Numerical Experiments	223
4.8.1	Differentiable Targets	223
4.8.2	Laplace Target	223
4.8.3	Mix of a Laplace and differentiable target	228
4.9	Taylor Expansions for the Results on Differentiable Targets	231
4.9.1	Coefficients of the Taylor Expansion	231
4.9.2	Taylor Expansions of the Log-acceptance Ratio	237
4.9.3	Derivatives of the Proximity Map for Differentiable Targets	242
4.10	Moments and Integrals for the Laplace Distribution	244
4.10.1	Moments of Acceptance Ratio for the Laplace Distribution	244
4.10.2	Bound on Second Moment of Acceptance Ratio for the Laplace Dis- tribution	250
4.10.3	Additional Integrals for the Laplace Distribution	254
4.10.4	Integrals for Moment Computations	255
A Introduction (en français)		263
A.1	Introduction à l’approximation stochastique	263
A.1.1	Motivation	263
A.1.2	Hypothèses sur le bruit	267
A.1.3	Techniques de preuve pour la convergence	267
A.1.4	Hypothèses pour la descente de gradient stochastique	268
A.1.5	Schéma projeté	268
A.2	Approximation stochastique sur une variété riemannienne	269
A.2.1	Schéma géodésique et schéma par rétraction	270
A.2.2	Premier résultat asymptotique de convergence	271
A.2.3	Analyse non-asymptotique	273
A.2.4	Méthodes par quantité de mouvement et la méthode de Nesterov d’accélération du gradient dans le cas déterministe	274
A.3	Introduction aux méthodes de Monte Carlo par chaînes de Markov	276
A.3.1	Le problème de simuler une loi cible	276
A.3.2	Algorithmes de Monte Carlo et leurs versions markoviennes	277
A.3.3	L’algorithme de Metropolis-Hastings pour simuler une loi cible	278

A.4	Analyse asymptotique des algorithmes de Metropolis-Hastings : échelonnement optimal	280
A.4.1	Introduction à l'échelonnement optimal : l'algorithme de Metropolis par marche aléatoire symétrique	280
A.4.2	Échelonnement optimal de MALA	283
A.5	Cadre de la Thèse et Contributions	285
A.5.1	Schémas riemanniens	285
A.5.2	Hypothèses sur le bruit	285
A.5.3	Techniques de convergence	286
A.5.4	Résultats des Chapitres 2 et 3	287
A.5.5	Résultats du chapitre 4	290
	Bibliographie	291

CHAPTER 1

Introduction

1.1 Introduction to stochastic approximation

Stochastic approximation (SA) appeared in 1951 following the article [RM51] in *Annals of Mathematical Statistics*. It is a family of iterative methods for finding the roots of a continuous function $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$ called mean field, when it is accessible only with noise. The scheme of these methods is as follows:

$$\theta_{n+1} = \theta_n + \eta_{n+1} H_{\theta_n}(X_{n+1}), \quad (1.1)$$

where $(\theta_n)_{n \in \mathbb{N}}$ is the sequence of iterates, with $\theta_n \in \mathbb{R}^d$ for any $n \in \mathbb{N}$ and which is initialized at θ_0 . The sequence $(\eta_n)_{n \in \mathbb{N}^*}$ is the sequence of step sizes, which is deterministic and chosen by the user under the condition $\eta_n > 0$ for any $n \in \mathbb{N}^*$. Finally, $H_{\theta_n}(X_{n+1})$ is an estimator of the continuous function h at the point θ_n . The noise is modeled by the sequence of random variables $(X_{n+1})_{n \in \mathbb{N}}$ with values in $(\mathcal{X}, \mathcal{X})$.

The scheme (1.1) has received interest for decades, analyzed mainly by the communities of stochastic approximation (see [BMP90; Bor09; BCZ11; KY03; PJ92; Rup88]) and optimization (see [dST+21; Nem+09; Nes04; NV08]). In the last twenty years, it has found a renewed interest within the statistical learning community, as its applications are at the heart of numerous successes (see [Bot10; Bot12; BB08b; BCN18; Cra+06; DHS11; Sha+12; SSS07; Xia09]). Both in optimization and statistical learning, the main application of (1.1) is stochastic gradient descent (SGD), which is obtained when $h = -\text{grad} f$ for a differentiable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$.

We introduce some examples where stochastic approximation is used.

1.1.1 Motivation

Example 1.1 (Estimation of the mean). *For $d \in \mathbb{N}^*$, consider a probability distribution π over \mathbb{R}^d with a finite second order moment. Consider a sequence of i.i.d. observations $(X_n)_{n \in \mathbb{N}^*}$ with distribution π . We can estimate the unknown quantity $\mu = \int_{\mathbb{R}^d} u \, d\pi(u)$ with the stochastic gradient descent algorithm. Consider the quadratic risk \mathcal{L}_2 defined for*

any $\theta \in \mathbb{R}^d$ by

$$\mathcal{L}_2(\theta) = \frac{1}{2} \int_{\mathbb{R}^d} \|\theta - u\|^2 d\pi(u).$$

By differentiating under the integral sign, we have that \mathcal{L}_2 is differentiable, with gradient given by $\text{grad } \mathcal{L}_2(\theta) = \theta - \mu$. Since we do not have access to μ , $\text{grad } \mathcal{L}_2$ is not computable, but we can use the observations $(X_n)_{n \in \mathbb{N}}$ as estimators of μ to obtain a cheap, yet inefficient, estimator of $-\text{grad } \mathcal{L}_2(\theta)$. Consider, for any $n \in \mathbb{N}^*$ and any $\theta \in \mathbb{R}^d$, $H_\theta(X_n) = -(\theta - X_n)$. We summarize its relevant estimation properties :

$$\mathbb{E}[H_\theta(X_n)] = -\text{grad } \mathcal{L}_2(\theta), \quad \text{and} \quad \mathbb{E}[\|H_\theta(X_n) + \text{grad } \mathcal{L}_2(\theta)\|^2] = \sigma^2 < +\infty,$$

where $\sigma^2 = \text{Var}(X_1)$. Consider the Robbins-Monro stochastic gradient descent scheme (1.1) to minimize \mathcal{L}_2 , initialized at $\theta_0 = 0$, with step sizes $\eta_n = 1/n$ for $n \in \mathbb{N}^*$:

$$\theta_{n+1} = \theta_n + \eta_{n+1} H_{\theta_n}(X_n) = \theta_n - \frac{1}{n+1} (\theta_n - X_n) = \frac{n}{n+1} \theta_n + \frac{1}{n+1} X_n.$$

A simple induction allows us to see that for any $n \in \mathbb{N}^*$, $\theta_n = n^{-1} \sum_{i=1}^n X_i$. By the strong law of large numbers, $\theta_n \xrightarrow[n \rightarrow \infty]{a.s.} \mu$.

Let us now look at two situations in supervised learning where stochastic gradient descent is used.

Supervised learning We have a set of n i.i.d. observations $(x_i, y_i) \in \mathbf{X} \times \mathbf{Y}$, for $i = 1, \dots, n$, and having a new observation x , we would like to give a prediction of an unseen y using a function $\theta^\top \Phi(x)$ which is linear in the features $\Phi(x) \in \mathbb{R}^d$. We suppose that the feature function $\Phi : \mathbf{X} \rightarrow \mathbb{R}^d$ is known and we want to learn the parameters $\theta \in \mathbb{R}^d$ such that $\theta^\top \Phi(x)$ resembles as closely as possible y . To evaluate the performance of the current parameter θ , we consider a *loss function* $\ell : \mathbf{Y} \times \mathbb{R} \rightarrow \mathbb{R}_+$.

One way to answer this problem is to consider the *empirical risk*.

Example 1.2 (Empirical risk minimization). *One approach is to optimize the empirical risk, that is,*

$$\min_{\theta \in \mathbb{R}^d} \mathcal{R}_{emp}(\theta), \quad \text{where} \quad \mathcal{R}_{emp} : \theta \mapsto \frac{1}{n} \sum_{i=1}^n \ell(y_i, \theta^\top \Phi(x_i)). \quad (1.2)$$

Computing analytically the minimizer of (1.2) is costly if not impossible in general (consider that the amount of observations n could be quite large), therefore, we use an iterative algorithm to approximate the minimizer $\theta^* \in \mathbb{R}^d$.

At iteration t , we produce a random variable i_{t+1} , uniform in $\{1, \dots, n\}$ and consider g_{t+1} , the gradient of $\theta \mapsto \ell(y_{i_{t+1}}, \theta^\top \Phi(x_{i_{t+1}}))$. Then, we produce the following scheme, initialized at $\theta_0 \in \mathbb{R}^d$,

$$\theta_{t+1} = \theta_t - \eta_{n+1} g_{t+1}(\theta_t),$$

1.1. Introduction to stochastic approximation

where $(\eta_n)_{n \in \mathbb{N}^*}$ is the sequence of positive step sizes. In this approach, the randomness comes from the choice of the indices, and since $(i_t)_{t \in \mathbb{N}^*}$ is sampled with replacement, we will use multiple times the same data point (x_i, y_i) . As $t \rightarrow \infty$, the sequence $(\theta_t)_{t \in \mathbb{N}}$ converges to a minimizer θ^* of the empirical risk (1.2).

Another risk function that is relevant to consider is the *expected risk*.

Example 1.3 (Expected risk minimization). We now consider is the minimization of the expected risk, that is,

$$\min_{\theta \in \mathbb{R}^d} \mathcal{R}_{exp}(\theta), \quad \text{where } \mathcal{R}_{exp} : \theta \mapsto \mathbb{E} \left[\ell \left(y, \theta^\top \Phi(x) \right) \right],$$

where (x, y) is a random variable with same law as the observations $(x_i, y_i)_{i \in \{1, \dots, n\}}$. This risk function encompasses all possible values for x , not just the ones that have been observed. It is therefore a better indicator of how good our prediction is. However, the law of the data points (x, y) is not known in practice, therefore the expected risk is not computable. We decide to iteratively optimize an approximation of \mathcal{R}_{exp} .

Just like in Example 1.1, the data points allow us to get simple, yet separately inefficient estimators of $\text{grad } \mathcal{R}_{exp}$. This means that if you denote g_t the gradient of $\theta \mapsto \ell(y_t, \theta^\top \Phi(x_t))$, then, by differentiating under the integral sign, $g_t(\theta)$ is an unbiased estimator of $\text{grad } \mathcal{R}_{exp}(\theta)$. This gives us the following scheme, initialized at $\theta_0 \in \mathbb{R}^d$,

$$\theta_{t+1} = \theta_t - \eta_{n+1} g_{t+1}(\theta_t), \tag{1.3}$$

where t takes now values in $\{0, \dots, n-1\}$ and $(\eta_{n+1})_{n \in \mathbb{N}}$ is the sequence of positive step sizes. In this approach, the randomness comes from the observations themselves. To preserve the independence of the estimators $(g_t)_{t \in \{1, \dots, n\}}$, we can only perform a single pass on the data. If the number of observations n goes to infinity, then the scheme (1.3) will converge to θ^* a minimizer of the expected risk \mathcal{R}_{exp} .

Note that Example 1.2 is well-suited for learning from static data sets $(x_i, y_i)_{i \in \{1, \dots, n\}}$ – also called *batch learning* – whereas Example 1.3 is better suited when data comes as a stream, meaning the user only receives one observation (x_t, y_t) at each step – also known as *online learning*.

The next example showcases the transformation of a batch learning algorithm into an online version using stochastic approximation.

The Expectation-Maximisation (EM) algorithm Let us first introduce the EM algorithm – see [DLR77] – which we present for the estimation of parameters in a latent variable model. This model encompasses missing data, censored observations and noisily observed data.

Consider that there exists a non-observable or latent variable X , that is distributed according to the density $f(x; \theta)$, where $\{f(x; \theta) : \theta \in \Theta\}$ is a parametric family of probability density functions, indexed by the parameter $\theta \in \Theta \subset \mathbb{R}^d$. Consider also an observation

$Y \in \mathcal{Y}$, which is a deterministic function of X . Denote by $g(y; \theta)$ the observed likelihood of the model.

Given n i.i.d. observations $Y = (Y_1, \dots, Y_n)$, the EM algorithm iteratively maximises the log-likelihood $\log g(Y; \theta) = \sum_{i=1}^n \log g(Y_i; \theta)$, despite the complicated formulation of g , resulting from the latent data model. Each iteration k consists of two steps. The E-step evaluates the conditional expectation

$$\theta \mapsto Q_{\theta_k}(Y; \theta) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\theta_k} [\log f(X_i; \theta) | Y_i] ,$$

where θ_k is the current estimate of θ and $\mathbb{E}_{\theta}[\cdot | Y]$ is the conditional expectation under the model parameterized by θ . The M-step consists of finding the value θ that maximises $Q_{\theta_k}(Y; \theta)$, which is taken as the new estimate θ_{k+1} . Interestingly, increasing the quantity $Q_{\theta_k}(Y; \theta)$ results in an increase of the log-likelihood $\log g(Y; \theta)$ – see [DLR77].

This algorithm uses the all data in batch in every iteration. There have been multiple online versions of the EM algorithm, we present the one of [CM09].

Example 1.4 (Online EM algorithm). *Within the context of the EM algorithm, the online E-step becomes: at iteration k , for $\theta \in \Theta$, evaluate the function*

$$\hat{Q}_{k+1}(\theta) = \hat{Q}_k(\theta) + \eta_{k+1} \left(\mathbb{E}_{\theta_k} [\log f(X_{k+1}; \theta) | Y_{k+1}] - \hat{Q}_k(\theta) \right) ,$$

where $(\eta_n)_{n \in \mathbb{N}^*}$ is a sequence of positive step-sizes. The M-step remains taking θ_{n+1} as the value in Θ which maximizes \hat{Q}_{n+1} .

[CM09] prove that this algorithm converges to a stationary point of the Kullback-Leibler distance to the solution (see Theorem 1 of [CM09]). It also has the nice property of decreasing the Kullback-Leibler distance to the solution at each iteration, provided the step-size is sufficiently small (see Proposition 2 of [CM09]).

Having presented a broad view of what stochastic approximation is capable of solving, we now look at the assumptions and proof techniques of this method.

1.1.2 Noise Assumptions

The estimator $H_{\theta_n}(X_{n+1})$ is rewritten as follows,

$$H_{\theta_n}(X_{n+1}) = h(\theta_n) + e_{\theta_n}(X_{n+1}) ,$$

where $e_{\theta_n}(X_{n+1})$ is the estimation error. We define the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$ as the natural filtration of $(\theta_n)_{n \in \mathbb{N}}$.

The classical assumption to make about the noise (see assumption A4 of [Kar+19], assumption A1 of [GL13] or assumption H1 of [MB11]), is to bound the first two moments of the noise, as done below.

A 1 (Martingale difference setting (Euclidean)). *There exists $\sigma^2 > 0$ such that for any $n \in \mathbb{N}$,*

$$\mathbb{E} [e_{\theta_n}(X_{n+1}) | \mathcal{F}_n] = 0 , \text{ and } \mathbb{E} [\|e_{\theta_n}(X_{n+1})\|^2 | \mathcal{F}_n] \leq \sigma^2 .$$

1.1. Introduction to stochastic approximation

This assumption on the noise means that the estimator $H_\theta(X)$ is unbiased and of uniformly bounded variance. This case is called *martingale difference* because the sequence of random variables $(\sum_{k=0}^n e_{\theta_k}(X_{k+1}))_{n \in \mathbb{N}}$ defines a martingale with respect to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$. A generalization of this assumption is to replace the inequality with $\mathbb{E}[\|e_{\theta_n}(X_{n+1})\|^2 | \mathcal{F}_n] \leq \sigma_0^2 + \sigma_1^2 \|h(\theta_n)\|^2$, to say that the bound on the variance can depend on the current point θ_n .

The other kind of noise that has been of interest to the community – see Case 2 in [Kar+19] and [BCN18; GL13] – is the Markovian setting, where the sequence of noise vectors $(e_{\theta_n}(X_{n+1}))_{n \in \mathbb{N}}$ is a state-dependent Markov noise.

A2 (Markovian setting (Euclidean)). *There exists a Markov kernel P on $(\Theta \times \mathbf{X}) \times \mathcal{X}$ such that for any $n \in \mathbb{N}$ and any bounded measurable function $\varphi : \mathbf{X} \rightarrow \mathbb{R}$, $\mathbb{E}[\varphi(X_{n+1}) | \mathcal{F}_n] = \int_{\mathbf{X}} \varphi(y) P_{\theta_n}(X_n, dy)$. Moreover, for any $\theta \in \Theta$, P_θ admits a unique invariant distribution π_θ such that $h(\theta) = \int_{\mathbf{X}} H_\theta(y) d\pi_\theta(y)$.*

This setting is relevant in reinforcement learning topics as Q-learning (see [JJS93]) and policy gradient (see [BB01]).

1.1.3 Convergence techniques

A fruitful approach to analyze scheme (1.1) is to find a non-negative function $V : \mathbb{R}^d \rightarrow \mathbb{R}_+$ that decreases with each iteration. This energy-like function V is called a Lyapunov function, much like in [Lya92]. The terminology comes from Lyapunov’s second method in [KB60], who use this technique to prove the global asymptotic stability of the solutions to the scheme (1.1). They do it by showing that all the trajectories of the solutions to the limiting ordinary differential equation (ODE) $\dot{\theta} = h(\theta)$ go to zero. This technique, coined the ODE approach, has been used with quite a lot of success as in [BFH13; Bor09; BM00; Kar+19]. The caveat is that it requires the user to design their own Lyapunov function, since there is no general way to find them. Most of the analysis has been of asymptotic nature (see [Fab68; Pf86; Sac58]), but non-asymptotic literature has been prolific more recently (see [BM11; Kar+19; Pf90; RRT17]), as non-asymptotic results have palpable implications on most common applications.

1.1.4 SGD assumptions

For the SGD case, Lyapunov conditions are met by different convexity assumptions. Two cases are important to showcase, the strongly convex case and the convex case. When f is strongly convex, under mild assumptions and having for any $n \in \mathbb{N}^*$, $\eta_n = c/n$ for some $c > 0$, we have the following non-asymptotic rate of convergence

$$\mathbb{E} \left[f(\bar{\theta}_n) - f(\theta^*) \right] \leq C/n ,$$

where $C > 0$ is a constant, $\bar{\theta}_n = (1/n) \sum_{k=0}^{n-1} \theta_k$ and θ^* is the unique solution (see [BM11]). This corresponds to the best rate possible for this class of problems (see [Aga+09]).

f is only convex, under similar mild assumptions and taking the same choice for the step size, we obtain

$$\mathbb{E} \left[f(\bar{\theta}_n) - f(\theta^*) \right] \leq C/\sqrt{n}.$$

This bound is also optimal for this class of problems (see [Aga+09]).

1.1.5 Projected scheme

There are multiple reasons to want the iterations of (1.1) to remain on a smaller set $\Theta \subset \mathbb{R}^d$: it is used to accelerate the convergence as in projected SVM [LSB12; SSS07], it can be used to prove the stability of the algorithm – for which it is commonly needed that the iterates remain in some compact set – as in [MMM18] and sometimes the projection is there to ensure that $(\theta_n)_{n \in \mathbb{N}}$ satisfies some constraints S as in [Pfl86]. For any of these cases, we would rewrite (1.1) as follows

$$\theta_{n+1} = \text{proj}_{\Theta} (\theta_n + \eta_{n+1} H_{\theta_n}(X_{n+1})) , \quad (1.4)$$

where proj_{Θ} is the orthogonal projection onto some convex and compact subset $\Theta \subset \mathbb{R}^d$.

Let us look at the following example to motivate further the study of the projected scheme.

Example 1.5 (Estimation of eigenvectors). *We observe a sequence of $d \times d$ independent random matrices $(A_k)_{k \in \mathbb{N}}$, with common expectation $\mathbb{E}[A_k] = A$, a symmetric matrix. We want to approximate an eigenvector of the mean matrix A , using the noisy observations. Consider the scheme found in [OK85],*

$$\theta_n = \frac{\theta_{n-1} + \eta_n A_n \theta_{n-1}}{\|\theta_{n-1} + \eta_n A_n \theta_{n-1}\|} , \quad (1.5)$$

where $\|\cdot\|$ is the Euclidean norm. Assume that $\sum_{i \in \mathbb{N}^*} \eta_i = \infty$, $\sum_{i \in \mathbb{N}^*} \eta_i^2 < \infty$ and **A1** holds – meaning $\sup_{n \in \mathbb{N}} \mathbb{E}[\|A_n\|^2] < \infty$ – then $(\theta_n)_{n \in \mathbb{N}}$ converges to some eigenspace of A (see Theorem 3 in [Del00]). If A has no multiple eigenvalue, then the sequence converges to an eigenvector of A .

It is important to remark that scheme (1.5) can be decomposed in two steps: first $\tilde{\theta}_n = \theta_{n-1} + \eta_n A_n \theta_{n-1}$ and second $\theta_n = \tilde{\theta}_n / \|\tilde{\theta}_n\| = \text{proj}_{\mathbb{S}^{d-1}}(\tilde{\theta}_n)$, where \mathbb{S}^{d-1} is the unit sphere (of dimension $d - 1$ embedded in \mathbb{R}^d) for the Euclidean norm $\|\cdot\|$. The first step corresponds to a scheme of the form (1.1). However, performing only this step would lead to a divergent sequence in many cases because $\|\theta_n\| \rightarrow \infty$ – e.g. in the case of a positive eigenvalue or a poorly chosen sequence $(\eta_n)_{n \in \mathbb{N}^*}$. Since the norm of the vector is not important when looking for an eigenvector, the second step removes an irrelevant dimension of the ambient space by projecting its iterates on the sphere.

In this case, the projection step solves the indetermination of $(\|\theta_n\|)_{n \in \mathbb{N}}$. However, since $\theta_n - \theta_{n-1}$ is no longer proportional to $A_n \theta_{n-1}$, by only looking at $A_n \theta_{n-1}$, we do not know if the next step is going to be large or small. We believe that a more satisfying answer to this indetermination would be to create a scheme similar to (1.1) in a constrained space

Θ , such that the distance between two consecutive iterates corresponds to a norm of the mean field. This new scheme must also address the other issues that led us to consider the projected scheme.

To do this, we restrict the parameter space \mathbb{R}^d to a differentiable manifold Θ , which we endow with a Riemannian metric g .

1.2 Stochastic approximation on Riemannian manifolds

For a Riemannian manifold, we adapt scheme (1.1) into one that remains inside of Θ intrinsically, that is without an ambient space \mathbb{R}^d and a projection step:

$$\theta_{n+1} = \text{Exp}_{\theta_n}(\eta_{n+1}H_{\theta_n}(X_{n+1})) , \quad (1.6)$$

where $\text{Exp}_{\theta} : T_{\theta} \rightarrow \Theta$ is the Riemannian exponential, which corresponds to the flow of a “straight line” in a Riemannian manifold, meaning $\text{Exp}_{\theta}(v) \approx \theta + v$. One might ask if (1.6) is really fundamentally different to the projected scheme (1.4). Intuitively, the path from θ_n to θ_{n+1} is not the projection of a straight line as in (1.4), but rather the “straightest path” that remains on Θ . This key difference causes that under simple conditions¹, the Riemannian distance between θ_{n+1} and θ_n is exactly the Riemannian norm of the mean field $H_{\theta_n}(X_{n+1})$ times the step-size η_{n+1} .

1.2.1 Geodesic versus retraction schemes

The geodesic scheme (1.6) has many advantages compared to the projective scheme (1.4), like the intrinsic belonging of the sequence $(\theta_n)_{n \in \mathbb{N}}$ to the manifold Θ and the very nice property that the Riemannian distance between two consecutive iterations $\rho_{\Theta}(\theta_n, \theta_{n+1})$ is equal to the step size η_{n+1} times the Riemannian norm of the noisy mean field $\|H_{\theta_n}(X_{n+1})\|_{\theta_n}$, both of which we have already mentioned above. However, there are some costs that come with these perks, which come from the computation of the exponential map, or equivalently, the computation of a geodesic $\gamma : [0, 1] \rightarrow \Theta$ with initial point $\gamma(0) = \theta \in \Theta$ and initial speed $\dot{\gamma}(0) = u \in T_{\theta}\Theta$. In the best of cases, the geodesic has a closed form that is easy to compute, like in the case below.

Example 1.6 (Geodesics on the sphere \mathbb{S}^d). *Consider the unit sphere $\mathbb{S}^d = \{\mathbf{x} \in \mathbb{R}^{d+1} : \|\mathbf{x}\| = 1\}$, where $\|\cdot\|$ stands for the standard Euclidean norm on \mathbb{R}^{d+1} . The standard Riemannian manifold structure of the sphere is the following. For any $\theta \in \mathbb{S}^d$, the tangent space is given by $T_{\theta}\mathbb{S}^d = \{\mathbf{u} \in \mathbb{R}^{d+1} : \mathbf{u}^{\top}\theta = 0\}$ – see [AMS08, Example 3.5.1]. The canonical metric on the sphere g is defined as the restriction of the Euclidean scalar product from \mathbb{R}^{d+1} to the tangent space $T_{\theta}\mathbb{S}^d$. The formulation of the exponential map is found in [Lee19, Proposition 5.27],*

$$\text{Exp}_{\theta}(\mathbf{u}) = \cos(\|\mathbf{u}\|)\theta + \sin(\|\mathbf{u}\|)(\mathbf{u}/\|\mathbf{u}\|) .$$

¹this holds for small enough vectors, when they remain inside the tangent injectivity domain $\text{TD}(\theta)$ (see Section 2.11.6)

In other cases, the geodesic has a closed form but is more expensive to compute numerically, like in this example.

Example 1.7 (Geodesics on the Grassmann manifold $\text{Gr}_r(\mathbb{R}^d)$). For $r \in \{1, \dots, d\}$, let $\text{Gr}_r(\mathbb{R}^d)$ be the set of r -dimensional subspaces over \mathbb{R}^d . We can represent $\text{Gr}_r(\mathbb{R}^d)$ as the quotient of the non-compact Stiefel manifold by the action of the linear group $\text{GL}_r(\mathbb{R})$. This representation corresponds to saying that a r -dimensional subspace E of \mathbb{R}^d corresponds to any r linearly independent vectors $\mathbf{v}_1, \dots, \mathbf{v}_r$ in \mathbb{R}^d such that $\text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_r) = E$. We then compute the quotient by identifying the families of vectors that have the same span. See Example 2.9 and Section 2.7.2 for a more detailed presentation of both the Grassmann and Stiefel manifolds, their tangent space and standard metric. For \mathbf{B} in the Stiefel manifold, we denote by $[\mathbf{B}]$ the corresponding class in $\text{Gr}_r(\mathbb{R}^d)$. For any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ and $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, one formulation of the Riemannian exponential map, that is found in [AMS04, Theorem 3.6], is given by

$$\text{Exp}_{[\mathbf{B}]}(\mathbf{D}) = \left[\left(\mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \quad \mathbf{U} \right) \begin{pmatrix} \cos(\boldsymbol{\Sigma}) \\ \sin(\boldsymbol{\Sigma}) \end{pmatrix} \right],$$

where \cos and \sin are given by their power series definitions and $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top$ is the compact singular value decomposition (denoted SVD in the future, see [HP14, Theorem 1.46]), with \mathbf{U} a $d \times r$ matrix with orthogonal columns, meaning $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$, \mathbf{V} a $r \times r$ orthogonal matrix, and $\boldsymbol{\Sigma}$ a $r \times r$ diagonal matrix with nonnegative elements.

In this case, on top of cheap computations as sums, multiplications and trigonometric functions, the computation of a geodesic requires a compact SVD which requires $\mathcal{O}(dr^2)$ computations (see [HGI07]). Such computations often prove too time-expensive for large dimension problems, because this cost is required for every iteration of a scheme such as (1.6). The last case is when there is no closed form for the solutions of the geodesic equation – see Section 2.11.3. In this case we can approximate the solutions numerically, but the scheme is no longer geodesic.

To deal with the problems present in the last two cases presented above, the community was very quickly interested in so-called retraction schemes. The idea is to replace the exponential map with a suitable approximation – which is easier to compute in practice. We call such approximations of the Riemannian exponential map *retractions*.

Definition 1.8 (Retraction). A continuously differentiable map $\text{Ret} : \mathbb{T}\Theta \rightarrow \Theta$, is called a retraction when for any $\theta \in \Theta$, $\text{Ret}_\theta(0_\theta) = \theta$ and $(\text{DRet})_\theta(0_\theta) = \text{Id}$, where 0_θ is the zero element of $\mathbb{T}_\theta\Theta$.

The resulting scheme is formulated as follows. Starting at $\theta_0 \in \Theta$, define recursively for any $n \in \mathbb{N}^*$,

$$\theta_{n+1} = \text{Ret}_{\theta_n}(\eta_{n+1} H_{\theta_n}(X_{n+1})) . \quad (1.7)$$

For examples of retractions, see Examples 2.8 and 2.9. The conditions in the definition imply that Exp_θ and Ret_θ have the same first-order Taylor expansion about 0_θ , which supports the idea that the two schemes may have a similar behavior.

Before looking at the work of the stochastic approximation community, it is necessary to underline the importance of the work of the deterministic optimization and numerical methods communities on Riemannian manifolds. [AMS08; EAS98] in particular have given the possibility to the optimization community to be interested in Riemannian geometry. Compared to other branches of Riemannian geometry, the introduction made in the optimization or numerical methods communities makes a considerable effort of sobriety, by introducing only the notions needed for optimization algorithms. Moreover, the examples treated in these articles have become actual case studies for the stochastic optimization community. To model the randomness into Riemannian manifolds, works like [Pen99] have paved the way for statistical theory on Riemannian manifolds. It is the joint effort of these three communities that has created the ecosystem necessary to make stochastic approximation on Riemannian manifolds a reality.

1.2.2 First asymptotic convergence result

In 2013, [Bon13] started the field of stochastic approximation on Riemannian manifolds by proving the following result.

Theorem 1.9 ([Bon13], Theorem 1). *Consider the scheme (1.6) with the mean field $h = -\text{grad } f$, where $f : \Theta \rightarrow \mathbb{R}$ is a three times differentiable map on a connected Riemannian manifold Θ with injectivity radius uniformly bounded from below (see Section 2.11.6). Assume the sequence of step sizes $(\eta_n)_{n \in \mathbb{N}^*}$ satisfy the standard condition $\sum_{n \in \mathbb{N}^*} \eta_n = \infty$ and $\sum_{n \in \mathbb{N}^*} \eta_n^2 < \infty$. Suppose there exists a compact set K such that $\theta_n \in K$ for all $n \in \mathbb{N}$, and that there exists $C > 0$ such that for any $\theta \in K$ and $x \in X$ we have $\|H_\theta(x)\|_\theta \leq C$. Then $(f(\theta_n))_{n \in \mathbb{N}}$ converges a.s. and $\lim_{n \rightarrow \infty} \text{grad } f(\theta_n) = 0$ a.s.*

For today’s standards, the assumptions on the manifold and the sequence of step sizes are perfectly usual. The existence of the compact set K holds if the manifold is compact, which is true in some of the most important applications, like the sphere, the Stiefel and Grassmann manifolds, see Examples 2.8 and 2.9. However, in general, there is no theoretical guarantee that such a compact exists and in practice, people just check empirically that the simulated trajectories remain in a compact that they choose. This assumption is very problematic, because (1) it restricts initialization inside said compact, meaning that the analysis is only local and never global, (2) the empirical check seems dubious and cannot prove that the trajectories remain inside the compact set *almost surely* and (3) the introduction of this assumption by [Bon13] has incentivized an important part of the community to adopt it, which has led to results on which the diameter of K appears in the non-asymptotic rates of convergence. One of our goals moving forward is to completely get rid of this assumption. This might require a different proof technique, as is seen in future sections. The assumption on the noise is quite restrictive, because it requires that the support of distribution of the noise is bounded. Riemannian equivalents to **A1**, where we only need bounds on the moments of the noise are much more common currently.

This result is in and of itself a stepping stone for the community, but much of the impact of [Bon13] comes from their study of the retraction scheme as well as the plethora

of examples treated. Indeed, in [Bon13, Theorem 2] obtains the same result as Theorem 1.9 for the retraction scheme (1.7), under the additional assumption that for any $\theta \in \Theta$, Ret_θ is twice continuously differentiable.

This second result, combined with illuminating applications in subspace tracking (on the Grassmann manifold), Karcher mean computation (on a hyperbolic space) and estimation on the manifold of symmetric and positive matrices, underline two key strengths of Riemannian analysis. Firstly, Riemannian analysis motivates the appearance of new algorithms by carefully choosing a problem-specific metric on the manifold of parameters and deriving the consequent geodesic or retraction scheme (1.6)-(1.4). These algorithms can be studied within the shared framework of Riemannian stochastic approximation. Secondly, one can also identify existing algorithms that did not start as geometrical ones, but that in fact approximate the geodesic scheme (1.6) closely enough, so that they can be analyzed as a retraction scheme within the Riemannian framework. Let us illustrate this statement with the first example in [Bon13, Section 4.1] and Oja's rule.

Example 1.10 (Online principal component estimation). *Consider \mathbf{A} a $d \times d$ covariance matrix, for which we aim to estimate the first r principal components, meaning the linear span of its principal r eigenvectors. As usual in an online framework, consider a sequence of i.i.d. samples $(X_n)_{n \in \mathbb{N}^*}$ in \mathbb{R}^d , for which we assume that $\mathbb{E}[X_1 X_1^\top] = \mathbf{A}$. To solve this problem, we look first at Oja's rule (see [Oja92]), where θ_0 is a $d \times r$ matrix such that $\theta_0^\top \theta_0 = \mathbf{I}_r$ and the scheme is given by*

$$\theta_{n+1} = \theta_n + \eta_{n+1} \left(\mathbf{I}_d - \theta_n \theta_n^\top \right) X_{n+1} X_{n+1}^\top \theta_n . \quad (1.8)$$

To simplify the geometry of this problem, let us look at the case $r = 1$. The general case is detailed in Section 2.4.1. One way of explaining this scheme is to introduce the cost function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, defined for any $\theta \in \mathbb{R}^d$ by $f(\theta) = -\theta^\top \mathbf{A} \theta / 2$. This cost function, when looking at unitary vectors, is minimized when θ is the leading eigenvector of \mathbf{A} , and its Euclidean gradient is given by $\text{grad}^{(E)} f(\theta) = -\mathbf{A} \theta$. However, Oja's rule does not use scheme (1.1) with this gradient, but rather it precomposes this gradient by the matrix $\mathbf{I}_d - \theta_n \theta_n^\top$. The geometric interpretation is that, since we do not care about the norm of θ_n , we only need to keep track of the parts of $\text{grad}^{(E)} f(\theta_n)$ that are orthogonal to θ_n . This is accomplished by the projection matrix on the orthogonal supplement to $\text{Span}(\theta_n)$, $\mathbf{I}_d - \theta_n \theta_n^\top$.

Now, to interpret (1.8) as a retraction scheme as in (1.7), we set the manifold to be the unit sphere embedded in \mathbb{R}^d , denoted by \mathbb{S}^{d-1} , and to compute the Riemannian gradient of f on the sphere, we use [AMS08, Example 3.6.1] $\text{grad}^{(S)} f(\theta) = (\mathbf{I}_d - \theta \theta^\top) \text{grad}^{(E)} f(\theta) = -(\mathbf{I}_d - \theta \theta^\top) \mathbf{A} \theta$. The simplest retraction is the one we used in Example 1.5, namely $\text{Ret}_\theta^{(S)} : v \in \mathbb{T}_\theta \mathbb{S}^{d-1} \mapsto (\theta + v) / \|\theta + v\| \in \mathbb{S}^{d-1}$, yielding the scheme

$$\theta_{n+1} = \text{Ret}_{\theta_n}^{(S)} \left(\eta_{n+1} \left(\mathbf{I}_d - \theta_n \theta_n^\top \right) X_{n+1} X_{n+1}^\top \theta_n \right) .$$

1.2.3 Non-asymptotic analysis

Since, the community has been looking to prove the properties of SA in Euclidean spaces for the Riemannian framework. As a Riemannian counterpart to [BM11], [ZS16] deal with non-asymptotic analysis for the SGD case for g -convex and g -strongly convex functions.

Definition 1.11 (Geodesically convex function). *A map $f : \Theta \rightarrow \mathbb{R}$ is said to be geodesically convex, or g -convex in short, when for any $\theta_0, \theta_1 \in \Theta$, a geodesic $\gamma : [0, 1] \rightarrow \Theta$ such that $\gamma(0) = \theta_0$ and $\gamma(1) = \theta_1$ and any $t \in [0, 1]$,*

$$f(\theta_1) \geq f(\theta_0) + \left\langle \text{grad } f(\theta_0), \text{Exp}_{\theta_0}^{-1}(\theta_1) \right\rangle_{\theta_0} .$$

Definition 1.12 (λ -strongly geodesically convex function). *Consider f a geodesically convex mapping and $\lambda > 0$. f is said to be λ -strongly geodesically convex if for any $\theta_0, \theta_1 \in \Theta$, there exists a geodesic $\gamma : [0, 1] \rightarrow \Theta$ such that $\gamma(0) = \theta_0$ and $\gamma(1) = \theta_1$ and any $t \in [0, 1]$,*

$$f(\theta_1) \geq f(\theta_0) + \left\langle \text{grad } f(\theta_0), \text{Exp}_{\theta_0}^{-1}(\theta_1) \right\rangle_{\theta_0} + \frac{\lambda}{2} \rho_{\Theta}^2(\theta_0, \theta_1) ,$$

where $\rho_{\Theta} : \Theta \times \Theta \rightarrow \mathbb{R}_+$ is the Riemannian distance on Θ (see Section 2.11.1 and (2.150)).

The results are equivalent bounds to the ones in the Euclidean setting, under the assumption that the iterates $(\theta_n)_{n \in \mathbb{N}}$ remain in a compact subset of Θ . Namely, for a strongly convex function f over Θ , [ZS16, Theorem 12] states that there exists $C > 0$ such that for any $n \in \mathbb{N}$,

$$\mathbb{E} \left[f(\bar{\theta}_n) - f(\theta^*) \right] \leq C/n ,$$

where $\bar{\theta}_0 = \theta_0$, $\bar{\theta}_{s+1} = \text{Exp}_{\bar{\theta}_n} \{ (2(n+1))^{-1} \text{Exp}_{\bar{\theta}_n}^{-1}(x_{s+1}) \}$ and θ^* is the unique solution. For the convex case, [ZS16, Theorem 10] gives under the same setting that there exists $C > 0$ such that for any $n \in \mathbb{N}$,

$$\mathbb{E} \left[f(\bar{\theta}_n) - f(\theta^*) \right] \leq C/\sqrt{n} .$$

1.2.4 Momentum and Nesterov's Accelerated Gradient in the deterministic case

For convex functions, there has been recent efforts to implement different kinds of stochastic gradient descent with momentum. Although the community has only analyzed the deterministic gradient case, meaning $H_{\theta}(X) = -\text{grad } f(\theta)$, their works are related and could be the avenue for new stochastic schemes. On vector spaces, the idea of momentum is a scheme like the following

$$v_{n+1} = \beta v_n + \eta H_{\theta_n}(X_{n+1}) , \quad \text{and} \quad \theta_{n+1} = \theta_n + v_{n+1} , \quad (1.9)$$

where $(v_n)_{n \in \mathbb{N}}$ is the sequence of momentum or velocity terms, $\beta \in (0, 1)$ is the decaying hyper-parameter usually taken equal to 0.9 and $\theta \mapsto H_{\theta}(X)$ is an estimator of the mean-field map $-\text{grad } f$. The intuitive idea is that, by keeping some of the speed of previous

iterations, it will smooth the trajectory, meaning that it will limit oscillation when close to the solution. In practice, momentum techniques allow to escape saddle points, and have been proven to obtain bounds of the form $f(\theta_n) - f(\theta^*) = \mathcal{O}(1/n^2)$ [Nes83] for convex functions. In particular, Nesterov's acceleration (see [Nes04; Nes83]) has had a lot of success on vector spaces. The scheme is the following,

$$\begin{aligned}\tilde{\theta}_n &= \theta_n + \beta v_n, \\ v_{n+1} &= \beta v_n + \eta H_{\tilde{\theta}_n}(X_{n+1}), \\ \theta_{n+1} &= \theta_n + v_{n+1}.\end{aligned}$$

The key difference with respect to scheme (1.9) being that the gradient is computed *after* the momentum step. An intuitive way of understanding the possible benefit of this scheme would be say the gradient comes afterwards the momentum step *correcting* its erroneous movements, instead of accentuating them.

In the Riemannian case, momentum methods require more computations than one would expect coming from the vector space point of view. In the deterministic setting, [Ali+21] proposes the following scheme, initialized for $\theta_0 = v_0 \in \Theta$,

$$\begin{aligned}\tilde{\theta}_n &= \text{Exp}_{v_n} \left(\beta_n \text{Exp}_{v_n}^{-1}(\theta_n) \right), \\ \theta_{n+1} &= \text{Exp}_{\tilde{\theta}_n} \left(-\eta \text{grad} f(\tilde{\theta}_n) \right), \\ v_{n+1} &= \text{Exp}_{v_n} \left(-a_{n+1} \text{T}_{\beta_n 0}^{\gamma_n} \text{grad} f(\tilde{\theta}_n) \right),\end{aligned}\tag{1.10}$$

where we set $\gamma_n : [0, 1] \rightarrow \Theta$ the geodesic starting at $\gamma_n(0) = v_n$ and with end point $\gamma_n(1) = \theta_n$, β_n is computed in $[0, 1]$ to minimize $f \circ \gamma_n$ and $(a_n)_{n \in \mathbb{N}}$ is given by $a_0 = 0$ and a_{n+1} is the positive solution to $c_0 a_{n+1}^2 - c_1 a_{n+1} - c_1 \sum_{k=0}^n a_k = 0$, with $c_0, c_1 > 0$ being some positive constants chosen as functions of the curvature and the smoothness of $\text{grad} f$ and lastly $\text{T}_{t_1 t_0}^{\gamma} : \text{T}_{\gamma(t_1)} \Theta \rightarrow \text{T}_{\gamma(t_0)} \Theta$ is the parallel transport map from $\text{T}_{\gamma(t_1)} \Theta$ to $\text{T}_{\gamma(t_0)} \Theta$ along γ (see Section 2.11.4 for more details on parallel transport). Some of the complications of the Riemannian scheme come from the fact that vectors can only be tangent to a specific point. This means that if $\text{grad} f(\tilde{\theta}_n)$ is used both to update θ_n and v_n , we will need to transport it from $\text{T}_{\tilde{\theta}_n} \Theta$ to $\text{T}_{v_n} \Theta$, thus the need of the parallel transport map. Their analysis applies to both convex functions and weak quasi-convex functions (see **F3** in Section 3.4 for a definition). They also assume that the iterates of the scheme stay inside some geodesically convex subset $K \subset \Theta$ with bounded diameter, and that the exponential is a global diffeomorphism on K . They obtain bounds of the form $f(\theta_n) - f(\theta^*) \leq C_0 \max(C_1/k^2, C_2/k) + \varepsilon$, which in the worst case scenario perform as well as scheme (1.1), but sometimes it outperforms it (see [Ali+20, Theorem 4]). However, as promising as their results are, the analysis of such a scheme has not yet been done in the stochastic case.

Also in the deterministic case, [ZS18] and [AS20] study a scheme of the general form of (1.10) but under strong-convexity assumptions. The other assumptions follow the ones of [Ali+20], except that [AS20] manage to remove the assumption of the existence of the subset K – meaning they perform a global analysis – by introducing a Lyapunov function.

The conclusions of [ZS18] are that acceleration (meaning an $\mathcal{O}(1/n^2)$ speed) is obtained, when the scheme is initialized close to the solution, which may take a while in badly conditioned problems. In [AS20], authors obtain that their algorithm is strictly faster than gradient descent using Lyapunov functions and a metric distortion inequality.

1.3 Introduction to Markov chain Monte Carlo methods

1.3.1 The problem of sampling from a target distribution

The second topic in this thesis is motivated by the problem of sampling from a probability distribution. Among all the possible settings where sampling is useful, let us introduce the case of Bayesian statistics. Basically, sampling is used in Bayesian statistics to explore the posterior distribution.

The main assumption in Bayesian statistics is that there are two probabilistic models, one for the observed data – which is usually a parametric model, say depending on the parameter θ – and another one for the parameter θ . The distribution of the parameter θ , which is called the *prior* distribution, encapsulates all the knowledge about said parameter *before* considering the observations at play. The goal is to combine knowledge from the observations and previous information on the parameter. We do this by computing the distribution of the parameter given the observation, which we call the posterior distribution.

More precisely, consider the measurable space (Θ, \mathcal{F}) , where the random variable θ lies, with distribution ν_θ . More often than not, the parameter space Θ is discrete or a subset of \mathbb{R}^d for $d \geq 1$. For that reason, we assume that ν_θ admits a density with respect to a measure ϖ , and we denote its density ν_θ still. Secondly, we observe w , a realization of the random variable W which takes its values in a measurable space (W, \mathcal{W}) . We then specify the family of conditional densities of W given θ ,

$$\mathcal{P}_\Theta = \{w \mapsto \mathcal{L}(w|\theta) \mid \theta \in \Theta\} ,$$

with respect to a reference measure μ on (W, \mathcal{W}) . We call likelihood of the model the function $\theta \mapsto \mathcal{L}(W|\theta)$.

Having described both the distribution of θ and of W given θ , we can access both the joint distribution, with density $(\theta, w) \mapsto \nu_\theta(\theta)\mathcal{L}(w|\theta)$ with respect to $\varpi \otimes \mu$ and the distribution of θ given W , the posterior. Using Bayes theorem ([Sch95, Theorem 1.31]), we have the following expression for the density of the posterior $f_{\theta|X}$, defined for ν_θ -almost any θ and μ -almost any w by

$$f_{\theta|X}(w, \theta) = \frac{\mathcal{L}(w|\theta)\nu_\theta(\theta)}{\int_\Theta \mathcal{L}(w|\tilde{\theta})\nu_\theta(\tilde{\theta})\varpi(d\tilde{\theta})} , \tag{1.11}$$

where for any w ,

$$f_X(w) = \int_\Theta \mathcal{L}(w|\theta)\nu_\theta(\theta)\varpi(d\theta) , \tag{1.12}$$

is the density of W with respect to μ .

1.3.2 Monte Carlo algorithms and their Markovian counterparts

Consider a probability measure π on a measurable space $(\mathbf{E}, \mathcal{E})$, which we call the target distribution. We want to estimate quantity of the form $\int_{\mathbf{E}} f(x)\pi(dx)$, where $f : \mathbf{E} \rightarrow \mathbb{R}$ is an integrable function with respect to π . The idea of Monte Carlo is to use estimators of the form

$$\hat{f}_N = \frac{1}{N+1} \sum_{i=0}^N f(Y_i),$$

where $(Y_i)_{i \in \mathbb{N}}$ is a sequence of random variables. The different Monte Carlo methods propose different settings for the sequence $(Y_i)_{i \in \mathbb{N}}$, to ensure that \hat{f}_N is a good estimator.

Traditional Monte Carlo methods rely on the strong law of large numbers, meaning they consider $(Y_i)_{i \in \mathbb{N}}$ to be a sequence of independent and identically distributed (in short i.i.d.) random variables distributed according to π . In this context, the strong law of large numbers states that \hat{f}_N converges almost surely, as N goes to infinity, towards $\int_{\mathbf{E}} f(x)\pi(dx)$. Moreover, in the case where f^2 is integrable with respect to π , we know that the asymptotic error is of order $\mathcal{O}(1/\sqrt{N})$ thanks to the central limit theorem. The simplicity of the method comes with the price of having to sample i.i.d. random variables with distribution π . Among the vast literature of standard methods, the inverse transform sampling and the accept-reject algorithm are the most popular ones – see [RC10, Section 2.1.2] and [RC10, Section 2.3].

An important inconvenience of these methods is that they suffer from the curse of dimensionality. Another problem is that sometimes we only know the density f up to a multiplicative constant. For example, in Bayesian inference, where π admits a density with respect to the prior distribution (1.11), given by:

$$x \mapsto \frac{\mathcal{L}(\mathbf{w}|x)}{f_X(\mathbf{w})},$$

where \mathbf{w} are the observed data. Indeed, except when the prior and the likelihood are conjugate distributions, the marginal density appearing in (1.12) does not have a closed form. This problem severely limits the applicability of both methods mentioned above.

A solution to some of these problems comes with Markov chain Monte Carlo (or MCMC in short). The idea is to no longer use i.i.d. samples, but rather generate the sequence $(Y_i)_{i \in \mathbb{N}}$ using a Markov chain with invariant distribution π . When the appropriate assumptions are met on π and the Markov chain, an equivalent to the law of large numbers holds – see [MT09, Chapter 17]. In particular [MT09, Theorem 17.1.7] shows that if the Markov kernel P is Harris recurrent and admits π as invariant distribution, then for any $x \in \mathbf{E}$,

$$\lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{i=0}^N f(Y_i) = \int_{\mathbf{E}} f(y)dy, \quad \mathbb{P}_x\text{-almost surely},$$

where \mathbb{P}_x is the induced law by the Markov kernel P with the initial distribution δ_x on $(\mathbf{E}^{\mathbb{N}}, \mathcal{E}^{\otimes \mathbb{N}})$ and $(Y_i)_{i \in \mathbb{N}}$ is the canonical chain on $(\mathbf{E}^{\mathbb{N}}, \mathcal{E}^{\otimes \mathbb{N}}, \mathbb{P}_x)$. In order to implement such methods, we need to build Markov kernels that admit π as invariant distribution. A

generic way to achieve this is using Metropolis-Hastings type algorithms, which we now present.

1.3.3 The Metropolis-Hastings procedure for sampling

The first introduction of Metropolis-Hastings algorithms dates back to [Met+53], which were later improved on by [Has70] and [Tie98]. Consider ν a measure on $(\mathbf{E}, \mathcal{E})$ that dominates π and consider the conditional density $x, y \in \mathbf{E}^2 \mapsto q(x, y)$, meaning that for any $x \in \mathbf{E}$, $y \mapsto q(x, y)$ is a density with respect to ν on \mathbf{E} and that for ν -almost every $y \in \mathbf{E}$, $x \mapsto q(x, y)$ is measurable. Then, given $X_0 \in \mathbf{E}$ with some initial distribution on $(\mathbf{E}, \mathcal{E})$, we define the Markov chain $(X_n)_{n \in \mathbb{N}}$, where for any $n \in \mathbb{N}$,

$$X_{n+1} = Y_{n+1} \mathbb{1}\{\mathbf{U}_{n+1} < \alpha(X_n, Y_{n+1})\} + X_n \mathbb{1}\{\mathbf{U}_{n+1} \geq \alpha(X_n, Y_{n+1})\}, \quad (1.13)$$

where Y_{n+1} is a random variable, which, conditionally to $\mathcal{F}_n = \sigma(X_0, \dots, X_n)$, the σ -algebra generated by the first $n + 1$ terms of the chain, is distributed according to the density $y \mapsto q(X_n, y)$ and $(\mathbf{U}_n)_{n \in \mathbb{N}^*}$ is an i.i.d. sequence of uniform random variables on $[0, 1]$. The function $\alpha : \mathbf{E} \times \mathbf{E} \rightarrow [0, 1]$ is assumed measurable and it is called the acceptance ratio. Denote $P : \mathbf{E} \times \mathcal{E} \rightarrow [0, 1]$ the Markov kernel of the chain $(X_n)_{n \in \mathbb{N}}$. Remark that the kernel P associated with the Markov chain defined in (1.13) is given by, for any $x \in \mathbf{E}$ and any $\mathbf{A} \in \mathcal{E}$,

$$P(x, \mathbf{A}) = \int_{\mathbf{A}} \alpha(x, y) q(x, y) \nu(dy) + \delta_x(\mathbf{A}) \int_{\mathbf{E}} (1 - \alpha(x, y)) q(x, y) \nu(dy).$$

This accept-reject procedure allows sampling of a Markov chain with Markov kernel P using the so-called family of proposal distributions $y \mapsto q(x, y)$. Denote π the density of target distribution π with respect to ν . Recall that we are interested in ensuring that the chain admits π as stationary distribution. However, P is reversible with respect to π if and only if for any bounded and measurable function $g : \mathbf{E}^2 \rightarrow \mathbb{R}$,

$$\int_{\mathbf{E}^2} g(x, y) \pi(x) P(x, dy) \nu(dx) = \int_{\mathbf{E}^2} g(x, y) \pi(y) P(y, dx) \nu(dy).$$

But, by remarking the presence of the Dirac measure in the second part of P , this is also equivalent to: for ν -almost any $x, y \in \mathbf{E}$,

$$\alpha(x, y) \pi(x) q(x, y) = \alpha(y, x) \pi(y) q(y, x).$$

This last condition holds, if the acceptance ratio is given by

$$\alpha(x, y) = \begin{cases} \min\left(1, \frac{\pi(y)q(y,x)}{\pi(x)q(x,y)}\right) & \text{if } \pi(x)q(x, y) > 0 \text{ and } \pi(y)q(y, x) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (1.14)$$

The Metropolis-Hastings algorithm is specially useful when it is easy to sample from the proposal distribution of density $y \mapsto q(x, y)$.

If we look for the simplest choice possible for the proposal, we can take for any $x, y \in \mathbf{E}^2$, $q(x, y) = q(y)$, so that the proposal distribution does not depend on the current state of the chain. This gives the following form for the acceptance ratio α ,

$$\alpha(x, y) = 1 \wedge \frac{\pi(y)q(x)}{\pi(x)q(y)} . \quad (1.15)$$

This version of the Metropolis sampler is called the Metropolis independent sampler, notably because the proposal distribution does not depend on the current point x . Another very classical choice of Metropolis-Hastings proposals relies on \mathbf{E} being an additive group and ν a translation invariant measure. In this setting, we consider the random walk where the increment has a symmetric probability measure on $(\mathbf{E}, \mathcal{E})$, with density \tilde{q} with respect to ν . This means that the sequence of proposals is given by

$$Y_{k+1} = X_k + Z_{k+1} ,$$

where the random variable X_k is the current state of the Markov chain at iteration k and $(Z_i)_{i \geq 1}$ are i.i.d. random variables with density \tilde{q} with respect to ν . Moreover, there is a relationship between the transition density q and the density of the increment \tilde{q} : for any $x, y \in \mathbf{E}$ $q(x, y) = \tilde{q}(x - y)$, which comes from the fact that \tilde{q} is symmetric and thus the acceptance ratio is of the form (1.15). The quintessential example in this setting is the symmetric random walk Metropolis algorithm. It is defined on \mathbb{R}^d with the Lebesgue measure as the dominating measure, and the distribution of the increment is the zero-mean Gaussian with covariance matrix $\sigma^2 \text{Id}$, where $\sigma^2 > 0$.

We finish this part mentioning the Metropolis Adjusted Langevin Algorithm (MALA). Consider here $\mathbf{E} = \mathbb{R}^d$ endowed with its Borel σ -field $\mathcal{B}(\mathbb{R}^d)$. In addition, suppose that π admits a positive smooth density with respect to the Lebesgue measure Leb , still denoted by π and set $U = -\log \pi$. MALA is based on the Langevin diffusion which is a Stochastic Differential Equation (SDE) which defines a Markov semigroup reversible with respect to π . However, sampling exact solutions of this SDE in continuous time is not possible in general. Therefore, we consider instead the corresponding Euler Maruyama discretization, which reads

$$Y_{k+1} = Y_k - \gamma \nabla U(Y_k) + \sqrt{2\gamma} Z_{k+1} , \quad (1.16)$$

where $\gamma > 0$ is the step size of the discretization and $(Z_k)_{k \geq 1}$ is a sequence of i.i.d. d dimensional standard normal random variables. This Markov chain is a discrete approximation of the Langevin diffusion. It is important to remark that $(Y_n)_{n \in \mathbb{N}}$ has a unique invariant distribution π_γ – which depends on the choice of the step size γ – and that under mild conditions on U , it is also ergodic. And even though π_γ is different from π most of the time, it is sometimes a small enough cost to pay for having such a simple scheme. Indeed, this algorithm for sampling was first introduced by [Erm75] and [Par81] for applications in molecular dynamics. It has also been of special interest in the machine learning community thanks to [Gre83], [GM94a] and in the computational statistics community thanks to [Nea93] and [RT96a].

Nonetheless, there are many applications where sampling from a distribution *close* to π is not good enough, since this will cause our estimators $(\hat{f}_N)_{N \in \mathbb{N}}$ to no longer converge

to the required integral $\int f d\pi$ in general. The idea of [RDF78] and [RT96a] is to use the Markov chain $(Y_n)_{n \in \mathbb{N}}$ in (1.16) as the proposal of the Metropolis-Hastings algorithm. As was introduced in [RT96a], we call this procedure Metropolis Adjusted Langevin Algorithm (MALA).

1.4 Asymptotic analysis of Metropolis-Hastings algorithms: optimal scaling

In this thesis, we are interested in analyzing of Metropolis-Hastings algorithms when the dimension goes to infinity. Optimal scaling, pioneered by [RGG97a], gives a good understanding of the convergence properties of MCMC in high dimension.

When using the symmetric random walk Metropolis or MALA in \mathbb{R}^d , we need to choose properly a parameter in the proposal distribution. For the case of the symmetric random walk Metropolis, it is the variance of the normal distribution and for MALA it is the step size of the discretization. The choice of the parameter influences the variance of the proposal and thus the behavior of the Markov chain: small increments explore slowly the support of π but have a high acceptance rate, whereas larger increments move more quickly to all the support of π but they are rejected more often. Optimal scaling defines a framework to formulate the question of the best possible choice for the parameter and answers it with respect to the dimension d . We present here the classical results in this field.

1.4.1 Introduction to optimal scaling: the symmetric random walk Metropolis algorithm

As said above, the study of optimal scaling started with the work of [RGG97a], where they study the symmetric random walk Metropolis algorithm (RWM in short) \mathbb{R}^d . For a positive target density π^d on \mathbb{R}^d , let us recall the expression of the Markov chain created by the RWM algorithm, which is of the form (1.13). For any initial value $X_0^d \in \mathbb{R}^d$ and for any $k \in \mathbb{N}$,

$$X_{k+1}^d = X_k^d + \sigma_d Z_{k+1}^d \mathbb{1}\{\mathbf{U}_{k+1}^d < \alpha_d(X_k^d, X_k^d + \sigma_d Z_{k+1}^d)\},$$

where $\sigma_d > 0$, $(Z_{k+1}^d)_{k \in \mathbb{N}}$ is a sequence of i.i.d. standard d -dimensional normal random variables and $(\mathbf{U}_{k+1}^d)_{k \in \mathbb{N}}$ is an i.i.d. sequence of uniform random variables on $[0, 1]$, with $(Z_{k+1}^d)_{k \in \mathbb{N}}$, $(\mathbf{U}_{k+1}^d)_{k \in \mathbb{N}}$ and X_0 are independent of one another. Moreover, the acceptance rate α_d is given by, for any $x, y \in \mathbb{R}^d$,

$$\alpha_d(x, y) = 1 \wedge \frac{\pi^d(y)}{\pi^d(x)}.$$

The user then can choose the parameter σ_d *i.e.* the variance of the proposal. A natural question that arises is “what is a *good* choice for σ_d ?” Indeed, a large variance is preferred in order to explore quickly the parts of \mathbb{R}^d where π^d takes higher values, but this will cause

the algorithm to reject more often, thus slowing down the exploration. In the case of an initial distribution different than π^d , the choice of σ_d would also impact the time needed so that the distribution of X_k^d is at distance ε with π^d (for example for the total variation distance on distributions). The starting point of optimal scaling is to find the values of the parameter σ_d so that the accept-reject step is not asymptotically trivial, meaning that the mean acceptance ratio under the stationary distribution converges inside $(0, 1)$, *i.e.*

$$\lim_{d \rightarrow +\infty} \mathbb{E}_{\pi^d} \left[\alpha_d \left(X_0^d, X_0^d + \sigma_d Z_1^d \right) \right] \in (0, 1) .$$

Answering this question requires a framework for each $d \in \mathbb{N}$, and the choice of π^d with respect to d influences greatly the behavior of the Markov chain. The first assumption for such a framework, the one in [RGG97a], relies on π^d being a product measure, of the following form. For any $x^d = (x_1^d, \dots, x_d^d) \in \mathbb{R}^d$,

$$\pi^d(x^d) = \prod_{i=1}^d \exp \left(-U(x_i^d) \right) , \quad (1.17)$$

where $U : \mathbb{R} \rightarrow \mathbb{R}$ is a map that defines the one-dimensional density $\pi(x) = \exp(-U(x))$. These functions satisfy the following assumptions,

H 1 U is three times continuously differentiable and U' is Lipschitz.

H 2 $\mathbb{E}[(U'(W))^8]$ and $\mathbb{E}[(U''(W))^4]$ are finite, where W is a real random variable with distribution π .

[RGG97a] propose to take $\sigma_d = \ell d^{-\alpha/2}$ as a general form for σ_d , with $\ell, \alpha \in \mathbb{R}_+^*$. Then, [RGG97a, corollary 1.2] proves that if for any $d \in \mathbb{N}$, X_0^d is distributed according to π^d , then for $\alpha = 1$ we have

$$\lim_{d \rightarrow +\infty} \mathbb{E}_{\pi^d} \left[\alpha_d \left(X_0^d, X_0^d + \ell d^{-1/2} Z_1^d \right) \right] = 2\Phi(-\ell\sqrt{K}/2) ,$$

where Φ is the standard normal cumulative distribution function

$$K = \int_{\mathbb{R}} (U')^2(x) \pi(x) dx . \quad (1.18)$$

Having found the scaling of $\sigma_d = \ell d^{-\alpha/2}$ such that the algorithm has non-degenerate acceptance, we can now investigate into the choice of the constant ℓ . [RGG97a] continues in this direction by studying the following continuous-time process $(V_t^d)_{t \geq 0}$ given by

$$V_t^d = X_{\Gamma_t^d, 1}^d , \quad (1.19)$$

where $(\Gamma_t^d)_{t \geq 0}$ is a Poisson process with rate d . A rough explanation of this choice is that, since the Poisson process jumps td times on average until time t , V_t^d is roughly the first component of the Markov chain X^d at time $\lfloor td \rfloor$. In [RGG97a, Theorem 1.1], the authors

1.4. Optimal scaling

prove that if U satisfies the conditions H1 and H2, then the sequence of processes $(V^d)_{d \in \mathbb{N}^*}$ converges weakly, in the Skorokhod topology, to a process V that satisfies the following: V_0 is distributed according to π and for any $t \geq 0$, V satisfies the overdamped Langevin stochastic differential equation,

$$dV_t = \sqrt{h(\ell)} dB_t^1 - \frac{1}{2} h(\ell) U'(V_t) dt ,$$

where $(B_t^1)_{t \geq 0}$ is a unidimensional Brownian motion, $h(\ell)$ is often called the speed of the diffusion and is given for all $\ell > 0$ by

$$h(\ell) = 2\ell^2 \Phi \left(-\frac{\ell\sqrt{K}}{2} \right) ,$$

and K is defined by (1.18). To understand how this result helps us choose the parameter ℓ , remark that, by defining a rescaling of the limiting process. Consider, for any $t \in \mathbb{R}_+$, $\tilde{V}_t = V_{h(\ell)t}$. Then \tilde{V} is the Langevin diffusion with unit speed:

$$d\tilde{V}_t = dB_t^1 - \frac{1}{2} U'(\tilde{V}_t) dt .$$

We now understand that by increasing the speed h , the Markov chain explores faster the trajectories of \tilde{V} and \mathbb{R}^d , therefore being more efficient for our initial sampling problem. Thus, we want to find ℓ such that h attains its maximum. [RGG97a] even gives a simple criterion to find such ℓ . Denote $a(\ell) = 2\Phi(-\ell\sqrt{K}/2)$ the limit acceptance, then, [RGG97a, Corollary 1.2] proves that one can express $h(\ell)$ as a function of $a(\ell)$ by inverting the expression of $a(\ell)$. Their conclusion is that choosing ℓ_0 such that $a(\ell_0) = 0.234$ (up to four digits) attains the maximal speed for the diffusion V . Another way of interpreting the speed of the diffusion is through the fact that $\lim_{d \rightarrow \infty} (d/2) \mathbb{E}[(X_0^d - X_1^d)^2] = h(\ell)$, where $\mathbb{E}[(X_0^d - X_1^d)^2]$ is sometimes called the expected square jumped distance of the chain X^d .

In the same context, the work of [Dur+17] consider a more general set of assumptions for U , namely the L - p mean differentiability.

H 3 There exists a measurable function $\dot{U} : \mathbb{R} \rightarrow \mathbb{R}$ such that:

- (i) There exist $p > 4$, $C > 0$ and $\beta > 1$ such that for any $\theta \in \mathbb{R}$,

$$\left(\int_{\mathbb{R}} |U(x + \theta) - U(x) - \theta \dot{U}(x)|^p \pi(x) dx \right)^{1/p} \leq C |\theta|^\beta .$$

- (ii) The function \dot{U} satisfies

$$\int_{\mathbb{R}} (\dot{U}(x))^6 \pi(x) dx < \infty .$$

The function \dot{U} is called the weak derivative of U . In this setting, [Dur+17, Theorem 2] proves the convergence of the acceptance rate towards $a(\ell) = 2\Phi(-\ell\sqrt{K}/2)$ when

$\sigma_d = \ell d^{-1/2}$. Then [Dur+17, Theorem 3] proves that under this setting, the random process V satisfies the same overdamped Langevin stochastic differential equation, with notably the same expression for the speed of the diffusion $h(\ell)$. These results are interesting to point out in this thesis because our contributions in the field of MCMC are very much a mirror to these statements, but for the optimal scaling of MALA.

1.4.2 Optimal scaling of MALA

We can ask the same question of “what is the optimal choice for σ_d ?” for MALA. Let us recall the setting of MALA: we have the positive target density π^d on \mathbb{R}^d that can be written as a product of the form (1.17) with $U : \mathbb{R} \rightarrow \mathbb{R}$ is continuously differentiable and, given an initial value $X_0^d \in \mathbb{R}^d$, the Markov chain is given by, for any $k \in \mathbb{N}$,

$$X_{k+1}^d = X_k^d + \tilde{Z}_{k+1}^d \mathbb{1}\{U_{k+1}^d < \alpha_d(X_k^d, X_k^d + \tilde{Z}_{k+1}^d)\},$$

where $\tilde{Z}_{k+1}^d = \sigma_d^2 \nabla \log \pi^d(X_k^d) + \sqrt{2} \sigma_d Z_{k+1}^d$ is the increment in the proposal, $\sigma_d \in \mathbb{R}_+^*$, $(Z_i^d)_{i \geq 1}$ is an i.i.d. sequence of d -dimensional standard normal random variables. Remark- ing that conditionally on X_k^d , \tilde{Z}_{k+1}^d has a normal distribution of mean $\sigma_d^2 \nabla \log \pi^d(X_k^d)$ and variance $2\sigma_d^2$, we can simplify the general formulation of α_d in (1.14). For any $x, y \in \mathbb{R}^d$, α_d is given by,

$$\alpha_d(x, y) = 1 \wedge \frac{\pi^d(y) \exp\left(-\|x - y - \sigma_d^2 \nabla \log \pi^d(y)\|^2 / (4\sigma_d^2)\right)}{\pi^d(x) \exp\left(-\|y - x - \sigma_d^2 \nabla \log \pi^d(x)\|^2 / (4\sigma_d^2)\right)}.$$

The resulting framework is very similar at the one for the symmetric random walk Metropo- lis, where σ_d is the only parameter to tune. After writing $\sigma_d = \ell d^{\alpha/2}$, what is the value of α such that $\mathbb{E}_{\pi^d}[\alpha_d(X_0^d, X_0^d + \tilde{Z}_1^d)]$ converges, when $d \rightarrow +\infty$ inside $(0, 1)$? The result is given in [RR98a], where it is proven that for $\alpha = 1/3$, the acceptance rate converges in $(0, 1)$. The assumptions are the following.

M 1 U is eight times continuously differentiable on \mathbb{R} and there exists a real polynomial P satisfying for any $i \in \{0, \dots, 8\}$ and $x \in \mathbb{R}$, $|U^{(i)}(x)| \leq P(x)$.

M 2 For any $p \geq 1$, $\int_{\mathbb{R}} |x|^p \pi_1(x) dx < \infty$.

More precisely, [RR98a, Theorem 1] shows that

$$\lim_{d \rightarrow +\infty} \mathbb{E}_{\pi^d} \left[\alpha^d \left(X_0^d, X_0^d + \ell^2 d^{-1/3} \nabla \log \pi^d(X_0^d) + \sqrt{2} \ell d^{-1/6} Z_1^d \right) \right] = 2\Phi(-\ell^3 \sqrt{J}/2), \quad (1.20)$$

where Φ is still the cumulative distribution function of the standard normal distribution and the new constant J is given by

$$J = (48)^{-1} \int_{\mathbb{R}} \left\{ 5(U^{(3)}(x))^2 - 3(U^{(2)}(x))^3 \right\} \pi(x) dx > 0. \quad (1.21)$$

Second, [RR98a] also prove that a limiting process satisfies the Langevin stochastic differential equation. The authors in [RR98a] consider again the unidimensional process $(V_t^d)_{t \geq 0}$ given by (1.19), but this time $(\Gamma_t^d)_{t \geq 0}$ is a Poisson process with rate $d^{1/3}$. In [RR98a, Theorem 2], the authors prove that if U satisfies M1 and M2, then the sequence of processes $(V^d)_{d \in \mathbb{N}^*}$ converges weakly, in the Skorokhod topology, to a process V that again satisfies the following: V_0 is distributed according to π and for any $t \geq 0$, V satisfies the overdamped Langevin stochastic differential equation,

$$dV_t = \sqrt{h(\ell)} dB_t^1 - \frac{1}{2} h(\ell) U'(V_t) dt ,$$

where $(B_t^1)_{t \geq 0}$ is a unidimensional Brownian motion, $h(\ell)$ is often called the speed of the diffusion and is given for all $\ell > 0$ by

$$h(\ell) = 2\ell^2 \Phi \left(-\frac{\ell^3 \sqrt{J}}{2} \right) ,$$

and J is defined by (1.21).

This result is interpreted with the same tools as the one for the symmetric random walk Metropolis, meaning that the diffusion limit informs us that maximizing the speed h leads to the most efficient choice of ℓ . Using (1.20) and calculating h as a function of the limit acceptance, we find that to maximize the function $\ell \mapsto g(\ell)$, the limit mean acceptance ratio given in (1.20) has to be 0.574 (up to three digits).

There are two important differences with respect to random walk Metropolis. Let us interpret them to compare both algorithms. The first one is that the scaling of MALA is $\sigma_d = \ell d^{-1/6}$ whereas for random walk Metropolis, we had $\sigma_d = \ell d^{-1}$. This difference means that, in general, we can use larger values for σ_d using MALA, compared to using RWM. The fact that we can build a jump process from each of these Markov chains that converge weakly to the same limit tells us that when MALA requires $\mathcal{O}(d^{1/3})$ steps to attain a given point of the trajectory, RWM needs $\mathcal{O}(d)$ steps. Therefore, we conclude that MALA is more efficient than RWM, in the sense of optimal scaling. The second point is that the optimal mean acceptance is higher for MALA than for RWM. All in all, when the assumptions allow MALA to converge, it is a more efficient choice than RWM. We explore in this thesis the optimal scaling of MALA for the Laplace distribution, with density $x \mapsto e^{-|x|}/2$, because the assumption M1 no longer holds since $x \mapsto |x|$ is not differentiable at zero.

1.5 Setting of the Thesis and Contributions

1.5.1 Riemannian schemes

In Chapter 2, we study the standard geodesic scheme (1.6) with the extra possibility of having a bias term on top of the unbiased estimator H , meaning we consider the sequence

$$\theta_{n+1} = \text{Exp}_{\theta_n} [\eta_{n+1} (H_{\theta_n} (X_{n+1}) + b_{\theta_n} (X_{n+1}))] ,$$

where $b : \Theta \times \mathsf{X} \rightarrow \mathsf{T}\Theta$ is a bi-measurable function. We believe we are the first to allow for a bias in Riemannian stochastic approximation. We work under the assumption that the bias is uniformly bounded, meaning $\sup_{\theta, x} \|b_\theta(x)\|_\theta \leq b_\infty$ – see **A 3**. We also look at more general retraction schemes like (1.7). The analysis of these schemes is very important because retractions are usually much simpler to compute in practice, and have been used for a long time – see [AM12]. The most common example is to see a projected scheme as a retraction scheme e.g. scheme (1.5) is a retraction scheme on the sphere.

We contribute to the study of retraction schemes in Section 2.3.1 by introducing first order (see **R 2**) and second order retractions (see **R 3**), for which we quantify the error made by the retraction. Since we see retractions as approximations of the Riemannian exponential map, we introduce the notion of *retraction error map* in (2.13). The function $\Phi_\theta = \text{Exp}_\theta^{-1} \circ \text{Ret}_\theta : \mathsf{T}_\theta\Theta \rightarrow \mathsf{T}_\theta\Theta$ defined for any θ . It defines a bundle map $\Phi : \mathsf{T}\Theta \rightarrow \mathsf{T}\Theta$, which one can compare to the identity. In Lemmas 2.7 and 2.17, we provide operator bounds for $\theta \mapsto \|\Phi_\theta - \text{Id}_\theta\|_\theta$ by means of bounding the terms in a Taylor expansion. We also check that usual retractions on the sphere and on the Grassmann manifold are both first and second order retractions: see Example 2.8 and Proposition 2.18 for the sphere and Example 2.9 and Proposition 2.10 for the Grassmann manifold.

In Chapter 3, we study a *constant* step-size scheme, with the novel assumption that the Riemannian scheme is constrained on a closed and geodesically convex set using a projection-like operator – see (3.2) and **A 6**. This encompasses the standard geodesic scheme (1.6) under **A 6-(ii)**.

1.5.2 Noise assumptions

The common noise assumption is the Martingale difference setting (or Martingale setting in short). Assumption **A 1** rewrites neatly into the Riemannian context.

A 3 (Martingale difference setting). *The sequence $(e_{\theta_n}(X_{n+1}))_{n \in \mathbb{N}}$ is a martingale difference sequence with respect to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}^*}$, i.e. $\mathbb{E}[e_{\theta_n}(X_{n+1}) | \mathcal{F}_n] = 0_{\theta_n}$ for any $n \in \mathbb{N}$. Moreover, there exist $\sigma_0^2, \sigma_1^2 \geq 0$ such that for any $n \in \mathbb{N}$*

$$\mathbb{E} \left[\|e_{\theta_n}(X_{n+1})\|_{\theta_n}^2 \mid \mathcal{F}_n \right] \leq \sigma_0^2 + \sigma_1^2 \|h(\theta_n)\|_{\theta_n}^2 .$$

In both Chapters 2 and 3, we study the Martingale difference setting **A 3**. However, in Chapter 2, we are the first to generalize **A 4** to the Riemannian setting.

A 4 (Markovian setting). *There exists a Markov kernel P on $(\Theta \times \mathsf{X}) \times \mathcal{X}$ such that for any $n \in \mathbb{N}$ and bounded and measurable function $\varphi : \mathsf{X} \rightarrow \mathbb{R}_+$, $\mathbb{E}[\varphi(X_{n+1}) | \mathcal{F}_n] = \int_{\mathsf{X}} \varphi(y) P_{\theta_n}(X_n, dy)$. In addition, for any $\theta \in \Theta$, P_θ admits a unique invariant distribution π_θ satisfying $h(\theta) = \int_{\mathsf{X}} H_\theta(y) d\pi_\theta(y)$.*

In this setting, the noise $(X_n)_{n \in \mathbb{N}^*}$ is a Markov chain, with every transition being controlled by the current value of the parameter θ_n . In many real-world applications, one cannot draw independent samples from the distribution π_θ that appears in (3.1), because it is too complicated. In this case, it remains possible to sample from an ergodic Markov

chain, with stationary distribution π_θ . Our Markovian assumptions are **MA2**(w), but we present two cases in which these conditions are met. Recall that a set $\mathbf{C} \in \mathcal{X}$ is called small for a Markov kernel Q in $\mathbf{X} \times \mathcal{X}$ when there exists $m > 0$ and a non-trivial measure ν_m on \mathcal{X} such that for any $x \in \mathbf{C}$, $\mathbf{B} \in \mathcal{X}$, $Q^m(x, \mathbf{B}) \geq \nu_m(\mathbf{B})$.

A5 (w). *The family of Markov kernels $(P_\theta)_{\theta \in \Theta}$ and the measurable function $w : \mathbf{X} \rightarrow [1, +\infty)$ satisfy one of the following conditions:*

- (a) *The family $(P_\theta)_{\theta \in \Theta}$ is uniformly ergodic with constant uniform in θ , i.e. there exist $\varepsilon_P \in (0, 1)$ and $C_P \geq 0$ such that for any $\theta \in \Theta$, P_θ has a unique stationary distribution π_θ and for any $x, x' \in \mathbf{X}$, and $k \in \mathbb{N}$, $\|\delta_x P_\theta^k - \delta_{x'} P_\theta^k\|_{\text{TV}} \leq C_P(1 - \varepsilon_P)^k$. In addition, there exists $C \geq 0$ such that for any $\theta_0, \theta_1 \in \Theta$, $x \in \mathbf{X}$ and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0 and θ_1 , $\|\delta_x P_{\theta_1} - \delta_x P_{\theta_0}\|_{\text{TV}} \leq C\ell(\gamma)$.*
- (b) *For any $\theta \in \Theta$, $P_\theta = P$, where P is a Markov kernel on $\mathbf{X} \times \mathcal{X}$. Moreover, there exists $w : \mathbf{X} \rightarrow [1, +\infty)$, $\lambda \in (0, 1)$, $b_w \geq 0$ such that for any $x \in \mathbf{X}$, $Pw(x) \leq \lambda w(x) + b_w \mathbb{1}_{\mathbf{C}}(x)$ where $\mathbf{C} \in \mathcal{X}$ is a small set for P .*

In Proposition 2.6, we prove that under simple conditions, when **A5**-(a) holds – meaning that the chain is uniformly ergodic, with constant uniform in $\theta \in \Theta$ – then **MA2**(w) holds, with the constant function $w : x \mapsto 1$. Moreover, if **A5**(w)-(b) holds instead – meaning the Markov kernel P does not depend on θ and that there is a drift condition – then **MA2**(w) holds. We believe that this noise setting is interesting as it can model scenarios that the Martingale difference setting does not encompass.

In each noise setting, we make sure not to assume that the iterations of the scheme remain in some arbitrary compact set K , contrary to [Bon13; ZS16], for the reasons mentioned in Section 1.2.2.

1.5.3 Convergence Techniques

In Chapter 2, we generalize Lyapunov’s second method – see [LaS60] – in the Riemannian case. Our assumption **A2** is satisfied when the Lyapunov function V is gradient Lipschitz or has a continuous Riemannian Hessian with bounded operator norm – see Lemma 2.38.

The same setting is used in Chapter 3, where we exchange **A7** for **A6**-**A7**. The problem Lyapunov functions on Riemannian manifolds is that the usual candidate $\theta \mapsto \rho_{\Theta}^2(\theta, \theta^*)$ – the square distance with respect to the solution – is not gradient Lipschitz whenever Θ has non-zero curvature and is non-compact. We address this in Section 3.2.2, where we introduce the novel, “Huberized” version of the usual candidate – see Proposition 3.5.

1.5.4 Results of Chapters 2 and 3

Both in an interest of teaching the broader machine learning community about Riemannian geometry for stochastic approximation, of reformulating existing formulas in different scenarios and of using intelligible proofs with minimal background, my coauthors and I have written examples, definitions and theorems that regard only Riemannian geometry

throughout Chapters 2 and 3. We recommend any reader that would like to familiarize themselves with the concepts and notation of Riemannian geometry used in this thesis to start by reading the point-by-point preliminary Section 2.11.

In Section 2.6.2, we detail the sphere which was discussed in Example 1.6 and Example 2.8. We also use these notions of Riemannian geometry to prove that the projective retraction satisfies assumptions **R1-R2-R3** – see Proposition 2.18.

Two of the geometric concepts that required our closest attention were Taylor formulas (see Theorems 2.37 and 3.29) and the computation of the cut locus in the special case of the Grassmann manifold (see Proposition 2.26). The development of Taylor formulas was both difficult and precious for our kind of analysis. We wrote Section 3.11.2 to develop the Taylor formula (3.96) in Theorem 3.29, and specifically because we needed Proposition 3.30 which we did not find anywhere else in the Riemannian geometry community. It is a n -th order Taylor formula with integral reminder for smooth vector fields along a geodesic, using the notion of total covariant derivatives of any order. Much of Section 3.11.2 was written as an adaptation of the classical results of connections in the special case of vector fields *along a geodesic*. If the reader is unfamiliar with tensor fields and connections, we recommend reading also [Lee19, Chapter 4], since Section 3.11.2 is inspired from it. After developing this Taylor formula, we adapted it to the simpler case of a Taylor formula for a real-valued function in Section 2.11.9 and Theorem 2.37.

Our last contributions that purely concern Riemannian geometry are a presentation the Grassmann manifold in a self-contained fashion. Our main goal was to express the formulas for the Riemannian exponential in [AMS04; BZA20; EAS98] as well as computing its cut locus. In order to help the reader, the summary of Section 2.7.2 is the following. We start by giving the definition of the Grassmann manifold via a quotient of the *non-compact* Stiefel manifold using the action of the linear group. Then, we want to give the differential structure of the Grassmann manifold. To do so, we start by building the *cross section* map $\sigma_{\mathbf{B}}$ in Lemma 2.21. The inverse of this map is then used to explicitly write the family of parameterizations – or, conversely, the *smooth atlas* – that define the differential structure of the Grassmann manifold, see Proposition 2.22. We continue by describing the tangent space, using the canonical decomposition in the case of quotient manifolds. We recall the definition of connection in the sense of [KN96, Chapter II, Section 1] in Definition 2.24, we check that (2.71) is a connection in this sense and, combined with Proposition 2.23, we finally obtain a description of the tangent space. With such tools, we give an alternative formulation of the Riemannian exponential in Lemma 2.25. This formula is used in our final result in this section, Proposition 2.26, in which we compute the cut locus of the Grassmann manifold. This last proof is the result of a long work and, despite being long and intricate, remains simpler than the only other proof that we know of this result – see [Sak76]. Indeed, the proof of [Sak76] uses more complicated tools of Riemannian geometry, whereas ours uses as little concepts as possible.

In Chapter 2, we derive global non-asymptotic convergence rates under all the different noise assumptions and Lyapunov assumptions given above. For the geodesic scheme studied in Section 2.2, we obtain in Theorems 2.4 and 2.5 that after n iterations, $\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] = \mathcal{O}(b_\infty + \log n/\sqrt{n})$, where b_∞ is the uniform bound on the bias of H and I_n is a discrete

random variable supported on $\{1, \dots, n\}$. A more precise discussion on the choice of step-sizes and its implications is given after said theorems. For the retraction schemes studied in Section 2.3.2, we obtain in Theorems 2.11, 2.12 and 2.29 that after n iterations, similarly to the geodesic scheme, $\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] = \mathcal{O}(b_\infty + \log n/\sqrt{n})$. The interest of this analysis compared to the works of [ZS16] is that 1) we do not require the iterations to remain in some bounded subset of the manifold, 2) we establish these bounds for the general stochastic approximation (not only for gradient descent), 3) the assumptions needed for the Lyapunov function allow for the optimization of non-convex functions.

Assumptions			Result
Base	Scheme	Noise	
A1-A2-A3	Geodesic (2.3)	Martingale Setting MD1	Theorem 2.4
A1-A2-A3	Geodesic (2.3)	Markovian Setting MA1-MA2(w)	Theorem 2.5
A1-A2-A3 A4	Retraction (2.12) R1 and (R2 or R3)	Martingale Setting MD1 and (MD2(4) or MD2(6))	Theorem 2.11
A1-A2-A3 A4	Retraction (2.12) R1 and (R2 or R3)	Markovian Setting MA1-MA2(w)-MA4(w)	Theorem 2.12

Table 1.1: Summary of the non-asymptotic results for each set of assumptions.

Lastly, we consider two example applications for these results. First, an online PCA which deals with the problem introduced in Example 1.10 but in the general case $r \geq 1$. In section Section 2.4.1, we prove that the schemes (2.19) and (2.20) satisfy the conditions of our general theorems – see Theorem 2.13. We complete the analysis with some numerical simulations – see Section 2.4.3. Second, we look at an online barycentre estimation in Section 2.4.2. In this application, we examine the same starting idea of Example 1.1, which we refine to define the concept of robust Riemannian barycentre of a probability distribution on a Hadamard manifold, for which we prove the existence and uniqueness in Proposition 2.14 and derive a scheme for its computation (2.23), which satisfies the general conditions of our theorems – see Proposition 2.15.

In Chapter 3, we focus on the constant step-size setting. We start by proving non-asymptotic bounds in Theorem 3.1, which mostly resemble Theorem 2.4. However, we have a stronger bound specifically in (3.6), where we observe the exponential forgetting of the initial conditions, much like for strongly convex functions in the Euclidean case. In Theorem 3.2, we prove that the sequence of iterates $(\theta_n)_{n \in \mathbb{N}}$ is a geometrically ergodic Markov chain, with a unique stationary distribution μ^η that depends on the step-size η . From this point forward, we seek to understand the behaviour of μ^η when $\eta \rightarrow 0$, both with asymptotic and non-asymptotic results. The first step in this direction is Theorem 3.3, where we prove that $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$ converges weakly to the Dirac measure at θ^* , where θ^* is the unique solution to the root finding problem (3.1). In Section 3.3, with the unconstrained scheme (1.6), we quantify the convergence derived in Theorem 3.3-(b).

First, we prove an asymptotic expansion for the bias $\int_{\Theta} g(\theta) d\mu^{\eta}(\theta) - g(\theta^*)$ with respect to the step size η for g belonging to a certain class of smooth functions from Θ to \mathbb{R} – see Theorem 3.6. By applying our result to SGD ($h = -\text{grad } f$), it implies an asymptotic expansion of $\int_{\Theta} \|\text{grad } f(\theta)\|_{\theta}^2 d\mu^{\eta}(\theta)$. Secondly, we establish that the convergence derived in Theorem 3.3-(b) occurs at a rate $\eta^{1/2}$, through a central limit theorem for $(\mu^{\eta})_{\eta \in (0, \bar{\eta}]}$ – see Theorem 3.7. These two results can be understood as a bias-variance decomposition in which both terms are of order η and are therefore weak counterparts of [DDB17, Proposition 3, Theorem 5], [Pfl86, Theorem 1] in a Riemannian setting. Next, in Section 3.4, we look at the consequences of these results on the SGD case. We consider two possible set of assumptions. When f is strongly geodesically convex – see Definition 1.12 and **F2**, we have strong results in Corollary 3.9. We can see, as in the Euclidean case, that it takes $\mathcal{O}(\log(1/\varepsilon)\varepsilon^{-1})$ steps to have $\mathbb{E}[f(\theta_n) - f(\theta^*)] \leq \varepsilon$. On the other hand, when f is geodesically quasi-convex, which Euclidean counterpart is $\langle \text{grad } f(\theta), \theta - \theta^* \rangle \geq \tau[f(\theta) - f(\theta^*)]$ as given in [HMR19], we obtain weaker results in Proposition 3.11. They are weaker not only because the right-hand bound is $C + \mathcal{O}(n^{-1})$ – which closely resembles the results in Chapter 2 – but because the left-hand term is expressed through the “distance-like” function given in (3.14). The most novel part of this result is that in this second setting, we do not require boundedness conditions on Θ . Lastly, we apply these SGD results in the particular case of the barycentre estimation. To use both convexity assumptions, we look at the estimation of the barycentre of a discrete distribution – which yields a strongly convex function f – and a continuous distribution – yielding a geodesically quasi-convex function. The second case is novel in the literature as it does not have boundedness conditions on the distribution. For the second case we introduce the novel scheme (3.17), for which we prove non-asymptotic bounds in Theorem 3.13. The barycentre estimation is then illustrated with numerical experiments in Section 3.5 showcasing both sets of assumptions.

1.5.5 Results of Chapter 4

In Chapter 4, we study a generalization of MALA, first studied in [Per16], which aims at alleviating the smoothness assumption of the target distribution π of the form (1.17). If $\pi \propto \exp(-G)$, then we assume that the function $G : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex, proper and lower semi-continuous. Instead of using the real gradient of G in the proposal, we approximate G by its Moreau-Yoshida envelope $G^{\lambda} : \mathbb{R}^d \rightarrow \mathbb{R}$ for $\lambda \geq 0$ defined by

$$G^{\lambda}(\mathbf{x}) = \min_{\mathbf{u} \in \mathbb{R}^d} [G(\mathbf{u}) + \|\mathbf{u} - \mathbf{x}\|^2 / (2\lambda)] .$$

By setting $\pi^{\lambda} = \exp(-G^{\lambda})$, we obtain a smooth approximation of π , where λ controls the regularity of π^{λ} . Under these assumptions, G^{λ} is continuously differentiable and therefore we use MALA to target π^{λ} , where the proposal is of the form

$$\xi_{k+1} = \xi_k - \frac{\sigma^2}{2} \nabla G^{\lambda}(\xi_k) + \sigma Z_{k+1} ,$$

see (4.4) for the expression of the gradient of G^{λ} . Our goal is to produce theoretical results for the class of proximal MALA algorithms, study its asymptotic behavior when the

dimension d goes to infinity. We also interpret our results to provide guidelines regarding the implementation of these proximal algorithms.

We start by looking at the particular case when π is actually differentiable – see Section 4.3.1 – in an interest of comparing gradient-based methods (like MALA) with proximal-based methods like proximal MALA. In our analysis, we let the step size σ^2 and the regularization parameter λ depend on the dimension d of the target. In Theorem 4.1 we derive that there are two possible regimes of scaling. First, in cases (b) and (c), we establish that if λ goes to 0 faster or at the same rate as σ^2 , the scaling is equivalent to the one of MALA, meaning that σ^2 is proportional to $d^{-1/3}$. However, when λ goes to 0 at a slower rate than σ^2 , we find a new scaling, where σ^2 is proportional to $d^{-1/2}$ – see (a). In this second scaling regime, proximal MALA is less efficient than MALA. In both cases, the convergence of the acceptance rate $a_d(\ell, r) \xrightarrow{d \rightarrow \infty} a(\ell, r)$ is accompanied with weak convergence of the linear interpolation of the first component of chain $(X_k^d)_{k \geq 0}$, defined by

$$L_t^d = (\lceil d^{2\alpha}t \rceil - d^{2\alpha}t)X_{\lfloor d^{2\alpha}t \rfloor, 1}^d + (d^{2\alpha}t - \lfloor d^{2\alpha}t \rfloor)X_{\lceil d^{2\alpha}t \rceil, 1}^d, \quad (1.22)$$

towards the Langevin diffusion

$$dL_t = h(\ell, r)^{1/2}dB_t - \frac{h(\ell, r)}{2}g'(x)dt,$$

where $(B_t)_{t \geq 0}$ denotes standard Brownian motion and $h(\ell, r) = \ell^2 a(\ell, r)$ is the speed of the diffusion. Although proximal MALA is at best as efficient as MALA, since sometimes the proximal operator for a given distribution π is cheaper to compute than $\nabla \log \pi$ [CP11; PB14; Per16], we believe that there are cases where for a properly chosen λ , proximal MALA should provide a computationally cheaper alternative to MALA with comparable efficiency.

Next, we study proximal MALA for a special non-smooth target, the Laplace distribution $x \mapsto 2^{-1} \exp(-|x|)$. Although we only prove it for the particular case of Laplace distributions, our proof technique is not unique to this case, and could be applied to other non-smooth distributions, provided some computations of integrals are carried out – see Section 4.10. We chose this particular non-smooth target because it is commonly used in practical applications of proximal MALA, including image deconvolution [DMP18; Per16; VDC19], LASSO, and sparse regression [Atc15; GSS21; ZCZ22]. We obtain the scaling to achieve convergence of the acceptance rate in Theorem 4.4. We observe that a single point of non-differentiability is sufficient to obtain a different scaling than MALA for differentiable targets. Indeed, for the Laplace distribution, the step size σ^2 needs to be proportional to $d^{-2/3}$ (whereas MALA for smooth targets is $d^{-1/3}$) to obtain a non-trivial limit for the acceptance rate. This optimal scaling is new in the study of MCMC algorithms, somewhere between the rate d^{-1} of RWM and $d^{-1/3}$ of (differentiable) MALA. For this study, we only investigated the case when λ converges to 0 at least as fast as σ^2 . As is usual for optimal scaling, we continue our study with the weak convergence of the interpolated first component of the chain $(L_t^d)_{t \geq 0}$ defined in (1.22) towards a Langevin

diffusion. We proceed in two steps, proving first the tightness of the sequence of laws of $\{(L_t^d)_{t \geq 0} : d \in \mathbb{N}\}$ – see Proposition 4.5, and then, proving that the only limit point possible for this sequence is indeed a solution of the Langevin diffusion

$$dL_t = [h^L(\ell)]^{1/2} dB_t - \frac{h^L(\ell)}{2} \operatorname{sgn}(L_t) dt .$$

This result is done using the Martingale problem associated with the diffusion – see Proposition 4.6. Our final result Theorem 4.7 combines these two into the actual optimal scaling result of the Laplace distribution for proximal MALA. It is important to remark that one can extend the definition of this method when $\lambda = 0$, leading to sub-gradient MALA or sG-MALA. All of our results regarding the Laplace distribution also hold with this generalization, meaning that we also have proven the proximal scaling of sG-MALA for the non-smooth Laplace distribution.

We complete our results with a discussion on the implications of our previous theorems – see Section 4.4 – as well as numerical experiments that illustrate both the scaling of differentiable targets and of the Laplace distribution for proximal MALA, with multiple choices for the rates of convergence to zero of the two parameters σ^2 and λ .

CHAPTER 2

Convergence Analysis of Riemannian Stochastic Approximation Schemes

ALAIN DURMUS,¹ PABLO JIMÉNEZ¹, ÉRIC MOULINES¹, SALEM SAID², HOI-TO WAI³.

This paper analyzes the convergence for a class of Riemannian stochastic approximation (SA) schemes, designed to tackle optimization problems on Riemannian manifolds. We consider SA schemes which use either the exponential map of the considered manifold or retraction maps used as a proxy for the exponential map. The relaxation to retraction maps is of considerable interest since in many problems the computation of the exponential map is either costly or even untractable. Our results are derived under mild assumptions. First, our results are global in the sense that we do not assume iterates to be a-priori bounded. Second, we allow the presence of a bias in SA schemes. To be more specific, we consider the case where the mean-field function can only be estimated up to a bias, and/or the case in which the samples are drawn from a controlled Markov chain. Third, the assumptions on retractions required to ensure convergence of the related SA schemes are weak and are satisfied in many settings. We illustrate our findings on applications to principal component analysis and Riemannian barycenter problems.

Contents

2.1 Introduction	31
2.2 Geodesic Stochastic Approximation Schemes	33
2.3 General Retraction SA Schemes	38
2.3.1 Retraction and Quantitative Estimates of the Retraction Error	38

¹CMAP, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France (alain.durmus@polytechnique.edu, pablo.jimenez-moreno@polytechnique.edu, eric.moulines@polytechnique.edu).

²Laboratoire Jean Kuntzmann, CNRS, Université Grenoble Alpes (salem.said@univ-grenoble-alpes.fr).

³Department of SEEM, The Chinese University of Hong Kong (htwai@se.cuhk.edu.hk).

2.3.2	Analysis of Retraction SA Schemes	40
2.4	Applications	41
2.4.1	Principal Component Analysis	42
2.4.2	Robust Barycenter in a Hadamard Manifold	43
2.4.3	Numerical Examples on the PCA Problem	45
2.5	Proofs for Section 2.2	46
2.5.1	Proof of Lemma 2.1	46
2.5.2	A technical lemma: proof of Lemma 2.3	46
2.5.3	Proof of Theorem 2.4	47
2.5.4	Proof of Proposition 2.6	48
2.5.5	Proof of Theorem 2.5	51
2.6	Proofs for Section 2.3.1	54
2.6.1	Proof of Lemma 2.7	54
2.6.2	Projective Retraction on the Sphere	56
2.6.3	Proof of Proposition 2.10	58
2.7	Complementary proofs for Section 2.3.1	60
2.7.1	Proof of Proposition 2.18	60
2.7.2	Details for the Grassmann manifold	61
2.8	Proofs for Section 2.3.2	81
2.8.1	A technical lemma	81
2.8.2	Proof of Theorem 2.11	83
2.8.3	Proof of Theorem 2.12	85
2.9	Proof of Section 2.4	92
2.9.1	Proofs of Section 2.4.1	92
2.9.2	Proofs of Section 2.4.2	96
2.9.3	Details and specification for Section 2.4.3	100
2.10	Useful results from matrix analysis	101
2.11	Preliminaries on Riemannian Geometry	103
2.11.1	Metric Tensor and Distance	104
2.11.2	Levi-Civita Connection	104
2.11.3	Geodesic Equation	105
2.11.4	Parallel Transport and Parallel Frames	105
2.11.5	Riemannian Exponential Map and Cut Locus	106
2.11.6	Injectivity Domain	107
2.11.7	Isometries and Homogeneous Spaces	107
2.11.8	Riemannian Gradient, Hessian	108
2.11.9	Taylor formula along a geodesic	109
2.11.10	Bounded Hessian Implies Lipschitz Gradient	111

2.1 Introduction

This paper is concerned with the root finding problem on a smooth Riemannian manifold Θ :

$$\text{find } \theta \in \Theta \text{ satisfying } h(\theta) = 0_\theta, \quad \text{where } h(\theta) = \int_{\mathcal{X}} H_\theta(x) d\pi_\theta(x), \quad (2.1)$$

such that $h : \Theta \rightarrow T\Theta$ is a smooth vector field, called the mean vector field, for any $\theta \in \Theta$, π_θ is a distribution over $(\mathcal{X}, \mathcal{X})$ and $H : \Theta \times \mathcal{X} \rightarrow T\Theta$ is a bimeasurable function satisfying for any $\theta \in \Theta$, $H_\theta : \mathcal{X} \rightarrow T_\theta\Theta$. This framework includes stochastic optimization problems related to a smooth (but possibly non-convex even in the geodesic sense) objective function $f : \Theta \rightarrow \mathbb{R}^*$, where $h = \text{grad } f$ and grad is the Riemannian gradient operator. Problem (2.1) arises in many applications such as principal component analysis (PCA) and the computation of geometric barycenter [ABY13; Oja92].

Our objective is to study Riemannian stochastic approximation (SA) schemes to compute an approximate solution of (2.1). In the case where the geodesic curves on Θ can be explicitly computed, the Riemannian SA scheme is defined through the recursion: for $n \in \mathbb{N}$,

$$\theta_{n+1} = \text{Exp}_{\theta_n} \{ \eta_{n+1} H_{\theta_n}(X_{n+1}) \}, \quad (2.2)$$

where Exp_{θ_n} is the exponential map at θ_n , $(X_n)_{n \in \mathbb{N}^*}$ is a sequence of *independent* random variables in $(\mathcal{X}, \mathcal{X})$ such that X_n has distribution π_{θ_n} for any $n \in \mathbb{N}^*$, and $(\eta_n)_{n \in \mathbb{N}^*}$ is a sequence of positive stepsizes. The recursion (2.2) is a natural extension of the Robbins-Monro algorithm originally applied on the Euclidean setting [RM51], where $\text{Exp}_\theta(y) = \theta + y$ for any $\theta, y \in \mathbb{R}^d$.

A number of results have been reported for Riemannian SA (2.2) under the stochastic gradient (SG) setting where $h = \text{grad } f$ [HS20]. The pioneering work of [Bon13] establishes the asymptotic convergence of (2.2) to a critical point using martingale techniques adapted from [BMP90; Bor09; KY03]. Later, [ZS16; ZRS16; ZZS18] provide non-asymptotic analysis under the assumption that the objective function f is geodesically convex (g-convex), and that all the iterates remain in a compact subset of Θ . Another related work is [Tri+18] which considers a *retraction* based averaging scheme and shows a central limit theorem for (2.2). Analysis of Riemannian SA for non-convex objective is addressed in [HS19]. However, the authors require a strong assumption on the retraction map, the objective function and the manifold. We also mention that a few other papers deal with deterministic optimization on Riemannian manifold and the convergence to local minimum in non-convex settings [BAC18; CB19; SFF19], or applying Nesterov's acceleration [AS20] for g-convex functions. To deal with relatively general conditions on the mean-field h , [Sha19] analyzes stochastic recursion schemes adapting the well-known ODE method in [BMP90; KY03] to the Riemannian setting to show asymptotic convergence where Θ is compact.

The present paper considers several relaxations of the scheme defined by (2.2). First, we assume that $(X_n)_{n \in \mathbb{N}^*}$ is a Markov chain whose transition is controlled by the current value of the parameter θ_n . This is motivated by the fact that, in many real-world applications, it is impossible to obtain independent samples from the distribution π_θ appearing in (2.1),

since the distribution is too complicated. Typically, in this case, it remains possible to sample from an ergodic Markov chain $(X_n)_{n \in \mathbb{N}^*}$ whose stationary distribution is π_θ . For example, in Bayesian inference, where π_θ is the posterior distribution, π_θ is practically never directly accessible, and may only be sampled by using Markov chain Monte Carlo (MCMC) techniques. Second, we deal with the use of retraction operators in place of the exponential map in the recursion (2.2). At last, we allow for the non-SG setting when $h \neq \text{grad } f$ as well as errors in the numerical evaluations of H – such as systematic errors that do not come from the randomness of the environment – resulting in a potentially biased vector field in (2.2). Note that the non-SG setting is common for SA: examples include, among many others, online Expectation Maximization algorithms [CM09] and policy gradient [BB01]. Hence, our work is in line with recent studies of SA in the Euclidean setting for minimizing a non-convex function in the controlled Markovian setting under relaxed conditions on the mean-field h , see e.g. [BCN18; GL13; Kar+19] and the references therein. Our contributions are:

- We perform a global convergence analysis of a *biased* geodesic Riemannian SA scheme without assuming a bounded domain for the iterates $(\theta_n)_{n \in \mathbb{N}}$ nor g -convexity. Under this relaxed setting, the iterates of the geodesic SA scheme satisfy $\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] = \mathcal{O}(b_\infty + \log n/\sqrt{n})$, where b_∞ is the asymptotic bias of the vector field H and I_n is a discrete random variable supported on $\{1, \dots, n\}$. We cover settings where the noise is a sequence of martingale increments or stems from a controlled Markovian dynamics, see Theorems 2.4 and 2.5.
- We consider Riemannian SA schemes with first or second order retraction. Note that retraction operators have met great success in practice (see e.g. [AM12]) since they are computationally cheaper to evaluate. Similar to the geodesic scheme, our methods find an $\mathcal{O}(b_\infty + \log n/\sqrt{n})$ -stationary point in n iterations in expectation, see Theorems 2.11 and 2.12. We illustrate that the conditions required on the retraction function by our results hold on several examples.
- We consider examples of applications on PCA (see Section 2.4.1 & Theorem 2.13) and barycenter estimation (see Section 2.4.2), for which we show that the required convergence conditions are satisfied. For the second application, we introduce the concept of robust Riemannian barycenter of a probability distribution defined on a Hadamard manifold, prove the existence and uniqueness of this robust barycenter (see Proposition 2.14), and derive an SA scheme for its computation (see (2.23) and Proposition 2.15).

Notation

For any two sequences of real numbers $(u_n)_{n \in \mathbb{N}}$ and $(v_n)_{n \in \mathbb{N}}$, we write $u_n = \mathcal{O}(v_n)$ when there exist $n_0 \in \mathbb{N}$ and $M_0 \in \mathbb{R}_+$ such that for any $n \geq n_0$, $|u_n| \leq M_0 v_n$. If $u_n = \mathcal{O}(v_n)$ and $v_n = \mathcal{O}(u_n)$, we write $u_n = \Omega(v_n)$. If $\mu = \mu^+ - \mu^-$ is a signed measure over a measurable space (E, \mathcal{A}) , where μ^+ and μ^- are non-negative measures, called the positive and negative part of μ , define the total variation $\|\mu\|_{\text{TV}} = \mu^+(E) + \mu^-(E)$. For any $d \in \mathbb{N}^*$, \mathbf{I}_d is the identity matrix of size $d \times d$. We denote the tangent space of Θ at θ by

$T_\theta\Theta$ and its tangent bundle $T\Theta$. If Θ_0 and Θ_1 are two smooth manifolds, for any smooth function $f : \Theta_0 \rightarrow \Theta_1$, we denote its derivative by $Df : T\Theta_0 \rightarrow T\Theta_1$. The Riemannian metric on Θ is denoted by g but for ease of notation and if there is no risk of confusion, for any $\theta \in \Theta$, $u, v \in T_\theta\Theta$, we should denote $g_\theta(u, v) = \langle u, v \rangle_\theta$ and $g_\theta(u, u) = \|u\|_\theta^2$. For $\theta \in \Theta$ and $k, l \in \mathbb{N}$, the set of (k, l) -tensors on $T_\theta\Theta$ is denoted by $T^{(k, l)}T_\theta\Theta$ and the set of smooth (k, l) -tensor fields on Θ is denoted $\Gamma(T^{(k, l)}T\Theta)$, see [Lee19, Appendix B] for a complete description of tensors and tensor fields. Let ∇ be the Levi-Civita connection of the metric g on Θ (see Section 2.11.2). We The map $T_{t_0 t_1}^\gamma : T_{\gamma(t_0)}\Theta \rightarrow T_{\gamma(t_1)}\Theta$ stands for the parallel transport associated to the Levi-Civita connection along a curve $\gamma : I \rightarrow \Theta$ from $\gamma(t_0)$ to $\gamma(t_1)$. In the interest of space, we leave detailed definitions and generalities on Riemannian geometry to Section 2.11.

2.2 Geodesic Stochastic Approximation Schemes

Let $(X, \mathcal{X}, \mathbb{P}, (\mathcal{F}_n)_{n \in \mathbb{N}})$ be a filtered probability space and consider $(X_n)_{n \in \mathbb{N}^*}$, an $(\mathcal{F}_n)_{n \in \mathbb{N}}$ -adapted stochastic process. The present paper studies stochastic approximation (SA) sequences $(\theta_n)_{n \in \mathbb{N}}$, used as approximate solutions of (2.1), starting from $\theta_0 \in \Theta$ and defined by the recursion

$$\theta_{n+1} = \text{Exp}_{\theta_n} \{ \eta_{n+1} (H_{\theta_n}(X_{n+1}) + b_{\theta_n}(X_{n+1})) \} , \quad (2.3)$$

where H is given in (2.1), $b : \Theta \times X \rightarrow T\Theta$ is a bi-measurable function satisfying for any $\theta \in \Theta$, $b_\theta : X \rightarrow T_\theta\Theta$, and $(\eta_n)_{n \in \mathbb{N}^*}$ is a sequence of positive stepsizes. Compared to the recursion given in the introduction (2.2), $b_{\theta_n}(X_{n+1})$ is an additional bias term used to explicitly model cases when h cannot be evaluated no matter how we sample.

We discuss some basic assumptions to be used throughout this paper. First, we ensure that the exponential map Exp is well defined over $T\Theta$ with the following condition (for a definition of the Riemannian exponential mapping, see Section 2.11.5).

A1. Θ is a geodesically complete Riemannian manifold of dimension $d \in \mathbb{N}^*$.

In addition, we assume the existence of a Lyapunov function V for the vector field h .

A2. There exists a continuously differentiable function $V : \Theta \rightarrow \mathbb{R}_+^*$ satisfying:

(a) There exist constants $\underline{c}_1, \underline{c}_2, \bar{c} > 0$ such that for any $\theta \in \Theta$,

$$\underline{c}_1 \|h(\theta)\|_\theta^2 \leq \underline{c}_2 - \langle \text{grad } V(\theta), h(\theta) \rangle_\theta , \quad \|\text{grad } V(\theta)\|_\theta \leq \bar{c} \|h(\theta)\|_\theta . \quad (2.4)$$

(b) The Riemannian gradient $\text{grad } V$ (see Section 2.11.8) is geodesically L -Lipschitz, i.e. there exists $L \geq 0$ such that for any $\theta_0, \theta_1 \in \Theta$, and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0 and θ_1 ,

$$\|\text{grad } V(\theta_1) - T_{\theta_0 \theta_1}^\gamma \text{grad } V(\theta_0)\|_{\theta_1} \leq L\ell(\gamma) ,$$

where $\ell(\gamma) = \|\dot{\gamma}(0)\|_{\theta_0}$ is the length of the geodesic.

The first inequality in (a) is a Lyapunov condition (see [LaS60]), and the constant c_2 underlines a mismatch between the orientations of the mean-field and the gradient of the Lyapunov function. In addition, **A 2-(a)** implies $\|h(\theta)\|_\theta \leq C_1 \|\text{grad } V(\theta)\|_\theta + C_2$ for any $\theta \in \Theta$, for some constants $C_1, C_2 \geq 0$. Meanwhile, (b) is satisfied if V has a continuous Riemannian Hessian $\text{Hess } V(\theta) : T_\theta\Theta \rightarrow T_\theta\Theta$ with a bounded operator norm for all $\theta \in \Theta$, see Lemma 2.38 in Section 2.11.10. Overall, conditions (a) and (b) ensure the stability of the recursion (2.3) as they imply that h is sublinear, *i.e.*, for any $\theta_0, \theta_1 \in \Theta$, there exists $C \geq 0$ such that for any geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0, θ_1 , it holds since $T_{\theta_0}^\gamma$ is an isometry by [Lee19, Proposition 5.5-(f)], $\|h(\theta_1)\|_{\theta_1} \leq C_1 \|\text{grad } V(\theta_1)\|_{\theta_1} + C_2 \leq C_3(\ell(\gamma) + 1)$, for $C_3 \geq 0$. Importantly, **A 1-A 2** allow us to generalize the descent lemma to the Riemannian setting, as follows (see [ZS16]).

Lemma 2.1. *Assume **A 1**, **A 2-(b)** hold. For any $\theta_0, \theta_1 \in \Theta$ and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0, θ_1 ,*

$$\left| V(\theta_1) - V(\theta_0) - \langle \text{grad } V(\theta_0), \dot{\gamma}(0) \rangle_{\theta_0} \right| \leq L\ell(\gamma)^2/2.$$

Proof. For completeness, the proof is given in Section 2.5.1. \square

This derivation is crucial for our proofs, and exists in different settings in the literature, *e.g.* [AMS08, Lemma 7.4.7], which differs with this result because Lemma 2.1 holds for any geodesic curves, and is not limited to the length-minimizing ones. Finally, we assume that the bias term in (2.3) is uniformly bounded.

A 3. *There exists a constant b_∞ such that $\sup_{(\theta, x) \in \Theta \times X} \|b_\theta(x)\|_\theta \leq b_\infty$.*

We provide non-asymptotic guarantees for the geodesic SA scheme (2.3) to find an (approximate) stationary point. Roughly, these results ensure the ability of the scheme (2.3) to produce an estimator $\tilde{\theta} \in \Theta$ based on $\{\theta_1, \dots, \theta_n\}$, satisfying $\mathbb{E}[\|h(\tilde{\theta})\|_{\tilde{\theta}}^2] \leq \epsilon + \tilde{c}b_\infty$, for a given target precision $\epsilon > 0$ and if n is large enough. More precisely, the estimator that we consider in this paper is the I_n -th iterate θ_{I_n} , where I_n is a random variable independent of θ_0 and \mathcal{F}_n , with distribution, for any $\ell \in \{0, \dots, n\}$,

$$\mathbb{P}(I_n = \ell) = \left(\sum_{i=0}^n \eta_{i+1} \right)^{-1} \eta_{\ell+1}. \quad (2.5)$$

To simplify notation in our subsequent discussions, define for any $n \in \mathbb{N}$ and $p \geq 1$,

$$\Gamma_n^{(p)} = \sum_{k=1}^n \eta_k^p, \quad \Gamma_n = \Gamma_n^{(1)}.$$

To gain intuition on the analysis for geodesic SA scheme (2.2), let us present a simple proof in a simplified setting.

Proposition 2.2. *Consider the sequence $(\theta_n)_{n \in \mathbb{N}}$ defined by the deterministic version of (2.2), *i.e.* $H_\theta(x) = h(\theta)$ and $b_\theta(x) = 0$ for any $x \in X$ and $\theta \in \Theta$. Assume **A 1-A 2** and $\sup_{k \in \mathbb{N}^*} \eta_k \leq c_1/L$. Then, for any $n \in \mathbb{N}$ we have,*

$$\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] \leq 2c_1^{-1} \{V(\theta_0) - V(\theta_{n+1})\} / \Gamma_{n+1} + 2c_1^{-1} c_2, \quad (2.6)$$

where I_n has distribution defined by (2.5).

2.2. Geodesic Stochastic Approximation Schemes

Proof. The analysis here is a consequence of Lemma 2.1. For any $k \geq 0$ and $t \in [0, 1]$, we observe that the geodesic segment $\gamma^{(k)}(t) = \text{Exp}_{\theta_k}(t\eta_{k+1}h(\theta_k))$ satisfies $\dot{\gamma}^{(k)}(0) = \eta_{k+1}h(\theta_k)$ and $\ell(\gamma^{(k)}) = \eta_{k+1} \|h(\theta_k)\|_{\theta_k}$; see [Lee19, Corollary 5.6-(b)]. Applying Lemma 2.1 yields,

$$\begin{aligned} \left| V(\theta_{k+1}) - V(\theta_k) - \eta_{k+1} \langle \text{grad } V(\theta_k), h(\theta_k) \rangle_{\theta_k} \right| &\leq (L/2)\ell(\gamma^{(k)})^2 \\ &= (L\eta_{k+1}^2/2) \|h(\theta_k)\|_{\theta_k}^2 . \end{aligned}$$

Using A2-(a), we obtain for any $k \in \mathbb{N}^*$

$$\sum_{k=0}^n \eta_{k+1}(\underline{c}_1 - (L/2)\eta_{k+1}) \|h(\theta_k)\|_{\theta_k}^2 \leq V(\theta_0) - V(\theta_{n+1}) + \underline{c}_2\Gamma_{n+1} . \quad (2.7)$$

Setting the step size as $\sup_{k \in \mathbb{N}^*} \eta_k \leq \underline{c}_1/L$ and using (2.5), we obtain by integrating with respect to I_n the result (2.6). \square

In the case were $\underline{c}_2 = 0$ in A2-(a) and taking a constant step size $\eta_k = \underline{c}_1/L$, Proposition 2.2 implies that $\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] = \mathcal{O}(1/n)$.

We now consider the the actual geodesic SA scheme (2.3). Define

$$e_\theta(x) = H_\theta(x) - h(\theta) . \quad (2.8)$$

Note that for any $x \in \mathsf{X}$, $\theta \mapsto e_\theta(x)$ is a vector field on Θ corresponding to the noise in the estimation of the mean field h . By Lemma 2.1 and arguments paralleling those used to derive (2.7), we show

Lemma 2.3. *Assume A1, A2. Setting $\Delta M_k = \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k}$ for any $k \in \mathbb{N}$, we have for any $n \in \mathbb{N}^*$ and $\varepsilon > 0$,*

$$\begin{aligned} &\sum_{k=0}^n \eta_{k+1}(\underline{c}_1 - (3L/2)\eta_{k+1} - \bar{c}^2\varepsilon) \|h(\theta_k)\|_{\theta_k}^2 \\ &\leq \sum_{k=0}^n \eta_{k+1} \Delta M_k + (3L/2) \sum_{k=0}^n \eta_{k+1}^2 \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \\ &\quad + V(\theta_0) + \underline{c}_2\Gamma_{n+1} + \sum_{k=0}^n \eta_{k+1} \{ (4\varepsilon)^{-1} + (3L/2)\eta_{k+1} \} \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2 . \end{aligned}$$

Proof. The proof is postponed to Section 2.5.2. \square

The terms in the right-hand side of the inequality given by Lemma 2.3 correspond to different sources of error that we can identify. First, the two terms $\sum_{k=0}^n \eta_{k+1} \Delta M_k$ and $\sum_{k=0}^n \eta_{k+1}^2 \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2$ come from the noise vector field, in the first and second-order approximation given by Lemma 2.1 respectively. The other term gathers the different sources of bias. More specifically, $V(\theta_0)$ corresponds to the initial conditions, $\underline{c}_2\Gamma_{n+1}$ is introduced by the constant \underline{c}_2 underlined in A2-(a) and the last term results from the bias on the measures of the mean-field function.

With a sufficiently small step size, the left-hand side in the inequality above can be lower bounded by $\sum_{k=0}^n \eta_{k+1} \|h(\theta_k)\|_{\theta_k}^2$. Meanwhile, the right-hand side can be controlled if $\sup_{n \in \mathbb{N}} \sum_{k=0}^n \eta_{k+1} \mathbb{E}[\Delta M_k] < +\infty$. To control this term, we consider two different settings.

MD1 (*Martingale Setting*). The sequence $(e_{\theta_n}(X_{n+1}))_{n \in \mathbb{N}}$ is a martingale difference sequence with respect to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$, i.e. $\mathbb{E}[e_{\theta_n}(X_{n+1})|\mathcal{F}_n] = 0$ for any $n \in \mathbb{N}^*$. Also, there exist $\sigma_0^2, \sigma_1^2 < +\infty$ such that for any $n \in \mathbb{N}^*$, $\mathbb{E}[\|e_{\theta_n}(X_{n+1})\|_{\theta_n}^2|\mathcal{F}_n] \leq \sigma_0^2 + \sigma_1^2 \|h(\theta_n)\|_{\theta_n}^2$.

Under **MD1**, we note that $\mathbb{E}[\sum_{k=0}^n \eta_{k+1} \Delta M_k] = 0$ in Lemma 2.3 and thus the following result can be easily obtained:

Theorem 2.4. Assume **A1-A2-A3-MD1**. Consider $(\theta_k)_{k \in \mathbb{N}}$ defined by (2.3). If $\sup_{k \in \mathbb{N}^*} \eta_k \leq \underline{c}_1/(6L(1 + \sigma_1^2))$, then for any $n \in \mathbb{N}^*$,

$$\begin{aligned} \mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] &\leq (\underline{c}_1 \Gamma_{n+1})^{-1} \left\{ 2\mathbb{E}[V(\theta_0)] + 3L(\sigma_0^2 + b_\infty^2) \Gamma_{n+1}^{(2)} \right\} \\ &\quad + 2(b_\infty \bar{c}/\underline{c}_1)^2 + 2\underline{c}_2/\underline{c}_1, \end{aligned}$$

where I_n has distribution defined by (2.5) and is independent of $(X_k)_{k \in \mathbb{N}^*}$.

Proof. The proof of Theorem 2.4 is postponed to Section 2.5.3. \square

We notice that the result is similar to (2.6) for a deterministic scheme, with additional constants related to the variance σ_0 and bias b_∞ .

To control the term $\Gamma_{n+1}^{(2)}$, it is desirable to select a nonincreasing step size. When the step size is $\eta_k = \square/(k + \Delta)^\alpha$, for some $\square, \Delta > 0$ and $\alpha \in (0, 1]$, Theorem 2.4 shows that $\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] \leq R(\square, \Delta, \alpha, n) + 2(b_\infty \bar{c}/\underline{c}_1)^2 + 2\underline{c}_2/\underline{c}_1$ where $R(\square, \Delta, \alpha, n) = \mathcal{O}(1/n^{\alpha \wedge (1-\alpha)})$ for $\alpha \in (0, 1/2) \cup (1/2, 1)$, $R(\square, \Delta, 1/2, n) = \mathcal{O}(\log(n)/\sqrt{n})$ and $R(\square, \Delta, 1, n) = \mathcal{O}(\log(n)^{-1})$.

We now turn to the Markovian setting.

MA1 (*Markovian Setting*). There exists a Markov kernel P on $(\Theta \times \mathsf{X}) \times \mathcal{X}$ such that for any $n \in \mathbb{N}$ and bounded and measurable function $\varphi : \mathsf{X} \rightarrow \mathbb{R}_+$, $\mathbb{E}[\varphi(X_{n+1})|\mathcal{F}_n] = \int_{\mathsf{X}} \varphi(y) P_{\theta_n}(X_n, dy)$. In addition, for any $\theta \in \Theta$, P_θ admits a unique invariant distribution π_θ satisfying $h(\theta) = \int_{\mathsf{X}} H_\theta(y) d\pi_\theta(y)$.

Consider a measurable function $w : \mathsf{X} \rightarrow [1, +\infty)$ and the following additional condition.

MA2 (w). The Markov kernel P and the measurable function w satisfy the following conditions.

- (a) There exists $e_\infty > 0$ such that for any $x \in \mathsf{X}$, $\sup_{\theta \in \Theta} \|e_\theta(x)\|_\theta \leq e_\infty w^{1/2}(x)$.
- (b) There exists $C_w \geq 1$ such that $\sup_{k \in \mathbb{N}} \mathbb{E}[w(X_{k+1})] \leq C_w$.
- (c) For any $\theta \in \Theta$, there exists a measurable function $\hat{e} : \Theta \times \mathsf{X} \rightarrow \mathbb{T}\Theta$ satisfying for any $x \in \mathsf{X}$, $\theta \in \Theta$, $\hat{e}_\theta(x) \in \mathbb{T}\theta$, and

$$\hat{e}_\theta(x) - \int_{\mathsf{X}} P_\theta(x, dy) \hat{e}_\theta(y) = e_\theta(x). \quad (2.9)$$

Moreover, there exists $\hat{e}_\infty > 0$ such that for any $\theta \in \Theta$, $x \in \mathsf{X}$ $\|\hat{e}_\theta(x)\|_\theta \leq \hat{e}_\infty w^{1/2}(x)$.

2.2. Geodesic Stochastic Approximation Schemes

(d) There exists $L_{\hat{\varepsilon}} \geq 0$ such that for any $\theta_0, \theta_1 \in \Theta$, $x \in \mathsf{X}$ and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0 and θ_1 ,

$$\left\| \int_{\mathsf{X}} P_{\theta_1}(x, dy) \hat{\varepsilon}_{\theta_1}(y) - \mathsf{T}_{01}^{\gamma} \left[\int_{\mathsf{X}} P_{\theta_0}(x, dy) \hat{\varepsilon}_{\theta_0}(y) \right] \right\|_{\theta_1} \leq L_{\hat{\varepsilon}} \ell(\gamma) w^{1/2}(x).$$

Note that (a) is a mild assumption on a Θ -uniform control (uniformly on the manifold) of the noise; (b) is automatically satisfied if w is bounded; (c) assumes the existence of solutions to the Poisson equation (2.9), which can be established if the Markov kernel P_{θ} is w -geometrically for any $\theta \in \Theta$. In addition, it is also required that $\hat{\varepsilon}_{\theta}$ (for fixed θ) to be uniformly bounded by $w^{1/2}$. Finally, assumption (d) is implied by smoothness conditions on the Markov kernel and the noise with respect to the SA parameter θ . We shall discuss the setting which implies **MA 2** just after the statement of our main result.

In the case w is not bounded, we need to consider the following condition which is automatically satisfied if Θ is compact, which is the case for the Grassmann manifold for example.

A 4. There exist $h_{\infty} \geq 0$ such that $\sup_{\theta \in \Theta} \|h(\theta)\|_{\theta} \leq h_{\infty}$.

We obtain the following result for geodesic SA scheme with Markovian noise:

Theorem 2.5. Assume **A 1-A 2-A 3-MA 1-MA 2**(w) hold for some measurable function $w : \mathsf{X} \rightarrow [1, +\infty)$. Assume in addition either w is bounded or **A 4**. Let $(\eta_k)_{k \in \mathbb{N}^*}$ be a sequence of stepsizes, $a_1, a_2 \geq 0$ satisfying $\sup_{k \in \mathbb{N}^*} \eta_k \leq c_1 / (4(3L/2 + D_{\hat{\varepsilon}}))$ and

$$\sup_{k \in \mathbb{N}^*} \{\eta_{k+1}/\eta_k\} \leq 1, \quad \sup_{k \in \mathbb{N}^*} \{\eta_k/\eta_{k+1}\} \leq a_1, \quad \sup_{k \in \mathbb{N}^*} \{|\eta_k - \eta_{k+1}|/\eta_k^2\} \leq a_2. \quad (2.10)$$

Consider $(\theta_k)_{k \in \mathbb{N}}$ defined by (2.3). Then for any $n \in \mathbb{N}^*$,

$$\mathbb{E} \left[\|\hat{h}(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq 2 \left\{ \mathbb{E} [V(\theta_0)] + C(\eta_1) + C_{\hat{\varepsilon}} \Gamma_{n+1}^{(2)} \right\} / (c_1 \Gamma_{n+1}) + 2(b_{\infty} \bar{c}/c_1)^2 + 2c_2/c_1.$$

where $I_n \in \{0, \dots, n\}$ is independent of \mathcal{F}_n and with distribution defined by (2.5) and the constants $C(\eta_1)$, $D_{\hat{\varepsilon}}$ and $C_{\hat{\varepsilon}}$ are given by (2.39), (2.40) and (2.41) in the appendix.

Proof. The proof of Theorem 2.5 is postponed to Section 2.5.5. \square

When the step size is $\eta_k = \square / (k + \Delta)^{\alpha}$, for some $\square, \Delta > 0$ and $\alpha \in (0, 1]$, the same conclusions of Theorem 2.4 can be drawn again.

We now give formal conditions that lead to **MA 2**(w). We gather them into two assumptions, the first one concerns the family of Markov kernels and the second one concerns the mean-field h and its observation H . A set $C \in \mathcal{X}$ is called small for a Markov kernel Q in $\mathsf{X} \times \mathcal{X}$ when there exists $m > 0$ and a non-trivial measure ν_m on \mathcal{X} such that for any $x \in C$, $B \in \mathcal{X}$, $Q^m(x, B) \geq \nu_m(B)$.

MA 3 (w). The family of Markov kernels $(P_{\theta})_{\theta \in \Theta}$ and the measurable function $w : \mathsf{X} \rightarrow [1, +\infty)$ satisfy one of the following conditions:

- (a) The family $(P_\theta)_{\theta \in \Theta}$ is uniformly ergodic with constant uniform in θ , i.e. there exist $\varepsilon_P \in (0, 1)$ and $C_P \geq 0$ such that for any $\theta \in \Theta$, P_θ has a unique stationary distribution π_θ and for any $x, x' \in \mathsf{X}$, and $k \in \mathbb{N}$, $\|\delta_x P_\theta^k - \delta_{x'} P_\theta^k\|_{\text{TV}} \leq C_P(1 - \varepsilon_P)^k$. In addition, there exists $C \geq 0$ such that for any $\theta_0, \theta_1 \in \Theta$, $x \in \mathsf{X}$ and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0 and θ_1 , $\|\delta_x P_{\theta_1} - \delta_x P_{\theta_0}\|_{\text{TV}} \leq C\ell(\gamma)$.
- (b) For any $\theta \in \Theta$, $P_\theta = P$, where P is a Markov kernel on $\mathsf{X} \times \mathcal{X}$. Moreover, there exists $w : \mathsf{X} \rightarrow [1, +\infty)$, $\lambda \in (0, 1)$, $b_w \geq 0$ such that for any $x \in \mathsf{X}$, $Pw(x) \leq \lambda w(x) + b_w \mathbb{1}_{\mathsf{C}}(x)$ where $\mathsf{C} \in \mathcal{X}$ is a small set for P .

Since **MA 3**(*w*)-(a) does not depend on the choice of w , we will simply denote it by **MA 3**-(a).

A 5 (*w*). There exists a constant $C \geq 0$ such that for any $x \in \mathsf{X}$, and any $\theta_0, \theta_1 \in \Theta$ and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0 and θ_1 ,

$$\sup_{\theta \in \Theta} \|e_\theta(x)\|_\theta \leq Cw^{1/2}(x), \quad \text{and} \quad \|e_{\theta_1}(x) - \mathsf{T}_{01}^\gamma e_{\theta_0}(x)\|_{\theta_1} \leq C\ell(\gamma)w^{1/2}(x). \quad (2.11)$$

Proposition 2.6. Assume **A 1**, **MA 1**. Firstly, if, in addition, **MA 3**(a) and **A 5**(1) hold, then **MA 2**(1) holds, where **1** is the constant function equal to one on X . Secondly, if, in addition, **MA 3**(*w*)-(b) and **A 5**(*w*) hold for a measurable function $w : \mathsf{X} \rightarrow [1, +\infty)$, then **MA 2**(*w*) holds.

Proof. The proof is postponed to Section 2.5.4. □

2.3 General Retraction SA Schemes

The previous section focused on the geodesic schemes (2.3) that require performing the Riemannian exponential map Exp at each iteration. Evaluating the Exp map is often computationally challenging. An alternative is to use a retraction map $\text{Ret} : \mathsf{T}\Theta \rightarrow \Theta$ that approximates the exponential. The main focus of this section is to analyze a retraction SA scheme:

$$\theta_{n+1} = \text{Ret}_{\theta_n} \{ \eta_{n+1} (H_{\theta_n}(X_{n+1}) + b_{\theta_n}(X_{n+1})) \}, \quad n \in \mathbb{N}, \quad (2.12)$$

given the initialization $\theta_0 \in \Theta$, and Ret_θ stands for the restriction of Ret to $\mathsf{T}_\theta\Theta$, $\theta \in \Theta$. We first discuss the properties of retraction operators with some examples below.

2.3.1 Retraction and Quantitative Estimates of the Retraction Error

For any $\theta \in \Theta$, we denote the cut locus of θ by $\text{Cut}(\theta) \subset \Theta$ (see Section 2.11.5). The following assumptions are made on the retraction maps of interest:

R1. (i) For any $\theta \in \Theta$, $\text{Ret}_\theta(0_\theta) = \theta$, where 0_θ is the zero element of $\mathsf{T}_\theta\Theta$ and $(\text{DRet})_\theta(0_\theta) = \text{Id}$.

2.3. General Retraction SA Schemes

(ii) For any $(\theta, u) \in T\Theta$, $\text{Ret}_\theta(u) \notin \text{Cut}(\theta)$.

(iii) Θ is a homogeneous Riemannian manifold (see Section 2.11.7) with isometry group \mathbf{G} and for any $g \in \mathbf{G}$, $(\theta, u) \in T\Theta$, $g \cdot \text{Ret}_\theta(u) = \text{Ret}_{g \cdot \theta}(g \cdot u)$, where $g \cdot \theta = g(\theta)$ and $g \cdot u = (Dg)_\theta(u)$ is a vector in $T_{g \cdot \theta}\Theta$.

R1-(ii) ensures that the inverse exponential map is defined on $\text{Ret}_\theta(T_\theta\Theta)$; see Section 2.11.6. Thus, the following function is well defined:

$$\Phi_\theta = \text{Exp}_\theta^{-1} \circ \text{Ret}_\theta : T_\theta\Theta \rightarrow T_\theta\Theta, \quad (2.13)$$

which defines a bundle map $\Phi : T\Theta \rightarrow T\Theta$. This mapping allows us to precisely measure the quality of the approximation of the exponential map by the retraction through quantitative bounds on the *retraction error map* $\Phi_\theta - \text{Id}_\theta$ on $T_\theta\Theta$. Note that in the special case when $\text{Ret} \equiv \text{Exp}$, we have $\Phi_\theta \equiv \text{Id}_\theta$, and the retraction error map is the null function.

R1-(i) roughly implies by a Taylor expansion that $\|\Phi_\theta(u) - u\|_\theta \leq \|u\|_\theta$, for $(\theta, u) \in T\Theta$. However as we shall see in the analysis, to ensure convergence of SA using retraction maps, we need to establish a tighter bound, *i.e.* $\|\Phi_\theta(u) - u\|_\theta \leq \|u\|_\theta^\beta$, for $(\theta, u) \in T\Theta$, $\beta \geq 2$. It boils down to showing that the first non-zero term in the Taylor expansion is of order higher than one. To obtain such a result, we consider the following assumptions, defining regular *first-order* **R2** and *second-order* **R3** retraction; see [AM12].

R2. For any $\theta \in \Theta$, there exists $\mathcal{L}^{(1)}(\theta) \geq 0$, such that $\|(D^2\Phi_\theta)_{tu}(u, u)\|_\theta \leq \mathcal{L}^{(1)}(\theta) \|u\|_\theta^2$, for any $u \in T_\theta\Theta$ and $t \in [0, 1]$, where the function $\Phi_\theta : T_\theta\Theta \rightarrow T_\theta\Theta$ is defined by (2.13).

R3. For any $\theta \in \Theta$, the following hold.

(i) Setting $\gamma(t) = \text{Ret}_\theta(tu)$ for $t \in \mathbb{R}$ and $u \in T_\theta\Theta$, the initial acceleration of the curve γ satisfies $D_t\dot{\gamma}(0) = 0_\theta$, where D_t is the covariant derivative along γ (see [Lee19, Theorem 4.24]).

(ii) There exists $\mathcal{L}^{(2)}(\theta) \geq 0$, such that $\|(D^3\Phi_\theta)_{tu}(u, u, u)\|_\theta \leq \mathcal{L}^{(2)}(\theta) \|u\|_\theta^3$, for any $u \in T_\theta\Theta$ and $t \in [0, 1]$, where the function $\Phi_\theta : T_\theta\Theta \rightarrow T_\theta\Theta$ is defined by (2.13).

Note **R3** does not imply **R2** nor vice versa. The following is a consequence of **R2** or **R3**:

Lemma 2.7. Assume **R1** and let $\tilde{\theta} \in \Theta$, it holds

(a) Under **R2**, $\|\Phi_\theta(u) - u\|_\theta \leq \mathcal{L}^{(1)}(\tilde{\theta}) \|u\|_\theta^2$ for any $(\theta, u) \in T\Theta$.

(b) Under **R3**, $\|\Phi_\theta(u) - u\|_\theta \leq \mathcal{L}^{(2)}(\tilde{\theta}) \|u\|_\theta^3$ for any $(\theta, u) \in T\Theta$.

Proof. The proof is postponed to Section 2.6.1. □

We conclude this subsection by illustrating two examples of matrix manifolds with retraction operators satisfying **R1-R2-R3**.

Example 2.8. (*d-dimensional sphere* \mathbb{S}^d) For $\Theta = \mathbb{S}^d = \{\mathbf{x} \in \mathbb{R}^{d+1} : \|\mathbf{x}\| = 1\}$, we may take $(\theta, \mathbf{u}) \in T\mathbb{S}^d \mapsto (\theta + \mathbf{u}) / \|\theta + \mathbf{u}\|$ as a retraction map. We show in Section 2.6.2 that **R1-R2-R3** are satisfied for this example.

Example 2.9. (*Grassmannian manifold* $\text{Gr}_r(\mathbb{R}^d)$) For $r \in \{1, \dots, d-1\}$, let $\Theta = \text{Gr}_r(\mathbb{R}^d)$ be the set of r -dimensional subspaces over \mathbb{R}^d . Note that $\text{Gr}_r(\mathbb{R}^d)$ is a compact $(d-r)r$ -dimensional manifold [Lee03, Example 21.21]. We represent it as the quotient of the non-compact Stiefel manifold by the action of the linear group $\text{GL}_r(\mathbb{R})$, see Section 2.7.2 for a full description of the manifold and its tangent space. We consider the canonical metric (see [AMS04, Section 3.3]) defined for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$, and $\mathbf{D}_1, \mathbf{D}_2 \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, by $g_{[\mathbf{B}]}^{\text{Gr}}(\mathbf{D}_1, \mathbf{D}_2) = \text{Tr}((\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}_1^\top \mathbf{D}_2)$. The exponential map on $\text{Gr}_r(\mathbb{R}^d)$ corresponding to this metric is given for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$, $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ by

$$\text{Exp}_{[\mathbf{B}]}(\mathbf{D}) = \left[\left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \quad \mathbf{U} \right) \begin{pmatrix} \cos(\boldsymbol{\Sigma}) \\ \sin(\boldsymbol{\Sigma}) \end{pmatrix} \right], \quad (2.14)$$

where $\exp(\cdot)$ is the matrix exponential and $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top$ is the compact singular value decomposition (denoted SVD in the future, see [HP14, Theorem 1.46]), with \mathbf{U} a $d \times r$ matrix with orthonormal columns, meaning $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$, \mathbf{V} a $r \times r$ orthogonal matrix, and $\boldsymbol{\Sigma}$ a $r \times r$ diagonal matrix with nonnegative elements. See [AMS04, Theorem 3.6] for the proof. A retraction map is then defined for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$, $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ by

$$\text{Ret}_{[\mathbf{B}]}(\mathbf{D}) = [\mathbf{B} + \mathbf{D}]. \quad (2.15)$$

This retraction operator is very natural from the geometrical point of view, as it is a formulation of the parametrization of $\text{Gr}_r(\mathbb{R}^d)$ in a neighborhood of $[\mathbf{B}]$ (see Proposition 2.22). It has also been considered in [AM12; Bon13] which discuss in details the computational benefits of this approximation. In practice, to simplify the formulations of geodesics and metric computations, we always consider orthonormal representatives for $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ and $[\mathbf{B} + \mathbf{D}]$, $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ by performing a Gram–Schmidt orthonormalization step or by choosing \mathbf{Q} in a QR decomposition $\mathbf{B} + \mathbf{D} = \mathbf{Q}\mathbf{R}$.

Proposition 2.10. *The projective retraction on the Grassmannian manifold Ret defined by (2.15) satisfies R.1–R.2–R.3.*

Proof. The proof is postponed to Section 2.6.3. \square

Example 2.9 is important in our application of retraction SA to PCA problems to be discussed in Section 2.4.1. Indeed, the retraction SA scheme that we consider with the retraction map (2.15) corresponds to the famous Oja’s algorithm.

2.3.2 Analysis of Retraction SA Schemes

To analyze the retraction SA scheme in (2.12), we observe that it can be re-written as:

$$\theta_{n+1} = \text{Exp}_{\theta_n} \left\{ \eta_{n+1} (H_{\theta_n}(X_{n+1}) + b_{\theta_n}(X_{n+1}) + \Delta_{\theta_n, \eta_{n+1}}(X_{n+1})) \right\}, \quad (2.16)$$

where $\Delta_{\theta, \eta}(x)$ is the ‘retraction bias’ defined for any $(\theta, x) \in \Theta \times \mathbf{X}$ and $\eta > 0$ by

$$\Delta_{\theta, \eta}(x) = \eta^{-1} \Phi_\theta(\eta \{H_\theta(x) + b_\theta(x)\}) - \{H_\theta(x) + b_\theta(x)\}. \quad (2.17)$$

Our strategy for analyzing (2.12) is to incorporate the retraction bias in the analysis as it is done for the geodesic SA scheme, using Lemma 2.7 and noting that $\Delta_{\theta, \eta}(x) = \eta^{-1} \{ \Phi_\theta(u) - u \}$ with $u = \eta \{H_\theta(x) + b_\theta(x)\}$ for any $\theta \in \Theta$, $x \in \mathbf{X}$ and $\eta > 0$. Again, we consider separately the martingale and Markovian settings.

2.4. Applications

Martingale setting We shall strengthen this assumption first and consider for some $a > 0$:

MD 2 (a). *There exists $m_{(a)} \geq 0$ such that a.s. $\mathbb{E}[\|e_{\theta_n}(X_{n+1})\|_{\theta_n}^a | \mathcal{F}_n] \leq m_{(a)}$, for any $n \in \mathbb{N}$.*

We acquire the following convergence result for $(\theta_n)_{n \in \mathbb{N}}$.

Theorem 2.11. *Assume **A 1-A 2-A 3-A 4**, **R 1**, **MD 1** hold. Suppose either **R 2**, **MD 2(4)** or **R 3**, **MD 2(6)**. Consider $(\theta_k)_{k \in \mathbb{N}}$ defined by (2.12). Then, if $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta}$, for $\bar{\eta}$ defined in (2.99), for any $n \in \mathbb{N}$*

$$\mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq 2(\underline{c}_1 \Gamma_{n+1})^{-1} \{ \mathbb{E}[V(\theta_0)] + C_R^{(2)} \Gamma_{n+1}^{(2)} + A_{n+1}^{(R)} \} + 2(b_\infty \bar{c} h_\infty + \underline{c}_2) / \underline{c}_1,$$

where $C_R^{(2)}$ is given by (2.100), $A_{n+1}^{(R)} = \mathcal{O}(\Gamma_{n+1}^{(3)})$ is given by (2.104) and $I_n \in \{0, \dots, n\}$ is a random variable independent of \mathcal{F}_n and distributed according to (2.5).

Proof. The proof is postponed to Section 2.8.2. □

Theorem 2.11 shows that the bounds for the retraction scheme (2.16) and the geodesic scheme (2.3) are (up to the precise definitions of constants) the same.

Markovian Setting : We strengthen **MA 2(w)-(b)** as follows:

MA 4 (w). *The Markov chain $(X_k)_{k \in \mathbb{N}^*}$ satisfies that $\sup_{k \in \mathbb{N}} \mathbb{E}[w^3(X_{k+1})] \leq C_w^{(3)}$ for $C_w^{(3)} \geq 0$.*

Theorem 2.12. *Assume **A 1-A 2-A 3-A 4**, **R 1**, **MA 1-MA 2(w)-MA 4(w)** hold for some measurable function $w : X \rightarrow [1, +\infty)$. Suppose in addition that **R 2** or **R 3** holds. Assume that $(\eta_k)_{k \in \mathbb{N}^*}$ is a sequence of stepsizes and $a_1, a_2 \geq 0$ satisfying (2.10). Consider $(\theta_k)_{k \in \mathbb{N}}$ defined by (2.12). Then, if $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta}$, for $\bar{\eta}$ defined in (2.106), for any $n \in \mathbb{N}$*

$$\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] \leq 2(\underline{c}_1 \Gamma_{n+1})^{-1} \{ \mathbb{E}[V(\theta_0)] + C(\eta_1) + D_R^{(2)} \Gamma_{n+1}^{(2)} + B_{n+1}^{(R)} \} + 2(b_\infty \bar{c} h_\infty + \underline{c}_2) / \underline{c}_1,$$

$C(\eta_1), D_R^{(2)}$ are given by (2.39)-(2.107), $B_{n+1}^{(R)} = \mathcal{O}(\Gamma_{n+1}^{(3)})$ by (2.112) and lastly $I_n \in \{0, \dots, n\}$ is a random variable independent of \mathcal{F}_n and distributed according to (2.5).

Proof. The proof is postponed to Section 2.8.3. □

Again, the bounds are matched with those obtained in the martingale setting.

2.4 Applications

In this section, we illustrate our convergence analysis results on two examples: subspace tracking method and robust barycenter problem.

2.4.1 Principal Component Analysis

We consider an online principal component analysis (PCA) problem in which we aim at estimating the r principal eigenvectors of a $d \times d$ covariance matrix \mathbf{A} ; see [CD18] and the references therein. We assume that we have access to noisy data $(X_n)_{n \in \mathbb{N}^*}$ in \mathbb{R}^d which have covariance matrices close in some sense to \mathbf{A} . More precisely, we consider the following two cases.

HP1. $(X_n)_{n \in \mathbb{N}^*}$ is i.i.d., $\mathbb{E}[X_1 X_1^\top] = \mathbf{A}$, and there exists $m_{\text{pca}}^{(12)} < \infty$ such that $\mathbb{E}[\|X_1\|^{12}] \leq m_{\text{pca}}^{(12)}$.

HP2. $(X_n)_{n \in \mathbb{N}^*}$ is a Markov chain on $\mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d)$ with Markov kernel P and stationary distribution π such that $\int_{\mathbb{R}^d} x d\pi(x) = 0$ and $\mathbf{A} = \int_{\mathbb{R}^d} x x^\top d\pi(x)$. In addition, there exists $w : \mathcal{X} \rightarrow [1, +\infty)$, $\lambda_w \in (0, 1)$, $b_w \geq 0$ such that for any $x \in \mathcal{X}$, $Pw(x) \leq \lambda_w w(x) + b_w \mathbb{1}_{\mathcal{C}}(x)$ where $\mathcal{C} \in \mathcal{B}(\mathbb{R}^d)$ is a small set for P . Finally, w satisfies $\sup_{x \in \mathbb{R}^d} \{\|x\|^{12}/w(x)\} < +\infty$.

For $\mathbf{B} \in \text{St}_r(\mathbb{R}^d)$ and $\mathbf{B}_\perp \in \text{St}_{d-r}(\mathbb{R}^d)$ such that $\mathbf{B}^\top \mathbf{B}_\perp = \mathbf{0}_{r \times (d-r)}$, consider $\text{Proj}_{[\mathbf{B}]} = \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top$ and $\text{Proj}_{[\mathbf{B}_\perp]} = \mathbf{B}_\perp(\mathbf{B}_\perp^\top \mathbf{B}_\perp)^{-1} \mathbf{B}_\perp^\top = \mathbf{I}_d - \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top$ the orthogonal projection matrices on $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ and $[\mathbf{B}_\perp] \in \text{Gr}_{d-r}(\mathbb{R}^d)$ respectively. Online PCA in the i.i.d. case has been tackled by, among others, e.g. [Bon13; Oja92; ZS16], [CD18, Section 5]. We consider two algorithms introduced in [Bon13, Section 4.1] to estimate \mathbf{A} . In both cases, PCA is considered as a stochastic minimization problem on the Grassmann manifold (see Example 2.9) of r -dimensional linear subspaces of \mathbb{R}^d , $\text{Gr}_r(\mathbb{R}^d)$, $r \in \{1, \dots, d-1\}$ with objective function

$$f([\mathbf{B}]) = -\text{Tr}(\mathbf{B}^\top \mathbf{A} \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1}) / 2. \quad (2.18)$$

Since $\text{Proj}_{[\mathbf{B}_\perp]}$ does not depend on the representative $\mathbf{B} \in [\mathbf{B}]$, f defines a map on $\text{Gr}_r(\mathbb{R}^d)$. Then, by [AMS04, Section 3.4], the gradient of f is given by

$$\text{grad} f([\mathbf{B}]) = -(\mathbf{I}_d - \text{Proj}_{[\mathbf{B}]}) \mathbf{A} \mathbf{B}, \quad \text{for any } [\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d),$$

Based on the exponential map (2.14) or the retraction (2.15), we consider the two recursions

$$[\mathbf{B}_{n+1}^{(E)}] = \text{Exp}_{[\mathbf{B}_n^{(E)}]}(\eta_{n+1} \{\mathbf{I}_d - \text{Proj}_{[\mathbf{B}_n^{(E)}]}\} \{X_{n+1} X_{n+1}^\top\} \mathbf{B}_n^{(E)}), \quad (2.19)$$

$$[\mathbf{B}_{n+1}^{(Re)}] = [\mathbf{B}_n^{(Re)} + \eta_{n+1} \{\mathbf{I}_d - \text{Proj}_{[\mathbf{B}_n^{(Re)}]}\} \{X_{n+1} X_{n+1}^\top\} \mathbf{B}_n^{(Re)}]. \quad (2.20)$$

Recall that by Proposition 2.10, the retraction used for (2.20) satisfies **R1-R2-R3**.

Theorem 2.13. *The schemes defined by (2.19) and (2.20), on the manifold $\Theta = \text{Gr}_r(\mathbb{R}^d)$ is of the form (2.2) with $H_{[\mathbf{B}]}(x) = \{\mathbf{I}_d - \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top\} \{x x^\top\} \mathbf{B}$, and **A1-A2-A3-A4** are satisfied with the mean-field $h \leftarrow -\text{grad} f$, Lyapunov function $V \leftarrow f$ and $\underline{c}_2 = b_\infty = 0$.*

(a) Under **HP1**, (2.19) and (2.20) satisfy **MD1-MD2(6)**. Thus, there exist some $C > 0$ and $\bar{\eta} > 0$ such that if $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta}$, then for any $n \in \mathbb{N}^*$,

$$\mathbb{E} \left[\|\text{grad} f(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq \Gamma_{n+1}^{-1} \left\{ 2\mathbb{E}[f(\theta_0)] + C\Gamma_{n+1}^{(2)} \right\}, \quad \text{when using (2.19)},$$

$$\mathbb{E} \left[\|\text{grad} f(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq \Gamma_{n+1}^{-1} \left\{ 2\mathbb{E}[f(\theta_0)] + C \left(\Gamma_{n+1}^{(2)} + \Gamma_{n+1}^{(3)} \right) \right\}, \quad \text{when using (2.20)}.$$

2.4. Applications

Here $I_n \in \{0, \dots, n\}$ is a random variable independent from $(X_k)_{k \in \mathbb{N}^*}$ and θ_0 , with distribution defined by (2.5). Moreover, the squared norm of the gradient can be expressed as $\|\text{grad } f(\theta_{I_n})\|_{\theta_{I_n}}^2 = \text{Tr}(\text{Proj}_{[\mathbf{B}_{I_n}]} \mathbf{A} \text{Proj}_{([\mathbf{B}_{I_n}]_{\perp})} \mathbf{A})$, where, for any $k \in \mathbb{N}$, $\mathbf{B}_k \in \text{St}_r(\mathbb{R}^d)$ and $(\mathbf{B}_k)_{\perp} \in \text{St}_{d-r}(\mathbb{R}^d)$ are such that $[\mathbf{B}_k] = \theta_k$ and $\mathbf{B}_k^{\top} (\mathbf{B}_k)_{\perp} = 0$.

(b) Consider **HP 2** with its function $w : \mathbf{X} \rightarrow [1, +\infty)$, (2.19) and (2.20) satisfy **MA 1-MA 2(w)-MA 4(w)**. Then, the same conclusion as (a) follows.

Proof. Note that **A 1** is satisfied following for example [EAS98, Section 2.5]. Setting $h \leftarrow -\text{grad } f$ and $V \leftarrow -f$, since f is smooth and $\text{Gr}_r(\mathbb{R}^d)$ is compact, it is easy to verify that **A 2** is satisfied with $c_2 = 0$.

In light of the above recursion, the noise vector field can be written as $e_{[\mathbf{B}]}(x) = \{\mathbf{I}_d - \mathbf{B}(\mathbf{B}^{\top} \mathbf{B})^{-1} \mathbf{B}^{\top}\}(xx^{\top} - \mathbf{A})\mathbf{B}$ and the bias function is the null function $b_{[\mathbf{B}]}(x) = \mathbf{0}_d$, for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$, $x \in \mathbb{R}^d$. We recall from Example 2.9 that the retraction operator in (2.20) satisfies **R 1-R 2-R 3**. We now verify the other assumptions to apply Theorems 2.4, 2.5, 2.11 and 2.12.

- Under **HP 1**, **MD 1**, **MD 2(6)** hold.
- Under **HP 2**, **MA 1** is automatically satisfied. We show in Proposition 2.32 that for any $x \in \mathbb{R}^d$, $[\mathbf{B}] \mapsto e_{[\mathbf{B}]}(x)$ satisfies (2.11). Finally, **HP 2** implies that for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ and $x \in \mathbb{R}^d$, $\|e_{[\mathbf{B}]}(x)\|_{[\mathbf{B}]} \leq Cw^{1/2}(x)^1$, for some constant $C \geq 0$. Therefore, **HP 2** and an application of Proposition 2.6 shows that **MA 2(w)** is satisfied. Also, the explicit expression for the noise e and **HP 2** implies that **MA 4** holds. \square

Theorem 2.13 shows that the schemes (2.19) and (2.20) converge to a zero of $\text{grad } f$ with convergence bounds provided by our main results since $c_2 = b_{\infty} = 0$.

2.4.2 Robust Barycenter in a Hadamard Manifold

Let Θ be a Hadamard manifold – a simply-connected, complete Riemannian manifold of non-positive sectional curvature [Lee19, Propositions 8.29, 12.9]. We assume that the sectional curvature of Θ is bounded below, $-\kappa^2 \leq \text{sec } \Theta \leq 0$. A common example of this situation is $\Theta = \text{S}_d^{++}(\mathbb{R})$, the space of real $d \times d$ symmetric positive-definite matrices, equipped with its affine-invariant metric [PFA06]. Consider data points $(X_n)_{n \in \mathbb{N}^*}$ lying on the Riemannian manifold Θ , i.e., $X_n \in \Theta$, that are drawn from a distribution π . A fundamental machine learning problem is to compute some kind of central value of π , defined as an optimal solution to

$$\min_{\Theta} f, \quad \text{where } f(\theta) = \int_{\Theta} \tilde{\rho}(\theta, x) \pi(dx) \text{ for } \theta \in \Theta, \quad (2.21)$$

where $\tilde{\rho}(\theta, x)$ is some Riemannian dissimilarity measure. For instance, the Riemannian barycenter, also called the Fréchet mean, is obtained by taking $\tilde{\rho}(\theta, x) = \rho_{\Theta}^2(\theta, x)$ where $\rho_{\Theta} : \Theta \times \Theta \rightarrow \mathbb{R}_+$ is the Riemannian distance of Θ [GV08; Mat+19].

The Riemannian barycenter is known to be sensitive to outliers, which motivated the idea of considering the Riemannian median, obtained by taking $\tilde{\rho}(\theta, x) = \rho_{\Theta}(\theta, x)$ [ABY13].

¹ $\|e_{[\mathbf{B}]}(x)\|_{[\mathbf{B}]}$ is the norm the tangent vector $e_{[\mathbf{B}]}(x)$ associated with the Riemannian metric of $\text{Gr}_r(\mathbb{R}^d)$ given in Example 2.9 and $\|x\|$ is the Euclidean norm of x .

We consider a *robust barycenter* by using a Huber-like dissimilarity measure, $\rho_H : \Theta \times \Theta \rightarrow \mathbb{R}_+$, defined for any $\theta_0, \theta_1 \in \Theta$ by $\rho_H(\theta_0, \theta_1) = \delta^2 [1 + \{\rho_\Theta(\theta_0, \theta_1)/\delta\}^2]^{1/2} - \delta^2$, where $\delta > 0$ is a cut-off constant. Observe that $\rho_H(\theta_0, \theta_1)$ behaves like $(1/2)\rho_\Theta^2(\theta_0, \theta_1)$ when $\rho_\Theta(\theta_0, \theta_1)$ is small compared to δ , and like $\delta\rho_\Theta(\theta_0, \theta_1)$ when $\rho_\Theta(\theta_0, \theta_1)$ is large compared to δ . Using ρ_H in the optimization problem (2.21) yields a *robust barycenter* problem, and the robust barycenter is a global minimizer of (2.21). We say that a continuously differentiable function $f : \Theta \rightarrow \mathbb{R}$ is *strictly geodesically convex*, or *strictly g-convex* in short, when for any $\theta_0, \theta_1 \in \Theta$ with $\theta_0 \neq \theta_1$, $f(\theta_1) > f(\theta_0) + \langle \text{Exp}_{\theta_0}^{-1}(\theta_1), \text{grad } f(\theta_0) \rangle_{\theta_0}$.

Proposition 2.14. *Let Θ be a Hadamard manifold with sectional curvature bounded from below, and let π be a probability distribution on Θ . Assume there exists some $\tau \in \Theta$ such that*

$$\int_{\Theta} \rho_\Theta(\tau, x) \pi(dx) < +\infty. \quad (2.22)$$

Then, the function $f : \Theta \rightarrow \mathbb{R}_+$ given by (2.21) taking the Huber-like dissimilarity $\rho_H : (\theta_0, \theta_1) \mapsto \delta^2 [1 + \{\rho_\Theta(\theta_0, \theta_1)/\delta\}^2]^{1/2} - \delta^2$ is geodesically strictly convex. Moreover, f has a unique global minimizer $\theta^ \in \Theta$. In other words, π has a unique robust barycenter θ^* .*

Proof. The proof is postponed to Section 2.9.2. □

In the case where $(X_n)_{n \in \mathbb{N}^*}$ is either an i.i.d. sequence or is given by an ergodic Markov chain with stationary distribution π satisfying **MA 2**, tackling the problem (2.21) can be done by considering the following geodesic SA scheme:

$$\theta_{n+1} = \text{Exp}_{\theta_n}(\eta_{n+1} \text{Exp}_{\theta_n}^{-1}(X_{n+1}) / [1 + \{\rho_\Theta(\theta_n, X_{n+1})/\delta\}^2]^{1/2}). \quad (2.23)$$

The above is the “recursive barycenter” scheme proposed by [Arn+12; Stu03], except that a move in the direction of a new observation X_{n+1} is attenuated when this new observation lies too far from the current estimate θ_n . This means that less weight is assigned to extreme observations. Note that this strategy and the scheme (2.23) have been also considered in [Cha+20] to solve PCA in a robust manner.

Proposition 2.15. *Let Θ be a Hadamard manifold with sectional curvature bounded below by $-\kappa^2$. Furthermore, consider f given by (2.21) as the Lyapunov function for the stochastic approximation scheme (2.23). Then **A 2**, **A 3** and **A 5(1)** are satisfied, with $\underline{c} = \bar{c} = 1$, $L = 1 + \delta\kappa$, and $b_\infty = 0$. Moreover, if the family $(X_n)_{n \in \mathbb{N}^*}$ is i.i.d. with distribution π , then **MD 1** is satisfied, with $\sigma_0 = \delta^2$ and $\sigma_1 = 0$.*

Proof. The proof is postponed to Section 2.9.2 □

This means that in the i.i.d. setting, Theorem 2.4 can be used, and in the Markovian case, we only require **MA 1** and **MA 3(w)-(b)** for some measurable function $w : \Theta \rightarrow [1, +\infty)$, to satisfy the assumptions of Theorem 2.5. Furthermore, (2.23) finds a unique and global minimizer to (2.21).

2.4. Applications

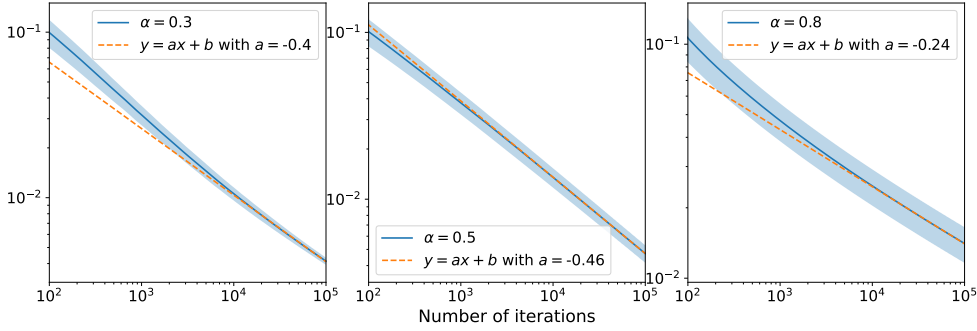


Figure 2.1: $\{\mathbb{E}[\|\text{grad } f(\mathbf{B}_{I_n}^{(Re)})\|_{[\mathbf{B}_n]}\}]_n$ for $\eta_n = \square/(n + \Delta)^\alpha$ in the i.i.d. setting

2.4.3 Numerical Examples on the PCA Problem

We illustrate our results on a PCA problem on $\text{Gr}_r(\mathbb{R}^d)$ with $r = 8$ and $d = 50$ using the two settings **HP 1** and **HP 2**. We use both the retraction map (2.15) and the recursion (2.20).

We first consider the case where $(X_n)_{n \in \mathbb{N}^*}$ is i.i.d. with zero-mean Gaussian distribution and covariance matrix \mathbf{A} . For our experiments, \mathbf{A} is a randomly sampled symmetric matrix with eigenvalues in $[10^{-6}, 10]$. Note that **HP 1** holds. In Figure 2.1 we illustrate the convergence bounds provided by Theorem 2.11. This figure shows a Monte Carlo estimation of $\{\mathbb{E}[\|\text{grad } f(\mathbf{B}_{I_n}^{(Re)})\|_{[\mathbf{B}_n]}\}] : n \in \{1, \dots, 10^5\}\}$ for the step-size sequences $\eta_n = \square/(n + \Delta)^\alpha$, $\alpha \in \{0.3, 0.5, 0.8\}$. We can observe that the convergence bounds provided by Theorem 2.11 which are of order $\mathcal{O}(1/n^{\alpha \wedge (1-\alpha)})$ for $\alpha \in (0, 1/2) \cup (1/2, 1)$, and $\mathcal{O}(\log(n)/\sqrt{n})$ are met.

In a second experiment, we assume that $(X_n)_{n \in \mathbb{N}^*}$ is a Markov chain, with state space \mathbb{R}^d , satisfying the recursion for any $k \in \mathbb{N}$, $X_{k+1} = \mathbf{M}X_k + \varepsilon_{k+1}$, where $\mathbf{M} \in \mathbb{R}^{d \times d}$ and $(\varepsilon_k)_{k \in \mathbb{N}^*}$ is a sequence of i.i.d. random variables with zero-mean Gaussian distribution and covariance matrix \mathbf{A} . The matrix \mathbf{M} is chosen so that its spectral radius is strictly less than 1, which implies that **HP 2** is satisfied; see Section 2.9.3. The stationary distribution is zero-mean Gaussian with covariance matrix satisfying the discrete Riccati equation: $\Sigma_\pi = \mathbf{M}\Sigma_\pi\mathbf{M}^\top + \mathbf{A}$. Similarly to the i.i.d. setting and with the same parameter choices, the convergence bounds provided by Theorem 2.12 are illustrated in Figure 2.2.




Figure 2.2: $\{\mathbb{E}[\|\text{grad } f(\mathbf{B}_{L_n}^{(Re)})\|_{[\mathbf{B}_n]}\}]_n$ for $\eta_n = \square/(n + \triangle)^\alpha$ in the Markovian setting

2.5 Proofs for Section 2.2

2.5.1 Proof of Lemma 2.1

The proof uses results on parallel transport and geodesics on a Riemannian manifold. For reader convenience, these concepts are recalled in Section 2.11.1 and 2.11.4.

Using $\gamma(0) = \theta_0$, $\gamma(1) = \theta_1$, by a Taylor expansion of $V \circ \gamma$ and the definition of the Riemannian gradient (see Section 2.11.8), we have

$$V(\theta_1) - V(\theta_0) = \int_0^1 \langle \text{grad } V(\gamma(t)), \dot{\gamma}(t) \rangle_{\gamma(t)} dt = \int_0^1 \langle \text{grad } V(\gamma(t)), T_{0t}^\gamma \dot{\gamma}(0) \rangle_{\gamma(t)} dt,$$

where we have used for the last equality the uniqueness of the parallel transport [Lee19, Theorem 4.32] and that γ is a geodesic. Therefore, we obtain, using that the parallel transport is a linear isometry [Lee19, Proposition 5.5], that

$$\begin{aligned} & \left| V(\theta_1) - V(\theta_0) - \langle \text{grad } V(\theta_0), \dot{\gamma}(0) \rangle_{\gamma(0)} \right| \\ & \leq \int_0^1 \left| \langle \text{grad } V(\gamma(t)), T_{0t}^\gamma \dot{\gamma}(0) \rangle_{\gamma(t)} - \langle \text{grad } V(\gamma(0)), \dot{\gamma}(0) \rangle_{\gamma(0)} \right| dt, \\ & = \int_0^1 \left| \langle \text{grad } V(\gamma(t)) - T_{0t}^\gamma \text{grad } V(\gamma(0)), T_{0t}^\gamma \dot{\gamma}(0) \rangle_{\gamma(t)} \right| dt, \\ & \leq L \int_0^1 \ell(\gamma|_{[0,t]}) \|T_{0t}^\gamma \dot{\gamma}(0)\|_{\gamma(t)} dt \leq L \ell(\gamma) \int_0^1 \ell(\gamma) t dt, \end{aligned}$$

where we have used that since γ is a geodesic $T_{0t}^\gamma \dot{\gamma}(0) = \dot{\gamma}(t)$ and by [Lee19, Corollary 5.6], $\|\dot{\gamma}(t)\|_{\gamma(t)} = \ell(\gamma)$, and $\ell(\gamma|_{[0,t]}) = t\ell(\gamma)$ by [Lee19, Lemma 5.18].

2.5.2 A technical lemma: proof of Lemma 2.3

For any $k \geq 0$ and $t \in [0, 1]$, consider $\gamma^{(k)}(t) = \text{Exp}_{\theta_k} \{t\eta_{k+1}(H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}))\}$. Note that $\dot{\gamma}^{(k)}(0) = \eta_{k+1}\{H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\}$ and $\ell(\gamma^{(k)}) = \eta_{k+1} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}$.

2.5. Proofs for Section 2.2

Then, by Lemma 2.1, (2.8) and using that for any $\theta \in \Theta$, $a, b, c \in \mathbb{T}_\theta \Theta$, $\|a + b + c\|_\theta^2 \leq 3(\|a\|_\theta^2 + \|b\|_\theta^2 + \|c\|_\theta^2)$, we get that for any $k \geq 0$,

$$\begin{aligned} \left| V(\theta_{k+1}) - V(\theta_k) - \eta_{k+1} \langle \text{grad } V(\theta_k), H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \right| &\leq (L/2)\ell(\gamma^{(k)})^2, \\ &= (L\eta_{k+1}^2/2) \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^2, \\ &\leq (3L/2)\eta_{k+1}^2 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \right\}. \end{aligned}$$

Therefore, we get that for any $k \in \mathbb{N}$,

$$\begin{aligned} & - \eta_{k+1} \langle \text{grad } V(\theta_k), h(\theta_k) \rangle_{\theta_k} \\ & \leq V(\theta_k) - V(\theta_{k+1}) + \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \quad (2.24) \\ & \quad + (3L/2)\eta_{k+1}^2 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \right\}. \end{aligned}$$

By **A2-(a)** and the Cauchy-Schwarz inequality, we obtain for any $k \in \mathbb{N}^*$ and $\varepsilon > 0$,

$$\langle \text{grad } V(\theta_k), b_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \leq \bar{c}^2 \varepsilon \|h(\theta_k)\|_{\theta_k}^2 + (1/4\varepsilon) \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2.$$

Thus, plugging this in (2.24) and using the first inequality in **A2-(a)** gives,

$$\begin{aligned} & \eta_{k+1}(\underline{c}_1 - (3L/2)\eta_{k+1} - \bar{c}^2 \varepsilon) \|h(\theta_k)\|_{\theta_k}^2 \\ & \leq V(\theta_k) - V(\theta_{k+1}) + \eta_{k+1} \underline{c}_2 + \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \\ & \quad + [\eta_{k+1}/4\varepsilon + (3L/2)\eta_{k+1}^2] \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + (3L/2)\eta_{k+1}^2 \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2. \end{aligned}$$

Adding these inequalities from 0 to n and rearranging terms concludes the proof.

2.5.3 Proof of Theorem 2.4

Let $n \in \mathbb{N}^*$. First note that for any $k \in \mathbb{N}$, $\mathbb{E}[\Delta M_k] = 0$, using that θ_k is \mathcal{F}_k -measurable and **MD1**. Therefore, taking the expectation in the inequality given by Lemma 2.3 and $\|b_{\theta_k}(X_{k+1})\|_{\theta_k} \leq b_\infty$ using **A3**, we obtain

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1}(\underline{c}_1 - (3L/2)\eta_{k+1} - \bar{c}^2 \varepsilon) \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] &\leq \mathbb{E}[V(\theta_0)] + \underline{c}_2 \Gamma_{n+1} \\ & \quad + (3L/2) \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E} \left[\|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \right] + b_\infty^2 \sum_{k=0}^n \eta_{k+1} \left\{ (4\varepsilon)^{-1} + (3L/2)\eta_{k+1} \right\}. \end{aligned}$$

Since for any $k \in \mathbb{N}$, $\mathbb{E}[\|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2] \leq \sigma_0^2 + \sigma_1^2 \|h(\theta_k)\|_{\theta_k}^2$ using that θ_k is \mathcal{F}_k -measurable and **MD1**, we get

$$\begin{aligned} & \sum_{k=0}^n \eta_{k+1}(\underline{c}_1 - \bar{c}^2 \varepsilon - (3L/2)\{1 + \sigma_1^2\}\eta_{k+1}) \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\ & \leq \mathbb{E}[V(\theta_0)] + \underline{c}_2 \Gamma_{n+1} + 3L\sigma_0^2 \Gamma_{n+1}^{(2)}/2 + b_\infty^2 \{\Gamma_{n+1}/(4\varepsilon) + 3L\Gamma_{n+1}^{(2)}/2\}. \end{aligned}$$

Taking $\varepsilon = \underline{c}_1/(4\bar{c}^2)$ and dividing by Γ_{n+1} , we get

$$\begin{aligned} & \Gamma_{n+1}^{-1} \sum_{k=0}^n \eta_{k+1} (3\underline{c}_1/4 - (3L/2)\{1 + \sigma_1^2\}\eta_{k+1}) \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\ & \leq \mathbb{E} [V(\theta_0)] / \Gamma_{n+1} + 3L\Gamma_{n+1}^{(2)} \{\sigma_0^2 + b_\infty^2\} / (2\Gamma_{n+1}) + \underline{c}_2 + b_\infty^2 \bar{c}^2 / \underline{c}_1. \end{aligned}$$

The proof is then completed using that for any $k \in \mathbb{N}^*$, $(\underline{c}_1/4 - (3L/2)(1 + \sigma_1^2)\eta_{k+1}) \geq 0$ and (2.5).

2.5.4 Proof of Proposition 2.6

Assume first **MA 1**, **MA 3(a)** and **A 5(1)**. Setting $w = \mathbf{1}$, we prove that **MA 2(1)** holds. Since $e_\infty = \sup_{x \in \mathbf{X}, \theta \in \Theta} \|e_\theta(x)\|_\theta$ is finite because of **A 5(1)**, then **MA 2(1)-(a)-(b)** hold. Note that the condition for any $\theta \in \Theta$, $x, x' \in \mathbf{X}$, $\|\delta_x P_\theta^k - \delta_{x'} P_\theta^k\|_{\text{TV}} \leq C_P(1 - \varepsilon_P)^k$ implies, since π_θ is the unique stationary distribution for P_θ , using Jensen inequality,

$$\|\delta_x P_\theta^k - \pi_\theta\|_{\text{TV}} \leq C_P(1 - \varepsilon_P)^k. \quad (2.25)$$

Therefore, the function \hat{e} given for any $\theta \in \Theta$ and $x \in \mathbf{X}$,

$$\hat{e}_\theta(x) = \sum_{k=0}^{+\infty} \{P_\theta^k e_\theta(x) - \pi_\theta(e_\theta)\}, \quad (2.26)$$

is well defined using the Minkowski's integral inequality for any $x \in \mathbf{X}$, $\theta \in \Theta$,

$$\|\hat{e}_\theta(x)\|_\theta \leq \sum_{k=0}^{+\infty} \|P_\theta^k e_\theta(x) - \pi_\theta(e_\theta)\|_\theta = e_\infty \sum_{k=0}^{+\infty} \|\delta_x P_\theta^k - \pi_\theta\|_{\text{TV}} \leq C_P e_\infty / \varepsilon_P. \quad (2.27)$$

Then, **MA 2(1)-(c)** holds.

To prove **MA 2(1)-(d)**, since \hat{e} satisfies (2.9) and e satisfies (2.11), it is sufficient to show that there exists $D_1 \geq 0$ such that

$$\|\hat{e}_{\theta_1}(x) - \mathsf{T}_{01}^\gamma \hat{e}_{\theta_0}(x)\|_{\theta_1} \leq D_1 \ell(\gamma),$$

for any $x \in \mathbf{X}$, $\theta_0, \theta_1 \in \Theta$ and γ a geodesic between θ_0 and θ_1 . Consider $\theta_0, \theta_1 \in \Theta$ and γ a geodesic between θ_0 and θ_1 . Note that by (2.26), **A 5(1)** and using that the parallel transport associated with the Levi-Civita connection is a linear isometry [Lee19, Proposition 5.5] and the same argument as proving (2.27) for any $x \in \mathbf{X}$, setting $\Delta e_\theta(x) = e_{\theta_1}(x) - \mathsf{T}_{01}^\gamma e_{\theta_0}(x)$ and using $\|\Delta e_\theta(x)\|_\theta \leq C\ell(\gamma)$,

$$\begin{aligned} & \|\hat{e}_{\theta_1}(x) - \mathsf{T}_{01}^\gamma \hat{e}_{\theta_0}(x)\|_{\theta_1} \\ & \leq \sum_{k=0}^{+\infty} \left\| P_{\theta_1}^k e_{\theta_1}(x) - \mathsf{T}_{01}^\gamma P_{\theta_0}^k e_{\theta_0}(x) - (\pi_{\theta_1}(e_{\theta_1}) - \mathsf{T}_{01}^\gamma \pi_{\theta_0}(e_{\theta_0})) \right\|_{\theta_1} \\ & \leq \sum_{k=0}^{+\infty} \left\| P_{\theta_1}^k \Delta e_\theta(x) - \pi_{\theta_1}(\Delta e_\theta) \right\|_{\theta_1} + \sum_{k=0}^{+\infty} \left\| [P_{\theta_1}^k - P_{\theta_0}^k](e_{\theta_0})(x) - [\pi_{\theta_1} - \pi_{\theta_0}](e_{\theta_0}) \right\|_{\theta_0} \\ & \leq CC_P \ell(\gamma) / \varepsilon_P + \sum_{k=0}^{+\infty} \left\| [P_{\theta_1}^k - P_{\theta_0}^k](e_{\theta_0})(x) - [\pi_{\theta_1} - \pi_{\theta_0}](e_{\theta_0}) \right\|_{\theta_0}. \end{aligned}$$

2.5. Proofs for Section 2.2

It remains to show that there exists $D_2 \geq 0$ such that

$$\sum_{k=0}^{+\infty} \left\| [P_{\theta_1}^k - P_{\theta_0}^k](e_{\theta_0})(x) - [\pi_{\theta_1} - \pi_{\theta_0}](e_{\theta_0}) \right\|_{\theta_0} \leq D_2 \ell(\gamma), \quad (2.28)$$

which follows from **MA 3-(a)**, (2.25) and the following decomposition,

$$\begin{aligned} [P_{\theta_1}^k - P_{\theta_0}^k](e_{\theta_0})(x) - [\pi_{\theta_1} - \pi_{\theta_0}](e_{\theta_0}) &= -\pi_{\theta_1} P_{\theta_0}^k e_{\theta_0} + \pi_{\theta_0}(e_{\theta_0}) \\ &\quad + \sum_{i=0}^{k-1} \{\delta_x P_{\theta_1}^{k-i-1} - \pi_{\theta_1}\} [P_{\theta_1} - P_{\theta_0}] \{P_{\theta_0}^i - \pi_{\theta_0}\} e_{\theta_0}. \end{aligned} \quad (2.29)$$

To obtain (2.29), expand the summation $\sum_{i=0}^{k-1} P_{\theta_1}^{k-i-1} (P_{\theta_1} - P_{\theta_0}) P_{\theta_0}^i$ and change the variables in the second sum, for any $k \in \mathbb{N}$, to get,

$$\sum_{i=0}^{k-1} P_{\theta_1}^{k-i-1} (P_{\theta_1} - P_{\theta_0}) P_{\theta_0}^i = \sum_{i=0}^{k-1} P_{\theta_1}^{k-i} P_{\theta_0}^i - \sum_{i=1}^k P_{\theta_1}^{k-i} P_{\theta_0}^i = P_{\theta_1}^k - P_{\theta_0}^k. \quad (2.30)$$

In order to use (2.25), we note that $\sum_{i=0}^{k-1} P_{\theta_1}^{k-i-1} (P_{\theta_1} - P_{\theta_0}) P_{\theta_0}^i = \sum_{i=0}^{k-1} P_{\theta_1}^{k-i-1} (P_{\theta_1} - P_{\theta_0}) \{P_{\theta_0}^i - \pi_{\theta_0}\}$ because $P_{\theta_1} \pi_{\theta_0} = P_{\theta_0} \pi_{\theta_0} = \pi_{\theta_0}$, and secondly that $\pi_{\theta_1} \sum_{i=0}^{k-1} (P_{\theta_1} - P_{\theta_0}) \{P_{\theta_0}^i - \pi_{\theta_0}\} = \pi_{\theta_1} - \pi_{\theta_1} P_{\theta_0}^k$. To bound (2.29) in order to obtain the announced bound (2.28), we use the triangle inequality and look at two separate terms. Let $A_k(x, \theta_0, \theta_1) = \sum_{i=0}^{k-1} \{\delta_x P_{\theta_1}^{k-i-1} - \pi_{\theta_1}\} [P_{\theta_1} - P_{\theta_0}] \{P_{\theta_0}^i - \pi_{\theta_0}\} e_{\theta_0}$. Using (2.25) and **MA 3-(a)**, for any $k \in \mathbb{N}, x \in \mathbf{X}, \theta_0, \theta_1 \in \Theta$,

$$\begin{aligned} \|A_k(x, \theta_0, \theta_1)\|_{\theta_0} &\leq \sum_{i=0}^{k-1} \|\delta_x P_{\theta_1}^{k-i-1} - \pi_{\theta_1}\|_{\text{TV}} \sup_{y \in \mathbf{X}} \left\| \delta_y [P_{\theta_1} - P_{\theta_0}] \{P_{\theta_0}^i - \pi_{\theta_0}\} e_{\theta_0} \right\|_{\theta_0} \\ &\leq \sum_{i=0}^{k-1} C_P (1 - \varepsilon_P)^{k-i-1} \sup_{y, z \in \mathbf{X}} \|\delta_y P_{\theta_1} - \delta_y P_{\theta_0}\|_{\text{TV}} \|\delta_z P_{\theta_0}^i - \pi_{\theta_0}\|_{\text{TV}} e_{\infty} \\ &\leq C C_P^2 e_{\infty} k (1 - \varepsilon_P)^{k-1} \ell(\gamma). \end{aligned} \quad (2.31)$$

Secondly, for any $k \in \mathbb{N}, x \in \mathbf{X}, \theta_0, \theta_1 \in \Theta$, using [Dou+18, Corollary D.2.5], for the measures $(\pi_{\theta_0}, \pi_{\theta_1})$ and the bounded function $x \mapsto \delta_x P_{\theta_0}^k(e_{\theta_0})$,

$$\begin{aligned} \left\| \pi_{\theta_0}(e_{\theta_0}) - \pi_{\theta_1} P_{\theta_0}^k(e_{\theta_0}) \right\|_{\theta_0} &= \left\| \pi_{\theta_0} P_{\theta_0}^k(e_{\theta_0}) - \pi_{\theta_1} P_{\theta_0}^k(e_{\theta_0}) \right\|_{\theta_0} \\ &\leq \|\pi_{\theta_0} - \pi_{\theta_1}\|_{\text{TV}} \sup_{x, y \in \mathbf{X}} \left\| \delta_x P_{\theta_0}^k(e_{\theta_0}) - \delta_y P_{\theta_0}^k(e_{\theta_0}) \right\|_{\theta_0} \\ &\leq \|\pi_{\theta_0} - \pi_{\theta_1}\|_{\text{TV}} C_P (1 - \varepsilon_P)^k e_{\infty}. \end{aligned} \quad (2.32)$$

Moreover, for any $\ell \in \mathbb{N}^*$ and $x \in \mathbf{X}$, using (2.25),

$$\begin{aligned} \|\pi_{\theta_0} - \pi_{\theta_1}\|_{\text{TV}} &\leq \|\pi_{\theta_0} - \delta_x P_{\theta_0}^{\ell}\|_{\text{TV}} + \|\delta_x P_{\theta_0}^{\ell} - \delta_x P_{\theta_1}^{\ell}\|_{\text{TV}} + \|\delta_x P_{\theta_1}^{\ell} - \pi_{\theta_1}\|_{\text{TV}} \\ &\leq 2C_P (1 - \varepsilon_P)^{\ell} + \|\delta_x P_{\theta_0}^{\ell} - \delta_x P_{\theta_1}^{\ell}\|_{\text{TV}}. \end{aligned} \quad (2.33)$$

Now, taking any function $f : \mathsf{X} \rightarrow \mathbb{R}$ such that $\sup_{\mathsf{X}} |f| \leq 1$ and using (2.30), (2.25) and **MA3-(a)** we have

$$\begin{aligned} \left| \left(\delta_x P_{\theta_1}^\ell - \delta_x P_{\theta_0}^\ell \right) f \right| &= \left| \sum_{i=0}^{\ell-1} \delta_x P_{\theta_1}^{\ell-i-1} (P_{\theta_1} - P_{\theta_0}) \{P_{\theta_0}^i - \pi_{\theta_0}\} f \right| \\ &\leq \sum_{i=0}^{\ell-1} \left\| \delta_x P_{\theta_1}^{\ell-i-1} \right\|_{\text{TV}} \sup_{y, z \in \mathsf{X}} \|\delta_y (P_{\theta_1} - P_{\theta_0})\|_{\text{TV}} \left\| \delta_z P_{\theta_0}^i - \pi_{\theta_0} \right\|_{\text{TV}} \\ &\leq \sum_{i=0}^{\ell-1} C \ell(\gamma) C_P (1 - \varepsilon_P)^i \leq C C_P \ell(\gamma) / \varepsilon_P . \end{aligned} \quad (2.34)$$

Taking the supremum over f in (2.34), we obtain

$$\|\delta_x P_{\theta_0}^\ell - \delta_x P_{\theta_1}^\ell\|_{\text{TV}} \leq C C_P \ell(\gamma) / \varepsilon_P . \quad (2.35)$$

Plugging (2.35) in (2.33) and taking $\ell \rightarrow \infty$, we have $\|\pi_{\theta_0} - \pi_{\theta_1}\|_{\text{TV}} \leq C C_P \ell(\gamma) / \varepsilon_P$ which, combined with (2.32), gives

$$\left\| \pi_{\theta_0}(e_{\theta_0}) - \pi_{\theta_1} P_{\theta_0}^k(e_{\theta_0}) \right\|_{\theta_0} \leq C C_P^2 e_\infty (1 - \varepsilon_P)^k \ell(\gamma) / \varepsilon_P . \quad (2.36)$$

Combining (2.31) and (2.36) into (2.29) brings the desired result (2.28) for $D_2 = C C_P e_\infty / \varepsilon_P^2$. This concludes the first point of Proposition 2.6.

Now, for the second point of Proposition 2.6, assume there exists a measurable function $w : \mathsf{X} \rightarrow [1, +\infty)$ such that **MA3(w)-(b)** and **A5(w)** hold. By assumption, **MA2(w)-(a)** holds. Besides, by a straightforward induction for any $x \in \mathsf{X}$, $k \in \mathbb{N}^*$, $P^k w(x) \leq \lambda^k w(x) + b_w \sum_{i=0}^{k-1} \lambda^i \leq \lambda^k w(x) + b_w / (1 - \lambda)$ and therefore **MA2(w)-(b)** holds. By **MA3(w)-(b)**, for any $\theta \in \Theta$, $P_\theta = P$, $\pi_\theta = \pi$ for some probability distribution π on $(\mathsf{X}, \mathcal{X})$. In addition, using Jensen inequality, for any $x \in \mathsf{X}$, $P w^{1/2}(x) \leq \lambda^{1/2} w^{1/2}(x) + b_w^{1/2} \mathbb{1}_{\mathsf{C}}(x)$. By [MT09, Theorem 16.0.1], there exist $D_2 \geq 0$ and $\rho \in [0, 1)$ such that for any $x \in \mathsf{X}$,

$$\|\delta_x P^k - \pi\|_{w^{1/2}} \leq C'_P \rho^k P w^{1/2}(x) . \quad (2.37)$$

Therefore, \hat{e} given by (2.26) is well defined since for any $\theta \in \Theta$ and $x \in \mathsf{X}$, using the Minkowski's inequality and (2.11),

$$\begin{aligned} \|\hat{e}_\theta(x)\|_\theta &\leq \sum_{k=0}^{+\infty} \|P^k e_\theta(x) - \pi_\theta(e_\theta)\|_\theta \\ &= e_\infty \sum_{k=0}^{+\infty} \|\delta_x P^k - \pi\|_{w^{1/2}} \\ &\leq C'_P e_\infty w^{1/2}(x) / (1 - \rho_P) . \end{aligned} \quad (2.38)$$

Then, **MA2(w)-(c)** holds.

To prove **MA 2(w)-(d)**, since \hat{e} satisfies (2.9), it is sufficient to show that there exists $D_3 \geq 0$ such that $\|\hat{e}_{\theta_1}(x) - T_{01}^Y \hat{e}_{\theta_0}(x)\|_{\theta_1} \leq D_3 w^{1/2}(x) \ell(\gamma)$. Note that by (2.26)-(2.11)-(2.37) and using that the parallel transport associated with the Levi-Civita connection is a linear isometry [Lee19, Proposition 5.5] and the same argument as proving (2.38) for any $x \in \mathsf{X}$, setting $\Delta e_\theta(x) = e_{\theta_1}(x) - T_{01}^Y e_{\theta_0}(x)$,

$$\begin{aligned} \|\hat{e}_{\theta_1}(x) - T_{01}^Y \hat{e}_{\theta_0}(x)\|_{\theta_1} &\leq \sum_{k=0}^{+\infty} \left\| P^k e_{\theta_1}(x) - T_{01}^Y P^k e_{\theta_0}(x) - (\pi(e_{\theta_1}) - T_{01}^Y \pi(e_{\theta_0})) \right\|_{\theta_1} \\ &= \sum_{k=0}^{+\infty} \left\| P^k \Delta e_\theta - \pi(\Delta e_\theta) \right\|_{\theta_1} \leq C'_P C \ell(\gamma) w^{1/2}(x) \sum_{k=0}^{+\infty} \rho_P^k, \end{aligned}$$

which completes the proof.

2.5.5 Proof of Theorem 2.5

Under **A 1-A 2-A 3-MA 2(w)** and for a sequence of step sizes $(\eta_k)_{k \in \mathbb{N}^*}$, consider the constants $D_{\hat{e}}$, $C(\eta_1)$ and $C_{\hat{e}}^0$ defined as follows:

$$C(\eta_1) = \bar{c} \hat{e}_\infty (\eta_1 + 2), \quad (2.39)$$

if w is bounded by $w_\infty = \sup_{\mathsf{X}} w$,

$$\begin{aligned} D_{\hat{e}} &= \hat{e}_\infty C_w L + 2\bar{c} a_2 + \bar{c} \hat{e}_\infty C_w + L_{\hat{e}} \bar{c} w_\infty [1 + a_1(e_\infty + b_\infty)], \\ C_{\hat{e}} &= C_{\hat{e}}^0 + (3L/2)(b_\infty^2 + e_\infty^2 C_w), \\ C_{\hat{e}}^0 &= L \hat{e}_\infty C_w (e_\infty + b_\infty) + L_{\hat{e}} \bar{c} C_w (e_\infty + b_\infty); \end{aligned} \quad (2.40)$$

if **A 4** holds,

$$\begin{aligned} D_{\hat{e}} &= \hat{e}_\infty C_w L + 2\bar{c} a_2 + \bar{c} \hat{e}_\infty C_w, \\ C_{\hat{e}} &= C_{\hat{e}}^0 + (3L/2)(b_\infty^2 + e_\infty^2 C_w), \\ C_{\hat{e}}^0 &= L \hat{e}_\infty C_w (e_\infty + b_\infty) + L_{\hat{e}} \bar{c} C_w [(e_\infty + b_\infty)\{a_1 h_\infty^2 + 1\} + h_\infty^2]. \end{aligned} \quad (2.41)$$

Lemma 2.16. *Assume **A 1-A 2-A 3-MA 1-MA 2(w)** hold for some measurable function $w : [1, +\infty)$. Assume in addition either $\sup_{x \in \mathsf{X}} w(x) = w_\infty < +\infty$ or **A 4**. Let $(\eta_k)_{k \in \mathbb{N}^*}$ be a sequence satisfying (2.10) and set for any $k \in \mathbb{N}$, $\Delta M_k = \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k}$. Consider $(\theta_k)_{k \in \mathbb{N}}$ defined by (2.3). Then,*

$$|\mathbb{E}[\sum_{k=0}^n \eta_{k+1} \Delta M_k]| \leq D_{\hat{e}} \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] + C_{\hat{e}}^0 \Gamma_{n+1}^{(2)} + C(\eta_1).$$

Proof. Consider the measurable function $\hat{e} : \Theta \times \mathsf{X} \rightarrow \text{T}\Theta$ which satisfies **MA 2(w)** and for any $k \in \mathbb{N}$, consider $\gamma^{(k+1)} : [0, 1] \rightarrow \Theta$ the geodesic between θ_k and θ_{k+1} defined by $\gamma^{(k+1)}(t) = \text{Exp}\{t \eta_{k+1} (H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}))\}$ for any $t \in [0, 1]$. Note that for any $k \in \mathbb{N}$,

$$\ell(\gamma^{(k+1)}) = \eta_{k+1} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}. \quad (2.42)$$

Using that the parallel transport associated with the Levi-Civita connection is a linear isometry [Lee19, Proposition 5.5] and $(T_{01}^\gamma)^{-1} = T_{10}^\gamma$ by uniqueness of parallel transport [Lee19, Theorem 4.32], we obtain the following decomposition

$$\mathbb{E} \left[- \sum_{k=0}^n \eta_{k+1} \Delta M_k \right] = - \mathbb{E} \left[\sum_{i=1}^5 A_i \right],$$

where

$$\begin{aligned} A_1 &= \sum_{k=1}^n \eta_{k+1} \langle \text{grad } V(\theta_k), \hat{e}_{\theta_k}(X_{k+1}) - P_{\theta_k} \hat{e}_{\theta_k}(X_k) \rangle_{\theta_k}, \\ A_2 &= \sum_{k=1}^n \eta_{k+1} \left\langle \text{grad } V(\theta_k), P_{\theta_k} \hat{e}_{\theta_k}(X_k) - T_{01}^{\gamma^{(k)}} P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \right\rangle_{\theta_k}, \\ A_3 &= \sum_{k=1}^n \eta_{k+1} \left\langle T_{10}^{\gamma^{(k)}} \text{grad } V(\theta_k) - \text{grad } V(\theta_{k-1}), P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \right\rangle_{\theta_{k-1}}, \\ A_4 &= \sum_{k=1}^n (\eta_{k+1} - \eta_k) \left\langle \text{grad } V(\theta_{k-1}), P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \right\rangle_{\theta_{k-1}}, \\ A_5 &= \eta_1 \langle \text{grad } V(\theta_0), \hat{e}_{\theta_0}(X_1) \rangle_{\theta_0} - \eta_{n+1} \langle \text{grad } V(\theta_n), P_{\theta_n} \hat{e}_{\theta_n}(X_{n+1}) \rangle_{\theta_n}. \end{aligned}$$

We now bound each term of this decomposition. Note that A_2 is the only term which will be bounded differently depending on either the assumption $\sup_{x \in \mathcal{X}} w(x) = w_\infty < +\infty$ or **A 4**. That is why we deal first with A_i , $i \in \{1, \dots, 5\} \setminus \{2\}$. Note that using **MA 1** and **MA 2(w)-(c)**, we obtain $\mathbb{E}[\langle \text{grad } V(\theta_k), \hat{e}_{\theta_k}(X_{k+1}) - P_{\theta_k} \hat{e}_{\theta_k}(X_k) \rangle_{\theta_k} | \mathcal{F}_k] = 0$ for any $k \in \{1, \dots, n\}$, and therefore

$$\mathbb{E}[A_1] = 0. \quad (2.43)$$

Using the Cauchy-Schwarz inequality, **A 2-(b)**, the definition of e (2.8), **A 3**, **MA 2(w)-(a)-(c)** and Jensen's inequality, we obtain

$$\begin{aligned} |A_3| &\leq L \sum_{k=1}^n \eta_{k+1} \ell(\gamma^{(k)}) \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \\ &\leq L \sum_{k=1}^n \eta_k \eta_{k+1} \left\{ \|H_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right\} \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}}, \\ &\leq L \hat{e}_\infty \sum_{k=1}^n \eta_k \eta_{k+1} \left(e_\infty w^{1/2}(X_k) + b_\infty + \|h(\theta_{k-1})\|_{\theta_{k-1}} \right) P_{\theta_{k-1}} w^{1/2}(X_k). \end{aligned}$$

Taking the expectation, using **MA 2(w)-(b)**, that $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (2.10) and the Cauchy-Schwarz inequality brings,

$$\mathbb{E}[|A_3|] \leq L \hat{e}_\infty C_w \{ [e_\infty + b_\infty] \Gamma_{n+1}^{(2)} + \sum_{k=1}^n \eta_k^2 \mathbb{E}[\|h(\theta_{k-1})\|_{\theta_{k-1}}^2] \}. \quad (2.44)$$

2.5. Proofs for Section 2.2

Using the Cauchy-Schwarz and Jensen inequality, **A2-(a)**, **MA2(w)-(c)**, for any $a, b \in \mathbb{R}$, $|ab| \leq (a^2 + b^2)/2$ and that $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (2.10), we have

$$\begin{aligned} |A_4| &\leq \bar{c} \sum_{k=1}^n |\eta_{k+1} - \eta_k| \|h(\theta_{k-1})\|_{\theta_{k-1}} \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}}, \\ &\leq 2\bar{c}\hat{e}_\infty^2 \sum_{k=1}^n |\eta_{k+1} - \eta_k| \{P_{\theta_{k-1}} w^{1/2}(X_k)\}^2 + 2\bar{c}a_2 \sum_{k=1}^n \eta_k^2 \|h(\theta_{k-1})\|_{\theta_{k-1}}^2. \end{aligned}$$

Taking the expectation, using **MA2(w)-(b)** and $(\eta_k)_{k \in \mathbb{N}^*}$ is non-increasing bring,

$$\mathbb{E}[|A_4|] \leq 2\bar{c}\hat{e}_\infty^2 C_w \eta_1 + 2\bar{c}a_2 \sum_{k=1}^n \eta_k^2 \mathbb{E}[\|h(\theta_{k-1})\|_{\theta_{k-1}}^2]. \quad (2.45)$$

Finally, using **A2-(a)**, **MA2(w)-(c)** and Jensen's inequality, we obtain,

$$\begin{aligned} |A_5| &\leq \eta_1 \bar{c} \|h(\theta_0)\|_{\theta_0} \|\hat{e}_{\theta_0}(X_1)\|_{\theta_0} + \eta_{n+1} \bar{c} \|h(\theta_n)\|_{\theta_n} P_{\theta_n} \|\hat{e}_{\theta_n}\|_{\theta_n}(X_{n+1}), \\ &\leq \bar{c}\hat{e}_\infty \left\{ \eta_1 w^{1/2}(X_1) \|h(\theta_0)\|_{\theta_0} + \eta_{n+1} P_{\theta_n} w^{1/2}(X_{n+1}) \|h(\theta_n)\|_{\theta_n} \right\}. \end{aligned}$$

Taking the expectation, using the Cauchy-Schwarz inequality, **MA2(w)-(b)** and that for any $a \in \mathbb{R}$, $a \leq a^2 + 1$ brings,

$$\begin{aligned} \mathbb{E}[|A_5|] &\leq \bar{c}\hat{e}_\infty (\eta_1 C_w^{1/2} \mathbb{E}[\|h(\theta_0)\|_{\theta_0}^2]^{1/2} + \eta_{n+1} C_w^{1/2} \mathbb{E}[\|h(\theta_n)\|_{\theta_n}^2]^{1/2}), \\ &\leq \bar{c}\hat{e}_\infty (2 + \eta_1^2 C_w \mathbb{E}[\|h(\theta_0)\|_{\theta_0}^2] + \eta_{n+1}^2 C_w \mathbb{E}[\|h(\theta_n)\|_{\theta_n}^2]), \\ &\leq \bar{c}\hat{e}_\infty (2 + C_w \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2]). \end{aligned} \quad (2.46)$$

It remains to treat A_2 depending on the additional two conditions we consider. We start by proving a general bound which hold. Using the Cauchy-Schwarz inequality, (2.42), **MA2(w)-(a)-(c)-(d)**, (2.42) and **A2-(a)**, we get

$$\begin{aligned} |A_2| &\leq L_{\hat{e}} \sum_{k=1}^n \eta_{k+1} \eta_k w^{1/2}(X_k) \|\text{grad } V(\theta_k)\|_{\theta_k} \left\{ \|H_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right\}, \\ &\leq L_{\hat{e}} \bar{c} \sum_{k=1}^n \eta_{k+1} \eta_k w^{1/2}(X_k) \|h(\theta_k)\|_{\theta_k} \left\{ e_\infty w^{1/2}(X_k) \right. \\ &\quad \left. + \|h(\theta_{k-1})\|_{\theta_{k-1}} + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right\}. \end{aligned}$$

Using that $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (2.10), **A3** and for any $a, b \in \mathbb{R}$, $|ab| \leq (a^2 + b^2)/2$, we obtain

$$\begin{aligned} |A_2| &\leq L_{\hat{e}} \bar{c} \sum_{k=1}^n \eta_k^2 \|h(\theta_k)\|_{\theta_k} \left(e_\infty w^{1/2}(X_k) + b_\infty w^{1/2}(X_k) \right) \\ &\quad + (L_{\hat{e}} \bar{c} / 2) \sum_{k=1}^n w^{1/2}(X_k) \left(\|h(\theta_k)\|_{\theta_k}^2 \eta_{k+1}^2 + \|h(\theta_{k-1})\|_{\theta_{k-1}}^2 \eta_k^2 \right). \end{aligned}$$

Changing the index in the second sum, using that for any $x \in \mathsf{X}$, $w^{1/2}(x) \leq w(x)$ since $w(x) \geq 1$, for any $a \in \mathbb{R}$, $a \leq a^2 + 1$ on $\|h(\theta_k)\|_{\theta_k}$ and the second inequality in (2.10) we get

$$\begin{aligned} |A_2| &\leq L_{\hat{e}}\bar{c} \sum_{k=1}^n \eta_k^2 w(X_k)(e_\infty + b_\infty) \\ &\quad + L_{\hat{e}}\bar{c} \sum_{k=0}^n \eta_{k+1}^2 \|h(\theta_k)\|_{\theta_k}^2 ([w(X_k) + w(X_{k+1})]/2 + a_1 w(X_k)[e_\infty + b_\infty]) . \end{aligned}$$

Now, taking the full expectation and using **MA2**(w)-(b) brings

$$\begin{aligned} \mathbb{E}[|A_2|] &\leq L_{\hat{e}}\bar{c}(e_\infty + b_\infty)C_w\Gamma_{n+1}^{(2)} \\ &\quad + L_{\hat{e}}\bar{c} \sum_{k=1}^n \eta_{k+1}^2 \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 (\{1/2 + a_1(e_\infty + b_\infty)\}w(X_k) + w(X_{k+1})/2) \right] . \end{aligned}$$

We now distinguish the two cases. If $\sup_{x \in \mathsf{X}} w(x) = w_\infty < +\infty$, we obtain

$$\mathbb{E}[|A_2|] \leq L_{\hat{e}}\bar{c}(e_\infty + b_\infty)C_w\Gamma_{n+1}^{(2)} + L_{\hat{e}}\bar{c}w_\infty\{1 + a_1(e_\infty + b_\infty)\} \sum_{k=1}^n \eta_{k+1}^2 \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] . \quad (2.47)$$

If **A4** holds, we get

$$\mathbb{E}[|A_2|] \leq L_{\hat{e}}\bar{c} \left[(e_\infty + b_\infty)\{h_\infty^2 a_1 + 1\} + h_\infty^2 \right] C_w\Gamma_{n+1}^{(2)} . \quad (2.48)$$

Combining (2.43)-(2.44)-(2.45)-(2.46)-(2.47)-(2.48), completes the proof. \square

Proof of Theorem 2.5. The proof only consists in applying Lemma 2.3 taking $\varepsilon = \underline{c}_1/(4\bar{c}^2)$, Lemma 2.16 using $\sup_{k \in \mathbb{N}^*} \eta_k \leq \underline{c}_1/(4(3L/2 + D_{\hat{e}}))$ and **A3-MA2**(w)-(a)-(b). \square

2.6 Proofs for Section 2.3.1

2.6.1 Proof of Lemma 2.7

We preface the proof of the Lemma by a preliminary result which does not assume **R1**-(iii).

Lemma 2.17. *Assume **R1**-(i)-(ii) hold.*

(a) *For any $(\theta, u) \in \text{T}\Theta$, setting $\gamma(t) = \text{Ret}_\theta(tu)$ for $t \in \mathbb{R}_+$, we have,*

$$D_t \dot{\gamma}(0) = (D^2 \Phi_\theta)_{0_\theta}(u, u) . \quad (2.49)$$

(b) *Under **R2**, $\|\Phi_\theta(u) - u\|_\theta \leq \mathcal{L}^{(1)}(\theta)\|u\|_\theta^2/2$ for any $(\theta, u) \in \text{T}\Theta$.*

(c) *Under **R3**, $\|\Phi_\theta(u) - u\|_\theta \leq \mathcal{L}^{(2)}(\theta)\|u\|_\theta^3/6$ for any $(\theta, u) \in \text{T}\Theta$.*

Proof. (a) The proof relies on the use of normal coordinates with origin at θ [Lee19, Chapter 5]. These coordinates are smooth and simply defined identifying $T_\theta\Theta$ with \mathbb{R}^d through Exp_θ^{-1} . More precisely, setting an orthonormal basis $\{\mathbf{b}_i : i \in \{1, \dots, d\}\}$ of $T_\theta\Theta$, define for any $\tilde{\theta} \notin \text{Cut}(\theta)$,

$$\varphi^i(\tilde{\theta}) = \langle \text{Exp}_\theta^{-1}(\tilde{\theta}), \mathbf{b}_i \rangle_\theta .$$

Then, $\varphi = \{\varphi^i : i \in \{1, \dots, d\}\}$ are smooth coordinates around θ . Therefore, by definition Φ_θ is simply Ret_θ read in these coordinates. Then by [Lee19, Equation (4.15)], setting for any $t \in \mathbb{R}_+$, $\gamma(t) = \text{Ret}_\theta(tu)$ for $u \in T_\theta\Theta$, we get that, in these coordinates,

$$D_t \dot{\gamma}(t) = (D^2\Phi_\theta)_{tu}(u, u) + \sum_{k=1}^d \sum_{i,j=1}^d D\Phi_\theta^i \gamma^j(t) \Gamma_{i,j}^k(\gamma(t)) \partial_k ,$$

where D_t is the covariant derivative along γ , $\{\Gamma_{i,j}^k : i, j, k \in \{1, \dots, d\}\}$ are the Christoffel symbols and $\{\partial_k : k \in \{1, \dots, d\}\}$ are the coordinate vector fields on $T\Theta$ corresponding to φ . But using [Lee19, Proposition 5.24], we get that $\Gamma_{i,j}^k(\gamma(0)) = \Gamma_{i,j}^k(\theta) = 0$ for any $i, j, k \in \{1, \dots, d\}$. Therefore,

$$D_t \dot{\gamma}(0) = (D^2\Phi_\theta)_{0_\theta}(u, u) .$$

(b) Let $\theta \in \Theta$ and for any $u \in T_\theta\Theta$, consider the first-order Taylor expansion of $\Phi_\theta : T_\theta\Theta \rightarrow T_\theta\Theta$, taken at 0_θ ,

$$\Phi_\theta(u) = \Phi_\theta(0_\theta) + (D\Phi_\theta)_{0_\theta}(u) + \int_0^1 (1-t) (D^2\Phi_\theta)_{tu}(u, u) dt , \quad (2.50)$$

where $D\Phi_\theta$ and $D^2\Phi_\theta$ denote the first and second derivative of Φ_θ . For the first term,

$$\Phi_\theta(0_\theta) = \text{Exp}_\theta^{-1} \circ \text{Ret}_\theta(0_\theta) = \text{Exp}_\theta^{-1}(\theta) = 0_\theta , \quad (2.51)$$

where the second equality follows because $\text{Ret}_\theta(0_\theta) = \theta$, by **R 1-(i)** and by definition that $\text{Exp}_\theta(0_\theta) = \theta$. For the second term, using that Exp_θ^{-1} and Ret_θ are continuously differentiable as function between smooth manifolds, $(D\text{Exp}_\theta)_{0_\theta} = \text{Id}$ by definition and **R 1-(i)**, we obtain that

$$(D\Phi_\theta)_{0_\theta} = (D\text{Exp}_\theta^{-1})_{\text{Ret}_\theta(0_\theta)}(D\text{Ret}_\theta)_{0_\theta} = (D\text{Exp}_\theta^{-1})_\theta(D\text{Ret}_\theta)_{0_\theta} = \text{Id} . \quad (2.52)$$

The proof is then completed using (2.51), (2.52) and **R 2** in (2.50).

(c) Let $\theta \in \Theta$ and consider the second-order Taylor expansion of $\Phi_\theta : T_\theta\Theta \rightarrow T_\theta\Theta$, taken at 0_θ :

$$\Phi_\theta(u) = u + (D^2\Phi_\theta)_{0_\theta}(u, u)/2 + 2^{-1} \int_0^1 (1-t)^2 (D^3\Phi_\theta)_{tu}(u, u, u) dt . \quad (2.53)$$

where $D^3\Phi_\theta$ is the third derivative of Φ_θ . Using (2.49) and **R 3**, $D_t \dot{\gamma}(0) = 0$, which implies that $(D^2\Phi_\theta)_{0_\theta}(u, u) = 0_\theta$. Plugging this result into (2.53) and using the bound on the third derivative of Φ_θ given by **R 3** completes the proof. \square

Proof of Lemma 2.7. Recall that for any $(\theta, u) \in T\Theta$ and $g \in \mathbf{G}$, we denote by $g \cdot u = (Dg)_\theta(u)$ and g is an isometry, i.e. $\|g \cdot u\|_{g \cdot \theta} = \|u\|_\theta$, see Section 2.11.7. The proof consists in showing that for any $(\theta, u) \in T\Theta$, and $g \in \mathbf{G}$,

$$\|\Phi_\theta(u) - u\|_\theta = \|\Phi_{g \cdot \theta}(g \cdot u) - g \cdot u\|_{g \cdot \theta} . \quad (2.54)$$

Indeed, suppose that this result holds and consider a fixed $\theta_0 \in \Theta$. Since Θ is assumed to be homogeneous **R1-(iii)**, for any $(\theta, u) \in T\Theta$, there exists $g \in \mathbf{G}$ such that $\theta_0 = g \cdot \theta$, which implies by (2.54),

$$\|\Phi_\theta(u) - u\|_\theta = \|\Phi_{\theta_0}(g \cdot u) - g \cdot u\|_{\theta_0} .$$

Using that $g \in \mathbf{G}$ is an isometry and Lemma 2.17 completes the proof.

We now show (2.54). Let $g \in \mathbf{G}$ and $(\theta, u) \in T\Theta$. Since g^{-1} is an isometry, we have $\|\Phi_\theta(u) - u\|_\theta = \|g^{-1} \cdot g \cdot \Phi_\theta(u) - g^{-1} \cdot g \cdot u\|_\theta = \|g \cdot \Phi_\theta(u) - g \cdot u\|_{g \cdot \theta}$, so we only need to show that $g \cdot \Phi_\theta(u) = \Phi_{g \cdot \theta}(g \cdot u)$. Using that g maps geodesics to geodesics (see (2.161) in Section 2.11.7) and **R1-(iii)** successively, we get that $g \cdot \Phi_\theta(u) = \text{Exp}_{g \cdot \theta}^{-1}(g \cdot \text{Ret}_\theta(u)) = \Phi_{g \cdot \theta}(g \cdot u)$ and the proof follows. \square

2.6.2 Projective Retraction on the Sphere

Consider the Euclidean unit sphere manifold $\Theta = \mathbb{S}^d = \{\mathbf{x} \in \mathbb{R}^{d+1} : \|\mathbf{x}\| = 1\}$, where $\|\cdot\|$ stands for the standard Euclidean norm on \mathbb{R}^{d+1} . By [AMS08, Example 3.5.1] for any $\theta \in \mathbb{S}^d$, $T_\theta \mathbb{S}^d = \{\mathbf{u} \in \mathbb{R}^{d+1} : \mathbf{u}^\top \theta = 0\}$. The Riemannian metric g is the canonical metric on the sphere, defined as the restriction of Euclidean scalar product from \mathbb{R}^{d+1} to the tangent space $T_\theta \mathbb{S}^d$. The corresponding Riemannian exponential is given by [Lee19, Proposition 5.27]:

$$\text{Exp}_\theta(\mathbf{u}) = \cos(\|\mathbf{u}\|)\theta + \sin(\|\mathbf{u}\|)(\mathbf{u}/\|\mathbf{u}\|) . \quad (2.55)$$

The following result holds.

Proposition 2.18. *The projective retraction Ret_θ defined for any $(\theta, \mathbf{u}) \in T\mathbb{S}^d$ by*

$$\text{Ret}_\theta(\mathbf{u}) = (\theta + \mathbf{u}) / \|\theta + \mathbf{u}\| \quad (2.56)$$

*satisfies **R1-R2-R3**.*

The retraction (2.56) is *both* a first-order and second-order retraction, and Φ_θ also has a bounded first-order derivative. Consequently, Theorem 2.11 and Theorem 2.12 can be applied according to conditions on the noise properties. We remark that by comparing (2.55) with (2.56), the retraction map Ret_θ has a better numerical stability as it does not involve evaluating the trigonometric functions.

Proof of Proposition 2.18. • Condition **R1-(i)** is satisfied by definition (2.56): this condition is just the definition of a retraction, as given in [AMS08].

- Condition **R1-(ii)** is satisfied: the cut locus of a point θ on the sphere \mathbb{S}^d is $\text{Cut}(\theta) = \{-\theta\}$ [Lee19, page 308]. The Riemannian (that is, spherical) distance between θ and $-\theta$ is $\rho_\Theta(\theta, -\theta) = \pi$. On the other hand, using (2.67) from Lemma 2.20, $\rho_\Theta(\theta, \text{Ret}_\theta(\mathbf{u})) < \pi/2$ because $\arctan(\|\mathbf{u}\|) < \pi/2$ for all $\mathbf{u} \in \mathbb{T}_\theta\Theta$. It is then clear that $\text{Ret}_\theta(\mathbf{u}) \notin \{-\theta\}$ for any $\mathbf{u} \in \mathbb{T}_\theta\Theta$.

- Condition **R1-(iii)** is satisfied: the isometry group of $\Theta = \mathbb{S}^d$ is $\mathbf{G} = \text{O}(d+1)$, the group of $(d+1) \times (d+1)$ orthogonal matrices. The action of \mathbf{G} on Θ is given by matrix-vector multiplication, if $\mathbf{G} \in \mathbf{G}$ is an orthogonal matrix, $\mathbf{G} \cdot \theta = \mathbf{G}\theta$ and $\mathbf{G} \cdot \mathbf{u} = \mathbf{G}\mathbf{u}$. From (2.56),

$$\mathbf{G} \cdot \text{Ret}_\theta(\mathbf{u}) = \mathbf{G} \cdot (\theta + \mathbf{u}) / \|\theta + \mathbf{u}\| . \quad (2.57)$$

However, since \mathbf{G} is an orthogonal matrix, \mathbf{G} preserves Euclidean norms, so $\|\theta + \mathbf{u}\| = \|\mathbf{G} \cdot (\theta + \mathbf{u})\|$. Replacing into (2.57) gives,

$$\mathbf{G} \cdot \text{Ret}_\theta(\mathbf{u}) = \mathbf{G} \cdot (\theta + \mathbf{u}) / \|\mathbf{G} \cdot (\theta + \mathbf{u})\| = (\mathbf{G} \cdot \theta + \mathbf{G} \cdot \mathbf{u}) / \|\mathbf{G} \cdot \theta + \mathbf{G} \cdot \mathbf{u}\| , \quad (2.58)$$

where the second equality follows since the action of \mathbf{G} is linear. Finally, the right-hand side of (2.58) is $\text{Ret}_{\mathbf{G} \cdot \theta}(\mathbf{G} \cdot \mathbf{u})$.

- Condition **R2** is satisfied: recall that $\Phi_\theta(\mathbf{u}) = \text{Exp}_\theta^{-1} \circ \text{Ret}_\theta(\mathbf{u})$. From (2.67) and **R1-(ii)**,

$$\Phi_\theta(\mathbf{u}) = \arctan(\|\mathbf{u}\|) \frac{\mathbf{u}}{\|\mathbf{u}\|} . \quad (2.59)$$

The required second derivative can now be computed, thanks to the identity, for any $t_0 \in [0, 1]$,

$$(\text{D}^2\Phi_\theta)_{t_0\mathbf{u}}(\mathbf{u}, \mathbf{u}) = \left. \frac{\text{d}^2}{\text{d}t^2} \Phi_\theta(t\mathbf{u}) \right|_{t=t_0} . \quad (2.60)$$

Indeed, using (2.59) and (2.60),

$$(\text{D}^2\Phi_\theta)_{t_0\mathbf{u}}(\mathbf{u}, \mathbf{u}) = \left. \frac{\text{d}^2}{\text{d}t^2} \arctan(t\|\mathbf{u}\|) \right|_{t=t_0} \frac{\mathbf{u}}{\|\mathbf{u}\|} = \|\mathbf{u}\|^2 \left(f_2(t_0\|\mathbf{u}\|) \frac{\mathbf{u}}{\|\mathbf{u}\|} \right) , \quad (2.61)$$

where $f_2 : x \mapsto -2x/(1+x^2)^2$ is the second derivative of the arctan function. However, for any $x \in \mathbb{R}$, $2x \leq 1+x^2 \leq (1+x^2)^2$ so $|f_2(x)| = |2x|/(1+x^2)^2 \leq 1$. Now, since $\Theta = \mathbb{S}^d$, here $\|\mathbf{u}\|_\theta = \|\mathbf{u}\|$. Thus, Condition **R2** is satisfied with $\mathcal{L}^{(1)}(\theta) = 1$.

- Condition **R3-(i)** is satisfied: recall (2.49) from Lemma 2.17-(a),

$$\text{D}_t\dot{\gamma}(0) = (\text{D}^2\Phi_\theta)_0(\mathbf{u}, \mathbf{u}) .$$

From (2.60), it then follows,

$$\text{D}_t\dot{\gamma}(0) = \left. \frac{\text{d}^2}{\text{d}t^2} \Phi_\theta(t\mathbf{u}) \right|_{t=0} = f_2(0)\|\mathbf{u}\|\mathbf{u} = 0 .$$

- Condition **R3-(ii)** is satisfied: the proof continues from (2.61) by differentiating once more. A straight-forward computation proves **R3-(i)** is satisfied with $\mathcal{L}^{(2)}(\theta) = 2$.
- Φ_θ has bounded first derivative: from (2.59), by differentiating,

$$(\mathbf{D}\Phi_\theta)_\mathbf{u}(\mathbf{v}) = \frac{1}{1 + \|\mathbf{u}\|^2} \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\|} \frac{\mathbf{u}}{\|\mathbf{u}\|} + \frac{\arctan(\|\mathbf{u}\|)}{\|\mathbf{u}\|} \left(\mathbf{v} - \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\|} \frac{\mathbf{u}}{\|\mathbf{u}\|} \right),$$

for any \mathbf{u} and \mathbf{v} in $T_\theta\Theta$. Then, by an elementary calculation, and recalling that, since $\Theta = \mathbb{S}^d$, Riemannian scalar products and norms are equal to Euclidean ones, $\|(\mathbf{D}\Phi_\theta)_\mathbf{u}(\mathbf{v})\|_\theta \leq 2\|\mathbf{v}\|_\theta$. Thus, the operator norm of $(\mathbf{D}\Phi_\theta)_\mathbf{u}$ is bounded by $\bar{\mathbf{D}} = 2$. \square

2.6.3 Proof of Proposition 2.10

First, recall by [AMS04, Theorem 3.6] that the Riemannian exponential map at $[\mathbf{B}]$ is given for $\mathbf{D} \in T_{[\mathbf{B}]}Gr_r(\mathbb{R}^d)$:

$$\text{Exp}_{[\mathbf{B}]}(\mathbf{D}) = \left[\left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \quad \mathbf{U} \right) \begin{pmatrix} \cos(\boldsymbol{\Sigma}) \\ \sin(\boldsymbol{\Sigma}) \end{pmatrix} \right], \quad (2.62)$$

where $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ is the compact singular value decomposition.

Lemma 2.19. *Let $[\mathbf{B}] \in Gr_r(\mathbb{R}^d)$, $\mathbf{D} \in T_{[\mathbf{B}]}Gr_r(\mathbb{R}^d)$ and the singular value decomposition $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$, where \mathbf{U} is $d \times r$ such that $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$, \mathbf{V} is $r \times r$ orthogonal and $\boldsymbol{\Sigma}$ is $r \times r$ diagonal with nonnegative elements. The retraction Ret defined by (2.15) can be written as $\text{Ret}_{[\mathbf{B}]} = \text{Exp}_{[\mathbf{B}]} \circ \Phi_{[\mathbf{B}]}$, where*

$$\Phi_{[\mathbf{B}]}(\mathbf{D}) = \mathbf{U} \arctan(\boldsymbol{\Sigma}) \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2}, \quad (2.63)$$

is well-defined using Lemma 2.35, and where $\arctan(\boldsymbol{\Sigma})$ is obtained by taking the arc tangent of each diagonal element of the matrix $\boldsymbol{\Sigma}$.

Proof. Let $\mathbf{B} \in St_r(\mathbb{R}^d)$, $\mathbf{D} \in T_{[\mathbf{B}]}Gr_r(\mathbb{R}^d)$ with singular value decomposition $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$, where \mathbf{U} is $d \times r$ with orthonormal columns i.e. $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$, $\mathbf{V} \in O_r(\mathbb{R})$ and $\boldsymbol{\Sigma}$ is a $r \times r$ nonnegative diagonal matrix. Consider the map $\Phi_{[\mathbf{B}]} : \mathbf{D} \mapsto \mathbf{U} \arctan(\boldsymbol{\Sigma}) \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2}$, which is well-defined using Lemma 2.35. Using (2.62), taking the exponential function brings,

$$\text{Exp}_{[\mathbf{B}]}(\Phi_{[\mathbf{B}]}(\mathbf{D})) = \left[\left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \quad \mathbf{U} \right) \begin{pmatrix} \cos(\arctan(\boldsymbol{\Sigma})) \\ \sin(\arctan(\boldsymbol{\Sigma})) \end{pmatrix} \right].$$

However, since $\cos(\arctan(x)) = 1/(1+x^2)^{1/2}$ and $\sin(\arctan(x)) = x/(1+x^2)^{1/2}$, using that $\mathbf{I}_r + \boldsymbol{\Sigma}^2$ is a diagonal and invertible matrix, this becomes

$$\begin{aligned} \text{Exp}_{[\mathbf{B}]}(\Phi_{[\mathbf{B}]}(\mathbf{D})) &= \left[\left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \quad \mathbf{U} \right) \begin{pmatrix} \mathbf{I}_r \\ \boldsymbol{\Sigma} \end{pmatrix} (\mathbf{I}_r + \boldsymbol{\Sigma}^2)^{-1/2} \right] \\ &= \left[\left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \quad \mathbf{U} \right) \begin{pmatrix} \mathbf{I}_r \\ \boldsymbol{\Sigma} \end{pmatrix} \right], \end{aligned}$$

where we have used that $[\mathbf{B}\mathbf{M}] = [\mathbf{B}]$ for any \mathbf{B} in $\text{St}_r(\mathbb{R}^d)$ and \mathbf{M} in $\text{GL}_r(\mathbb{R})$. The block matrix computation brings

$$\text{Exp}_{[\mathbf{B}]}(\Phi_{[\mathbf{B}]}(\mathbf{D})) = \left[\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} + \mathbf{U}\Sigma \right] .$$

By right-multiplication of the invertible matrix $\mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2}$ and using that $\mathbf{V}\mathbf{V}^\top = \mathbf{I}_r$, we have

$$\text{Exp}_{[\mathbf{B}]}(\Phi_{[\mathbf{B}]}(\mathbf{D})) = \left[\mathbf{B} + \mathbf{U}\Sigma\mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2} \right] = [\mathbf{B} + \mathbf{D}] ,$$

which completes the proof. \square

Using Lemma 2.19, it is possible to show that the retraction given by (2.15) satisfies the conditions **R1-R2-R3**.

- Condition **R1-(i)** is satisfied: the first equality is straightforward. For the second one, one can check that $\text{Ret}_{[\mathbf{B}]}(\mathbf{D}) = \varphi_{(\mathbf{B}, \mathbf{B}_\perp)}((\mathbf{B}_\perp^\top \mathbf{B}_\perp)^{-1} \mathbf{B}_\perp^\top \mathbf{D})$, where $\varphi_{(\mathbf{B}, \mathbf{B}_\perp)}$ is the parameterization given in Proposition 2.22. Since these parameterizations define the differential structure, we have $(D\text{Ret}_{[\mathbf{B}]})_{\mathbf{D}} = \text{Id}$ for any (\mathbf{B}, \mathbf{D}) in $\text{TGr}_r(\mathbb{R}^d)$. It holds in particular for $\mathbf{D} = \mathbf{0}_{d \times r}$.

- Condition **R1-(ii)** is satisfied: when $\Theta = \text{Gr}_r(\mathbb{R}^d)$, using Proposition 2.26, the cut locus of each $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ is given by,

$$\text{Cut}([\mathbf{B}]) = \left\{ \text{Exp}_{[\mathbf{B}]}(\mathbf{D}) : \left\| \mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} \right\| = \pi/2 \right\} ,$$

where $\|\cdot\|$ is introduced in (2.83). Using Lemma 2.19, for any $\mathbf{D} \in \text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, one has

$$\text{Ret}_{[\mathbf{B}]}(\mathbf{D}) = \text{Exp}_{[\mathbf{B}]} \left(\mathbf{U} \arctan(\Sigma) \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2} \right) .$$

As explained right after (2.83), the operator norm of $\mathbf{U} \arctan(\Sigma) \mathbf{V}^\top$ is equal to the largest diagonal value of $\arctan(\Sigma)$, which is always strictly inferior to $\pi/2$. It follows that $\text{Ret}_{[\mathbf{B}]}(\mathbf{D}) \notin \text{Cut}([\mathbf{B}])$.

- Condition **R1-(iii)** is satisfied: the isometry group of $\Theta = \text{Gr}_r(\mathbb{R}^d)$ is $\mathbf{G} = \text{O}_d(\mathbb{R})$, the group of $d \times d$ orthogonal matrices. The action of $\text{O}_d(\mathbb{R})$ on $\text{Gr}_r(\mathbb{R}^d)$ is given by $\mathbf{U} \cdot [\mathbf{B}] = [\mathbf{U}\mathbf{B}]$ (the image of the subspace $[\mathbf{B}]$ of \mathbb{R}^d by the orthogonal transformation \mathbf{U}). If $\mathbf{D} \in \text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, then $\mathbf{U} \cdot \mathbf{D} = \mathbf{U}\mathbf{D}$ is a matrix product. Applying these identities in (2.15),

$$\mathbf{U} \cdot \text{Ret}_{[\mathbf{B}]}(\mathbf{D}) = \mathbf{U}([\mathbf{B} + \mathbf{D}]) = [\mathbf{U}\mathbf{B} + \mathbf{U}\mathbf{D}] = [\mathbf{U} \cdot \mathbf{B} + \mathbf{U} \cdot \mathbf{D}] = \text{Ret}_{[\mathbf{U} \cdot \mathbf{B}]}(\mathbf{U} \cdot \mathbf{D}) ,$$

which is exactly what is required for **R1-(iii)**.

- Condition **R2** is satisfied: the required second derivative is computed using the identity (this is a repetition of (2.60)), for any $(\theta, \mathbf{D}) \in \text{TGr}_r(\mathbb{R}^d)$, choosing a representative $\mathbf{B} \in \theta$ such that $\mathbf{B}^\top \mathbf{B} = \mathbf{I}_r$ and a $[\mathbf{B}_\perp] \in \text{Gr}_{d-r}(\mathbb{R}^d)$ such that $\mathbf{B}_\perp^\top \mathbf{B}_\perp = \mathbf{I}_{d-r}$, $\mathbf{B}^\top \mathbf{B}_\perp = \mathbf{0}_{r \times (d-r)}$ and $t_0 \in [0, 1]$,

$$(D^2 \Phi_{[\mathbf{B}]})_{t_0 \mathbf{D}}(\mathbf{D}, \mathbf{D}) = \left. \frac{d^2}{dt^2} \Phi_{[\mathbf{B}]}(t\mathbf{D}) \right|_{t=t_0} . \quad (2.64)$$

Using (2.63) and (2.64),

$$(\mathbf{D}^2\Phi_{[\mathbf{B}]})_{t_0\mathbf{D}}(\mathbf{D}, \mathbf{D}) = \mathbf{B}_\perp \mathbf{U} \{\boldsymbol{\Sigma} \odot \boldsymbol{\Sigma} \odot f_2(t_0\boldsymbol{\Sigma})\} \mathbf{V}^\top, \quad (2.65)$$

where \odot denotes the Hadamard product, and $f_2 : x \mapsto -2x/(1+x^2)^2$ is the second derivative of the arctan function (again, this is applied to each element of the matrix $t_0\boldsymbol{\Sigma}$), meaning $\boldsymbol{\Sigma} \odot \boldsymbol{\Sigma} \odot f_2(t_0\boldsymbol{\Sigma}) = \text{diag}(-2t_0\Sigma_i^3/(1+t_0^2\Sigma_i^2)^2)$. From [EAS98, page 314],

$$\left\| (\mathbf{D}^2\Phi_{[\mathbf{B}]})_{t_0\mathbf{D}}(\mathbf{D}, \mathbf{D}) \right\|_{[\mathbf{B}]}^2 = \text{Tr} \left(\{\boldsymbol{\Sigma} \odot \boldsymbol{\Sigma} \odot f_2(t_0\boldsymbol{\Sigma})\} \{\boldsymbol{\Sigma} \odot \boldsymbol{\Sigma} \odot f_2(t_0\boldsymbol{\Sigma})\}^\top \right),$$

where Tr denotes the trace. Then, using the fact that $|f_2(x)| \leq 1$ for any real x , the right-hand side is less than $\text{Tr}(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^\top)$, which is equal to $\|\mathbf{D}\|_{[\mathbf{B}]}^2$. Thus, Condition **R2** is satisfied with $\mathcal{L}^{(1)}(\theta) = 1$.

- Condition **R3-(i)** is satisfied: recall (2.49) from Lemma 2.17-(a). This states,

$$\mathbf{D}_t\dot{\gamma}(0) = (\mathbf{D}^2\Phi_{[\mathbf{B}]})_{0[\mathbf{B}]}(\mathbf{D}, \mathbf{D}).$$

Setting $t = 0$ in (2.65), it then follows

$$\mathbf{D}_t\dot{\gamma}(0) = \mathbf{B}_\perp \mathbf{U} (\boldsymbol{\Sigma} \odot \boldsymbol{\Sigma} \odot f_2(0)) \mathbf{V}^\top,$$

which is equal to zero since $f_2(0) = 0$.

- Condition **R3-(ii)** is satisfied: the proof is similar to the above one for **R2**. Here, instead of (2.64), it is enough to use

$$(\mathbf{D}^3\Phi_{[\mathbf{B}]})_{t_0\mathbf{D}}(\mathbf{D}, \mathbf{D}, \mathbf{D}) = \left. \frac{d^3}{dt^3} \Phi_{[\mathbf{B}]}(t\mathbf{D}) \right|_{t=t_0}$$

by computing the derivative, as in (2.65), it can be shown that **R3** is satisfied with $\mathcal{L}^{(2)}([\mathbf{B}]) = 2$.

2.7 Complementary proofs for Section 2.3.1

2.7.1 Proof of Proposition 2.18

In order to prove Proposition 2.18, we introduce the following lemma.

Lemma 2.20. *Let $(\theta, \mathbf{u}) \in \text{TS}^d$. We have the following identities:*

$$\cos(\arctan(\|\mathbf{u}\|)) = 1 / \sqrt{1 + \|\mathbf{u}\|^2}, \quad \sin(\arctan(\|\mathbf{u}\|)) = \|\mathbf{u}\| / \sqrt{1 + \|\mathbf{u}\|^2}. \quad (2.66)$$

Moreover, the retraction given by (2.56) can be written as:

$$\text{Ret}_\theta(\mathbf{u}) = \text{Exp}_\theta \left(\arctan(\|\mathbf{u}\|) \frac{\mathbf{u}}{\|\mathbf{u}\|} \right). \quad (2.67)$$

2.7. Complementary proofs for Section 2.3.1

Proof of Lemma 2.20. First, to prove both identities in (2.66), consider $\alpha = \arctan(\|\mathbf{u}\|)$. Combining $\tan(\alpha) = \|\mathbf{u}\|$ and $\tan^2(\alpha) = (1 - \cos^2(\alpha))/\cos^2(\alpha)$, we obtain $\cos^2(\alpha) = 1/(1 + \|\mathbf{u}\|^2)$. Since $\alpha \in [0, \pi/2)$, then $\cos(\alpha) \geq 0$ and $\sin(\alpha) \geq 0$. Therefore $\cos(\alpha) = 1/\sqrt{1 + \|\mathbf{u}\|^2}$. The second equality in (2.66) is a direct result of the first one, using $\sin^2(\alpha) + \cos^2(\alpha) = 1$ and $\sin(\alpha) \geq 0$.

Secondly, we compute $\text{Exp}_\theta(\arctan(\|\mathbf{u}\|)\mathbf{u}/\|\mathbf{u}\|)$ using (2.55) and (2.66), which gives

$$\text{Exp}_\theta\left(\arctan(\|\mathbf{u}\|)\frac{\mathbf{u}}{\|\mathbf{u}\|}\right) = \theta \left/ \sqrt{1 + \|\mathbf{u}\|^2} + \mathbf{u} \right/ \sqrt{1 + \|\mathbf{u}\|^2} . \quad (2.68)$$

To see that this is equal to $\text{Ret}_\theta(\mathbf{u})$, note that \mathbf{u} is orthogonal to θ because $\mathbf{u} \in T_\theta\mathbb{S}^d$, thus $\|\theta + \mathbf{u}\|^2 = \|\theta\|^2 + \|\mathbf{u}\|^2$. Moreover, $\|\theta\| = 1$. Then, (2.67) follows from (2.68) and (2.56). \square

2.7.2 Details for the Grassmann manifold

To give a complete and comprehensive description of the Grassmann manifold $\text{Gr}_r(\mathbb{R}^d)$ introduced in the Example 2.9, we follow the steps of [AMS04, Sections 2, 3.1, 3.3, 3.4 and 3.7].

Definition of the Grassmann manifold as a set For $r \in \{1, \dots, d-1\}$, consider the set of real matrices whose columns span an r -dimensional subspace of \mathbb{R}^d , *i.e.* with r independent columns of size d . We call this set the *non-compact Stiefel manifold*, and denote it by $\text{St}_r(\mathbb{R}^d) = \{\mathbf{B} \in \mathbb{R}^{d \times r} : \det(\mathbf{B}^\top \mathbf{B}) \neq 0\}$. It is a submanifold of $\mathbb{R}^{d \times r}$ as an open subset of it. The set of r -dimensional subspaces of \mathbb{R}^d is called the *Grassmann manifold*, denoted by $\text{Gr}_r(\mathbb{R}^d)$.

Each matrix in $\text{St}_r(\mathbb{R}^d)$ represents a point in $\text{Gr}_r(\mathbb{R}^d)$ by considering the linear subspace spanned by its columns. However, the representation is not unique. Indeed, given \mathbf{B} in $\text{St}_r(\mathbb{R}^d)$, the set of matrices that have the same span as \mathbf{B} is $\mathbf{B}\text{GL}_r(\mathbb{R}) = \{\mathbf{B}\mathbf{M} : \mathbf{M} \in \text{GL}_r(\mathbb{R})\}$, which we denote by $[\mathbf{B}]$, where $\text{GL}_r(\mathbb{R})$ is the set of r -by- r invertible real matrices.

Thus, we can identify the Grassmann manifold $\text{Gr}_r(\mathbb{R}^d)$ with the quotient space $\text{St}_r(\mathbb{R}^d)/\text{GL}_r(\mathbb{R})$, which is defined as $\{[\mathbf{B}] : \mathbf{B} \in \text{St}_r(\mathbb{R}^d)\}$. This quotient is a principal fiber bundle over $\text{Gr}_r(\mathbb{R}^d)$ with group $\text{GL}_r(\mathbb{R})$ [KN96, Chapter I, Section 5], endowed with the group action of $\text{GL}_r(\mathbb{R})$ on $\text{St}_r(\mathbb{R}^d)$ and the projection map π given by

$$\begin{cases} \text{St}_r(\mathbb{R}^d) \times \text{GL}_r(\mathbb{R}) & \longrightarrow & \text{St}_r(\mathbb{R}^d) \\ (\mathbf{B}, \mathbf{M}) & \longmapsto & \mathbf{B}\mathbf{M} \end{cases}, \text{ and } \pi : \begin{cases} \text{St}_r(\mathbb{R}^d) & \longrightarrow & \text{Gr}_r(\mathbb{R}^d) \\ \mathbf{B} & \longmapsto & \text{Span}(\mathbf{B}) \end{cases} .$$

The differentiable structure of the Grassmann manifold Next, to endow $\text{Gr}_r(\mathbb{R}^d)$ with a manifold structure, we give an explicit parameterization using cross sections. We consider $S_{\mathbf{B}}$ as the intersection of the Stiefel manifold with the affine subspace that is orthogonal to $\mathbf{B}\text{GL}_r(\mathbb{R})$ and that intersects $\mathbf{B}\text{GL}_r(\mathbb{R})$ at \mathbf{B} . These two conditions result in

$$S_{\mathbf{B}} = \left\{ \mathbf{Y} \in \text{St}_r(\mathbb{R}^d) : \mathbf{B}^\top (\mathbf{Y} - \mathbf{B}) = \mathbf{0}_{r \times r} \right\} .$$

Lemma 2.21 (Cross section). *Let $\mathbf{B} \in \text{St}_r(\mathbb{R}^d)$ and $\mathcal{U}_{\mathbf{B}} = \{[\mathbf{Y}] \in \text{Gr}_r(\mathbb{R}^d) : \det(\mathbf{B}^\top \mathbf{Y}) \neq 0\}$. Then the mapping*

$$\sigma_{\mathbf{B}} : \begin{cases} \mathcal{U}_{\mathbf{B}} & \longrightarrow & S_{\mathbf{B}} \\ [\mathbf{Y}] & \longmapsto & \mathbf{Y}(\mathbf{B}^\top \mathbf{Y})^{-1} \mathbf{B}^\top \mathbf{B} \end{cases}, \quad (2.69)$$

is well-defined and is a bijection between $\mathcal{U}_{\mathbf{B}} \subset \text{Gr}_r(\mathbb{R}^d)$ and $S_{\mathbf{B}} \subset \text{St}_r(\mathbb{R}^d)$. It also satisfies the property $\pi \circ \sigma_{\mathbf{B}} = \text{Id}$ on $\mathcal{U}_{\mathbf{B}}$. We call it the cross section map, since $S_{\mathbf{B}} \cap \mathbf{YGL}_r(\mathbb{R}) = \{\sigma_{\mathbf{B}}([\mathbf{Y}])\}$ if $[\mathbf{Y}] \in \mathcal{U}_{\mathbf{B}}$, and $S_{\mathbf{B}} \cap \mathbf{YGL}_r(\mathbb{R}) = \emptyset$ otherwise.

Proof. First, we can see that if a fiber $[\mathbf{Y}]$, for \mathbf{Y} in $\text{St}_r(\mathbb{R}^d)$, intersects the section $S_{\mathbf{B}}$, then the intersection is unique. Indeed, if there exists $\mathbf{M} \in \text{GL}_r(\mathbb{R})$ such that $\mathbf{Y} \in S_{\mathbf{B}}$ and $\mathbf{Y}\mathbf{M} \in S_{\mathbf{B}}$, then $\mathbf{B}^\top (\mathbf{Y} - \mathbf{Y}\mathbf{M}) = \mathbf{0}$, meaning $\mathbf{B}^\top \mathbf{Y} = \mathbf{B}^\top \mathbf{Y}\mathbf{M}$ and the same computations for $\mathbf{Y}\mathbf{M}$ bring $\mathbf{B}^\top \mathbf{Y} = \mathbf{B}^\top \mathbf{Y}\mathbf{M}$. In conclusion $\mathbf{M} = \mathbf{I}_r$.

Secondly, we look at which fibers intersect the cross section $S_{\mathbf{B}}$. Let \mathbf{Y} in $\text{St}_r(\mathbb{R}^d)$, if $\mathbf{B}^\top \mathbf{Y}$ is not invertible, then for any $\mathbf{M} \in \text{GL}_r(\mathbb{R})$, $\det(\mathbf{B}^\top \mathbf{Y}\mathbf{M}) = 0$. Since $\det(\mathbf{B}^\top \mathbf{B}) \neq 0$ by definition, $\mathbf{B}^\top \mathbf{Y}\mathbf{M} \neq \mathbf{B}^\top \mathbf{B}$, meaning $\mathbf{B}^\top (\mathbf{Y}\mathbf{M} - \mathbf{B}) \neq \mathbf{0}_{r \times r}$. In this case, the fiber $[\mathbf{Y}]$ does not intersect the cross section $S_{\mathbf{B}}$. On the other hand, if $\mathbf{B}^\top \mathbf{Y}$ is invertible, then taking $\mathbf{M} = (\mathbf{B}^\top \mathbf{Y})^{-1} \mathbf{B}^\top \mathbf{B} \in \text{GL}_r(\mathbb{R})$ brings $\mathbf{B}^\top \mathbf{Y}\mathbf{M} = \mathbf{B}^\top \mathbf{Y}(\mathbf{B}^\top \mathbf{Y})^{-1} \mathbf{B}^\top \mathbf{B} = \mathbf{B}^\top \mathbf{B}$ and in conclusion $\mathbf{Y}\mathbf{M} \in S_{\mathbf{B}}$. These two points bring $S_{\mathbf{B}} \cap \mathbf{YGL}_r(\mathbb{R}) = \{\sigma_{\mathbf{B}}([\mathbf{Y}])\}$ and the fact that $\sigma_{\mathbf{B}}$ is well-defined.

Finally, $\sigma_{\mathbf{B}}$ is a bijection because π is its inverse on $S_{\mathbf{B}}$. \square

The classical manifold structure of $\text{Gr}_r(\mathbb{R}^d)$ [AMS04] is then obtained by considering the atlas formed by the following family of parameterizations.

Proposition 2.22. *Let $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ and $\mathbf{B}_\perp \in \text{St}_{d-r}(\mathbb{R}^d)$, with $\mathbf{B}^\top \mathbf{B}_\perp = \mathbf{0}_{r \times (d-r)}$. The mapping defined by*

$$\varphi_{(\mathbf{B}, \mathbf{B}_\perp)} : \begin{cases} \mathbb{R}^{(d-r) \times r} & \longrightarrow & \mathcal{U}_{\mathbf{B}} \\ \mathbf{C} & \longmapsto & [\mathbf{B} + \mathbf{B}_\perp \mathbf{C}] \end{cases},$$

is invertible. Consider $\mathcal{A} = \{(\mathbf{B}, \mathbf{B}_\perp) : [\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d), \mathbf{B}_\perp \in \text{St}_{d-r}(\mathbb{R}^d), \mathbf{B}^\top \mathbf{B}_\perp = \mathbf{0}_{r \times (d-r)}\}$. The family $(\mathcal{U}_{\mathbf{B}}, \varphi_\alpha^{-1})_{\alpha \in \mathcal{A}}$, is a smooth atlas of $\text{Gr}_r(\mathbb{R}^d)$.

Proof. Let $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ and $\mathbf{B}_\perp \in \text{St}_{d-r}(\mathbb{R}^d)$, with $\mathbf{B}^\top \mathbf{B}_\perp = \mathbf{0}_{r \times (d-r)}$. Let $\alpha = (\mathbf{B}, \mathbf{B}_\perp)$. We begin by proving that φ_α is invertible. Let $\mathbf{C} \in \mathbb{R}^{(d-r) \times r}$, since $\mathbf{B}^\top (\mathbf{B} + \mathbf{B}_\perp \mathbf{C}) = \mathbf{B}^\top \mathbf{B}$ is invertible, $\varphi_\alpha(\mathbf{C})$ is in $\mathcal{U}_{\mathbf{B}}$. We use the cross-section $\sigma_{\mathbf{B}}$ to compute the inverse of φ_α ,

$$\begin{aligned} \sigma_{\mathbf{B}} \circ \varphi_\alpha(\mathbf{C}) &= (\mathbf{B} + \mathbf{B}_\perp \mathbf{C}) \left(\mathbf{B}^\top (\mathbf{B} + \mathbf{B}_\perp \mathbf{C}) \right)^{-1} \mathbf{B}^\top \mathbf{B} \\ &= (\mathbf{B} + \mathbf{B}_\perp \mathbf{C}) \left(\mathbf{B}^\top \mathbf{B} \right)^{-1} \mathbf{B}^\top \mathbf{B} \\ &= \mathbf{B} + \mathbf{B}_\perp \mathbf{C}. \end{aligned}$$

Therefore, it is easy to check that the mapping defined by

$$\begin{cases} \mathcal{U}_{\mathbf{B}} & \longrightarrow & \mathbb{R}^{(d-r) \times r} \\ [\mathbf{Y}] & \longmapsto & (\mathbf{B}_{\perp}^{\top} \mathbf{B}_{\perp})^{-1} \mathbf{B}_{\perp}^{\top} (\sigma_{\mathbf{B}}([\mathbf{Y}]) - \mathbf{B}) \end{cases}$$

is the inverse of φ_{α} .

To check the smooth compatibility of the parameterizations, take $\alpha, \beta \in \mathcal{A}$, with $\alpha = (\mathbf{B}, \mathbf{B}_{\perp}), \beta = (\mathbf{Y}, \mathbf{Y}_{\perp})$ and compute $\varphi_{\alpha}^{-1} \circ \varphi_{\beta}$. The domain of $\varphi_{\alpha}^{-1} \circ \varphi_{\beta}$ is exactly $\mathcal{E}_{\alpha, \beta} = \{\mathbf{C} \in \mathbb{R}^{(d-r) \times r} : \varphi_{\beta}(\mathbf{C}) \in \mathcal{U}_{\mathbf{B}}\}$, which can be rewritten as $\mathcal{E}_{\alpha, \beta} = \{\mathbf{C} \in \mathbb{R}^{(d-r) \times r} : \det(\mathbf{B}^{\top}(\mathbf{Y} + \mathbf{Y}_{\perp} \mathbf{C})) \neq 0\}$. For any $\mathbf{C} \in \mathcal{E}_{\alpha, \beta}$,

$$\begin{aligned} \varphi_{\alpha}^{-1} \circ \varphi_{\beta}(\mathbf{C}) &= (\mathbf{B}_{\perp}^{\top} \mathbf{B}_{\perp})^{-1} \mathbf{B}_{\perp}^{\top} (\sigma_{\mathbf{B}}(\varphi_{\beta}(\mathbf{C})) - \mathbf{B}) \\ &= (\mathbf{B}_{\perp}^{\top} \mathbf{B}_{\perp})^{-1} \mathbf{B}_{\perp}^{\top} \left\{ (\mathbf{Y} + \mathbf{Y}_{\perp} \mathbf{C}) (\mathbf{B}^{\top} \mathbf{Y} + \mathbf{B}^{\top} \mathbf{Y}_{\perp} \mathbf{C})^{-1} \mathbf{B}^{\top} \mathbf{B} - \mathbf{B} \right\}. \end{aligned}$$

We observe that $\varphi_{\alpha}^{-1} \circ \varphi_{\beta}$ is a smooth function on its domain $\mathcal{E}_{\alpha, \beta}$, thus proving the smooth compatibility.

Finally, we need to prove that the images of the parameterizations cover the whole manifold. This is easy to check since $[\mathbf{B}] \in \mathcal{U}_{\mathbf{B}}$, therefore $\text{Gr}_r(\mathbb{R}^d) \subset \cup_{(\mathbf{B}, \mathbf{B}_{\perp}) \in \mathcal{A}} \mathcal{U}_{\mathbf{B}}$. This concludes the proof. \square

This choice of atlas for $\text{Gr}_r(\mathbb{R}^d)$ makes the cross-section mappings $(\sigma_{\mathbf{B}})_{\mathbf{B} \in \text{St}_r(\mathbb{R}^d)}$ smooth by construction.

Describing the tangent space Let $\mathbf{B} \in \text{St}_r(\mathbb{R}^d)$. The tangent space $\text{T}_{\mathbf{B}} \text{St}_r(\mathbb{R}^d) = \mathbb{R}^{d \times r}$ can be decomposed according to the quotient $\text{Gr}_r(\mathbb{R}^d) = \text{St}_r(\mathbb{R}^d) / \text{GL}_r(\mathbb{R})$ by a connection in the principal fibre bundle, in the sense of [KN96, Chapter II, Section 1]. First, the vertical space $\mathbf{V}_{\mathbf{B}} \subset \text{T}_{\mathbf{B}} \text{St}_r(\mathbb{R}^d)$ is defined as the tangent vectors that do not modify $[\mathbf{B}]$, namely the tangent space to the fiber $[\mathbf{B}] = \mathbf{B} \text{GL}_r(\mathbb{R})$,

$$\mathbf{V}_{\mathbf{B}} = \{\dot{\gamma}(0) \mid \gamma : [0, 1] \rightarrow \mathbf{B} \text{GL}_r(\mathbb{R}) \text{ such that } \gamma \text{ is smooth}\}.$$

Proposition 2.23. *Let $\mathbf{B} \in \text{St}_r(\mathbb{R}^d)$. The vertical space at \mathbf{B} can be written as*

$$\mathbf{V}_{\mathbf{B}} = \{\mathbf{B}\mathbf{C} : \mathbf{C} \in \mathbb{R}^{r \times r}\}, \quad (2.70)$$

and we also have $\text{Ker}((D\pi)_{\mathbf{B}}) = \mathbf{V}_{\mathbf{B}}$, where $(D\pi)_{\mathbf{B}} : \text{T}_{\mathbf{B}} \text{St}_r(\mathbb{R}^d) \rightarrow \text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ is the derivative of the canonical projection $\pi : \text{St}_r(\mathbb{R}^d) \rightarrow \text{Gr}_r(\mathbb{R}^d)$ at \mathbf{B} .

Proof. By the definition of the vertical space, for any $t \in [0, 1]$, there exists $\mathbf{M}_t \in \text{GL}_r(\mathbb{R})$ such that $\gamma(t) = \mathbf{B}\mathbf{M}_t$. Since γ is smooth, so is $t \mapsto (\mathbf{B}^{\top} \mathbf{B})^{-1} \mathbf{B}^{\top} \gamma(t) = \mathbf{M}_t$. Computing the derivative gives $\dot{\gamma}(t) = \mathbf{B}\dot{\mathbf{M}}_t$. This gives the characterization $\mathbf{V}_{\mathbf{B}} = \{\mathbf{B}\dot{\gamma}(0) \mid \gamma : [0, 1] \rightarrow \text{GL}_r(\mathbb{R}) \text{ such that } \gamma \text{ is smooth}\}$. Since $\text{GL}_r(\mathbb{R})$ is an open set of $\mathbb{R}^{r \times r}$, its tangent space is trivial, therefore (2.70) holds.

Secondly, since π is smooth [Lee03, Theorem 7.10], by differentiating $[\mathbf{B}] = \pi \circ \gamma$ we get,

$$\mathbf{0}_{T_{[\mathbf{B}]}Gr_r(\mathbb{R}^d)} = \frac{d(\pi \circ \gamma)}{dt}(t) = (D\pi)_{\gamma(t)} \circ \dot{\gamma}(t),$$

where the last equality comes from the chain rule. This gives that $\mathbf{V}_{\mathbf{B}} \subset \text{Ker}((D\pi)_{\mathbf{B}})$. Finally, by [Lee03, Theorem 7.10], $d(T_{[\mathbf{B}]}Gr_r(\mathbb{R}^d)) = d(T_{\mathbf{B}}St_r(\mathbb{R}^d)) - d(GL_r)$, where $d(GL_r) = r^2$, and $d(T_{\mathbf{B}}St_r(\mathbb{R}^d)) = dr$. Therefore $d(T_{[\mathbf{B}]}Gr_r(\mathbb{R}^d)) = (d-r)r$ and by the rank-nullity theorem applied to $(D\pi)_{\mathbf{B}}$, $d(\text{Ker}((D\pi)_{\mathbf{B}})) = r^2$. This gives $\mathbf{V}_{\mathbf{B}} = \text{Ker}((D\pi)_{\mathbf{B}})$ using (2.70). \square

Definition 2.24 (Connection in the sense of [KN96]). *A connection in the sense of [KN96, Chapter II, Section 1] is the assignment of a subspace $\mathbf{H}_{\mathbf{B}}$ of $T_{\mathbf{B}}St_r(\mathbb{R}^d)$, called horizontal space, to each $\mathbf{B} \in St_r(\mathbb{R}^d)$ such that: (i) $T_{\mathbf{B}}St_r(\mathbb{R}^d) = \mathbf{V}_{\mathbf{B}} \oplus \mathbf{H}_{\mathbf{B}}$, (ii) the group action of $GL_r(\mathbb{R})$ is compatible with the assignment, meaning that for any $\mathbf{M} \in GL_r(\mathbb{R})$, $\mathbf{H}_{\mathbf{B}\mathbf{M}} = \mathbf{H}_{\mathbf{B}}\mathbf{M}$, where $\mathbf{H}_{\mathbf{B}}\mathbf{M} = \{\mathbf{D}\mathbf{M} : \mathbf{D} \in \mathbf{H}_{\mathbf{B}}\}$; (iii) the assignment is smooth, meaning that if \mathfrak{X} is a smooth vector field on $St_r(\mathbb{R}^d)$, then its horizontal and vertical projections are also smooth vector fields.*

Let us verify now that the assignment

$$\mathbf{B} \mapsto \mathbf{H}_{\mathbf{B}}, \text{ where } \mathbf{H}_{\mathbf{B}} = \{\mathbf{D} \in T_{\mathbf{B}}St_r(\mathbb{R}^d) : \mathbf{B}^{\top}\mathbf{D} = \mathbf{0}_{r \times r}\}, \quad (2.71)$$

is a connection in the sense of [KN96].

To check (i), use that $\mathbf{B}^{\top}\mathbf{B}$ is invertible and the rank-nullity theorem for the map $\mathbf{D} \mapsto \mathbf{B}^{\top}\mathbf{D}$ on $\mathbb{R}^{d \times r}$, which gives $d(\mathbf{H}_{\mathbf{B}}) = (d-r)r$ as well as $\mathbf{H}_{\mathbf{B}} \cap \mathbf{V}_{\mathbf{B}} = \{\mathbf{0}_{d \times r}\}$. To prove (ii), see that for any $\mathbf{D} \in \mathbf{H}_{\mathbf{B}}$ and any $\mathbf{M} \in GL_r(\mathbb{R})$, $\mathbf{B}^{\top}\mathbf{D} = \mathbf{M}^{\top}\mathbf{B}^{\top}\mathbf{D}\mathbf{M} = \mathbf{0}_{r \times r}$ and therefore $\mathbf{H}_{\mathbf{B}} = \mathbf{H}_{\mathbf{B}\mathbf{M}} = \mathbf{H}_{\mathbf{B}}\mathbf{M}$. Finally, (iii) is verified because the horizontal space is orthogonal to the vertical space. Indeed, for any $\mathbf{D} \in T_{\mathbf{B}}St_r(\mathbb{R}^d)$, its vertical projection is $\mathbf{B}(\mathbf{B}^{\top}\mathbf{B})^{-1}\mathbf{B}^{\top}\mathbf{D} \in \mathbf{V}_{\mathbf{B}}$ and its horizontal projection is $(\mathbf{I}_d - \mathbf{B}(\mathbf{B}^{\top}\mathbf{B})^{-1}\mathbf{B}^{\top})\mathbf{D} \in \mathbf{H}_{\mathbf{B}}$. Since $\mathbf{B} \mapsto \mathbf{B}(\mathbf{B}^{\top}\mathbf{B})^{-1}\mathbf{B}^{\top}$ is smooth on $St_r(\mathbb{R}^d)$, for any smooth vector field \mathfrak{X} on $St_r(\mathbb{R}^d)$, its vertical projection $\mathbf{B} \mapsto \mathbf{B}(\mathbf{B}^{\top}\mathbf{B})^{-1}\mathbf{B}^{\top}\mathfrak{X}(\mathbf{B})$ is smooth, and so is therefore its horizontal projection, as a difference of smooth vector fields.

Since we have a connection in the sense of [KN96], using Proposition 2.23 and that $\mathbf{H}_{\mathbf{B}}$ is a supplement of $\mathbf{V}_{\mathbf{B}} = \text{Ker}((D\pi)_{\mathbf{B}})$, we can conclude that $(D\pi)_{\mathbf{B}}$ is an isomorphism between $\mathbf{H}_{\mathbf{B}}$ and $T_{[\mathbf{B}]}Gr_r(\mathbb{R}^d)$. This allows us to define the horizontal lift of a vector $\xi \in T_{[\mathbf{B}]}Gr_r(\mathbb{R}^d)$: there exists a unique vector $\mathbf{D} \in \mathbf{H}_{\mathbf{B}}$ such that $(D\pi)_{\mathbf{B}}(\mathbf{D}) = \xi$. The same can be done with a vector field: for any smooth vector field \mathfrak{X} on $Gr_r(\mathbb{R}^d)$, there exists a unique smooth vector field $\tilde{\mathfrak{X}}$ on $St_r(\mathbb{R}^d)$ such that for any $\mathbf{B} \in St_r(\mathbb{R}^d)$ $\tilde{\mathfrak{X}}(\mathbf{B})$ is horizontal, meaning $\tilde{\mathfrak{X}}(\mathbf{B}) \in \mathbf{H}_{\mathbf{B}}$ and that projects into \mathfrak{X} , meaning that for any $\mathbf{B} \in St_r(\mathbb{R}^d)$ we have $(D\pi)_{\mathbf{B}}(\tilde{\mathfrak{X}}(\mathbf{B})) = \mathfrak{X}(\mathbf{B})$, according to principal fiber bundle theory [KN96, Chapter II, Section 1, Proposition 1.2]. To see how the horizontal lift changes when the representative \mathbf{B} changes into $\mathbf{B}\mathbf{M}$, where \mathbf{M} is a $r \times r$ -invertible matrix, let $\mathbf{B}_t : [0, 1] \mapsto St_r(\mathbb{R}^d)$ be a smooth curve, with $\mathbf{B} = \mathbf{B}_0$ and $\mathbf{D} = \dot{\mathbf{B}}_0$. For any $t \in [0, 1]$, we have $\pi(\mathbf{B}_t) = \pi(\mathbf{B}_t\mathbf{M})$ and by differentiating this equality, we obtain $(D\pi)_{\mathbf{B}_t}(\dot{\mathbf{B}}_t) = (D\pi)_{\mathbf{B}_t\mathbf{M}}(\dot{\mathbf{B}}_t\mathbf{M})$ by using

that the derivative of the curve $t \mapsto \mathbf{B}_t \mathbf{M}$ is $t \mapsto \dot{\mathbf{B}}_t \mathbf{M}$. Evaluating this equation at $t = 0$ brings $(D\pi)_{\mathbf{B}} \mathbf{D} = (D\pi)_{\mathbf{B}} \mathbf{M} (\mathbf{D}\mathbf{M})$. This result was proven in [AMS04, Proposition 3.1] using the cross section $\sigma_{\mathbf{B}}$ introduced in Lemma 2.21.

However $(D\pi)_{\mathbf{B}}(\mathbf{D})$ does not have a matrix form, thus we use the matrix \mathbf{D} for any computation (e.g. the metric or Riemannian exponential). Therefore, for ease of notation, for any $\mathbf{D} \in \mathbf{H}_{\mathbf{B}}$, we denote $(D\pi)_{\mathbf{B}} \mathbf{D}$ by \mathbf{D} when there is no confusion.

Action of the orthogonal group on the Grassmann manifold Recall that the orthogonal group $O_d(\mathbb{R})$ acts on $\text{Gr}_r(\mathbb{R}^d)$, in a way that is compatible with the projection $\pi : \text{St}_r(\mathbb{R}^d) \rightarrow \text{Gr}_r(\mathbb{R}^d)$, meaning that for any $\mathbf{O} \in O_d(\mathbb{R})$ and $\mathbf{B} \in \text{St}_r(\mathbb{R}^d)$, the action of \mathbf{O} over $\pi(\mathbf{B})$ is given by

$$\mathbf{O} \cdot \pi(\mathbf{B}) = \pi(\mathbf{O}\mathbf{B}). \quad (2.72)$$

Moreover, the orthogonal group also acts on the tangent bundle, in a way that is compatible with the derivative of the projection. For any $\mathbf{D} \in \mathbf{T}_{\mathbf{B}}\text{St}_r(\mathbb{R}^d)$, the action of \mathbf{O} over the tangent vector $(D\pi)_{\mathbf{B}}(\mathbf{D})$ is given by

$$\mathbf{O} \cdot (D\pi)_{\mathbf{B}}(\mathbf{D}) = (D\pi)_{\mathbf{O}\cdot\mathbf{B}}(\mathbf{O}\mathbf{D}). \quad (2.73)$$

Alternative formulation of the Riemannian exponential We consider the following definition: for any integers $n, m, p, q \in \mathbb{N} \setminus \{0\}$ and any matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$, denote the reshaping (adding zeros if necessary) of \mathbf{A} to a $p \times q$ matrix by $\mathbf{A}_{[p,q]}$, where its coefficient $(i, j) \in \{1, \dots, p\} \times \{1, \dots, q\}$ is given by

$$(\mathbf{A}_{[p,q]})_{i,j} = \begin{cases} \mathbf{A}_{i,j} & \text{if } i \leq n \text{ and } j \leq m \\ 0 & \text{otherwise} \end{cases}. \quad (2.74)$$

We simply denote $\mathbf{A}_{[p,p]}$ by $\mathbf{A}_{[p]}$.

We provide a lemma that gives another formula for the Riemannian exponential on $\text{Gr}_r(\mathbb{R}^d)$, generalizing the formulation of [EAS98, Theorem 2.3] for non-orthonormal representations of the manifold.

Lemma 2.25. *Let $s = d - r$ and denote $r \wedge s = \min(r, s)$.*

(a) *Let \mathbf{N} be a $(r \wedge s) \times (r \wedge s)$ real diagonal matrix then, we have that*

$$\exp \begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} = \begin{pmatrix} \cos(\mathbf{N}_{[r]}) \\ \sin(\mathbf{N}_{[s,r]}) \end{pmatrix}, \quad (2.75)$$

where \exp is the usual matrix exponential defined by its power series, the matrices $\mathbf{N}_{[.,.]}$ are given by the definition (2.74) and \sin and \cos are applied componentwise to the diagonal of each matrix.

(b) Let $\mathbf{B} \in \text{St}_r(\mathbb{R}^d)$, $\mathbf{B}_\perp \in \text{St}_s(\mathbb{R}^d)$, $\mathbf{D} \in \text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ with $\mathbf{B}^\top \mathbf{B}_\perp = \mathbf{0}_{r \times s}$ and $\mathbf{D} = \mathbf{B}_\perp \mathbf{C}$, where $\mathbf{C} \in \mathbb{R}^{s \times r}$. The Riemannian exponential, given by (2.62), can be written as

$$\begin{aligned} & \text{Exp}_{[\mathbf{B}]}(\mathbf{D}) \\ &= \left[\begin{pmatrix} \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} & \mathbf{B}_\perp \end{pmatrix} \exp \begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{C}^\top \\ \mathbf{C}(\mathbf{B}^\top \mathbf{B})^{-1/2} & \mathbf{0}_{(d-r) \times (d-r)} \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{(d-r) \times r} \end{pmatrix} \right]. \end{aligned} \quad (2.76)$$

Proof. (a) First, denote the diagonal of \mathbf{N} by $(\mathbf{N}_i)_{i \in \{1, \dots, r \wedge s\}}$, let $\mathbf{e}_{j,k}$ be the $d \times d$ matrix all of whose elements are zero, except the one at row j and column k , which is equal to 1. For $i \in \{1, \dots, r \wedge s\}$, let $\mathbf{b}_{r+i} = \mathbf{e}_{r+i,i} - \mathbf{e}_{i,r+i}$. We separate the cases $r \leq s$ and $r > s$ to explicitly write the matrices at hand. In the case $r \leq s$,

$$\begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} = \begin{pmatrix} \mathbf{0}_{r \times r} & -\mathbf{N} & \mathbf{0}_{r \times (s-r)} \\ \mathbf{N} & \mathbf{0}_{r \times r} & \mathbf{0}_{r \times (s-r)} \\ \mathbf{0}_{(s-r) \times r} & \mathbf{0}_{(s-r) \times r} & \mathbf{0}_{(s-r) \times (s-r)} \end{pmatrix} = \sum_{i=1}^r \mathbf{N}_i \mathbf{b}_{r+i}.$$

And in the case $r > s$,

$$\begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} = \begin{pmatrix} \mathbf{0}_{s \times s} & \mathbf{0}_{s \times (r-s)} & -\mathbf{N} \\ \mathbf{0}_{(r-s) \times s} & \mathbf{0}_{(r-s) \times (r-s)} & \mathbf{0}_{(r-s) \times s} \\ \mathbf{N} & \mathbf{0}_{s \times (r-s)} & \mathbf{0}_{s \times s} \end{pmatrix} = \sum_{i=1}^s \mathbf{N}_i \mathbf{b}_{r+i}.$$

In both cases we obtain,

$$\begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} = \sum_{i=1}^{r \wedge s} \mathbf{N}_i \mathbf{b}_{r+i}. \quad (2.77)$$

Using the computation, for any $i, j, k, l \in \{1, \dots, d\}$, $\mathbf{e}_{i,j} \mathbf{e}_{k,l} = \delta_{j,k} \mathbf{e}_{i,l}$, one easily checks that for any $i, j \in \{1, \dots, r \wedge s\}$, $\mathbf{b}_{r+i} \mathbf{b}_{r+j} = -\delta_{i,j} (\mathbf{e}_{i,j} + \mathbf{e}_{r+i,r+j})$, thus the matrices $(\mathbf{b}_{r+i})_{i \in \{1, \dots, r\}}$ commute with each other. Therefore, (2.77) implies

$$\exp \begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} = \prod_{i=1}^{r \wedge s} \exp(\mathbf{N}_i \mathbf{b}_{r+i}). \quad (2.78)$$

Setting $\mathbf{a}_{r+i} = \mathbf{e}_{i,i} + \mathbf{e}_{r+i,r+i}$, and using that $\mathbf{b}_{r+i}^2 = -\mathbf{a}_{r+i}$, $\mathbf{b}_{r+i}^3 = -\mathbf{b}_{r+i}$, we can check that for any $k \in \mathbb{N}^*$, $\mathbf{b}_{r+i}^{2k} = (-1)^k \mathbf{a}_{r+i}$ and $\mathbf{b}_{r+i}^{2k+1} = (-1)^k \mathbf{b}_{r+i}$. Thus, the power series expansion of $\exp(\mathbf{N}_i \mathbf{b}_{r+i})$ is computed in the same way as the one for $\exp(i\theta) = \cos(\theta) + i \sin(\theta)$, where i is the imaginary unit. Namely,

$$\exp(\mathbf{N}_i \mathbf{b}_{r+i}) = \mathbf{I}_d + (\cos(\mathbf{N}_i) - 1) \mathbf{a}_{r+i} + \sin(\mathbf{N}_i) \mathbf{b}_{r+i},$$

where \mathbf{I}_d is the $d \times d$ identity matrix. Since we have, for $i, k \in \{1, \dots, r \wedge s\}$, $i \neq k$,

$$\begin{aligned} \mathbf{a}_{r+i} \mathbf{a}_{r+k} &= \mathbf{0}_{d \times d}, & \mathbf{a}_{r+i} \mathbf{b}_{r+k} &= \mathbf{0}_{d \times d}, \\ \mathbf{b}_{r+i} \mathbf{a}_{r+k} &= \mathbf{0}_{d \times d}, & \mathbf{b}_{r+i} \mathbf{b}_{r+k} &= \mathbf{0}_{d \times d}, \end{aligned}$$

2.7. Complementary proofs for Section 2.3.1

we get,

$$\prod_{i=1}^{r \wedge s} \exp(\mathbf{N}_i \mathbf{b}_{r+i}) = \mathbf{I}_d + \sum_{i=1}^{r \wedge s} [(\cos(\mathbf{N}_i) - 1) \mathbf{a}_{r+i} + \sin(\mathbf{N}_i) \mathbf{b}_{r+i}] .$$

Now, combining the previous equation with (2.78) brings, in the case $r \leq s$,

$$\begin{aligned} \exp \begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} &= \mathbf{I}_d + \begin{pmatrix} \cos(\mathbf{N}) - \mathbf{I}_r & \mathbf{0}_{r \times r} & \mathbf{0}_{r \times (s-r)} \\ \mathbf{0}_{r \times r} & \cos(\mathbf{N}) - \mathbf{I}_r & \mathbf{0}_{r \times (s-r)} \\ \mathbf{0}_{(s-r) \times r} & \mathbf{0}_{(s-r) \times r} & \mathbf{0}_{(s-r) \times (s-r)} \end{pmatrix} \\ &+ \begin{pmatrix} \mathbf{0}_{r \times r} & -\sin(\mathbf{N}) & \mathbf{0}_{r \times (s-r)} \\ \sin(\mathbf{N}) & \mathbf{0}_{r \times r} & \mathbf{0}_{r \times (s-r)} \\ \mathbf{0}_{(s-r) \times r} & \mathbf{0}_{(s-r) \times r} & \mathbf{0}_{(s-r) \times (s-r)} \end{pmatrix} , \\ &= \begin{pmatrix} \cos(\mathbf{N}) & -\sin(\mathbf{N}) & \mathbf{0}_{r \times (s-r)} \\ \sin(\mathbf{N}) & \cos(\mathbf{N}) & \mathbf{0}_{r \times (s-r)} \\ \mathbf{0}_{(s-r) \times r} & \mathbf{0}_{(s-r) \times r} & \mathbf{I}_{s-r} \end{pmatrix} . \end{aligned}$$

Now, extracting the first r columns when $r \leq s$ implies,

$$\exp \begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} = \begin{pmatrix} \cos(\mathbf{N}) \\ \sin(\mathbf{N}) \\ \mathbf{0}_{(s-r) \times r} \end{pmatrix} = \begin{pmatrix} \cos(\mathbf{N}_{[r]}) \\ \sin(\mathbf{N}_{[s,r]}) \end{pmatrix} . \quad (2.79)$$

In the case $r > s$, we have,

$$\begin{aligned} \exp \begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} &= \mathbf{I}_d + \begin{pmatrix} \cos(\mathbf{N}) - \mathbf{I}_s & \mathbf{0}_{s \times (r-s)} & \mathbf{0}_{s \times s} \\ \mathbf{0}_{(r-s) \times s} & \mathbf{0}_{(r-s) \times (r-s)} & \mathbf{0}_{(r-s) \times s} \\ \mathbf{0}_{s \times s} & \mathbf{0}_{s \times (r-s)} & \cos(\mathbf{N}) - \mathbf{I}_s \end{pmatrix} \\ &+ \begin{pmatrix} \mathbf{0}_{s \times s} & \mathbf{0}_{s \times (r-s)} & -\sin(\mathbf{N}) \\ \mathbf{0}_{(r-s) \times s} & \mathbf{0}_{(r-s) \times (r-s)} & \mathbf{0}_{(r-s) \times s} \\ \sin(\mathbf{N}) & \mathbf{0}_{s \times (r-s)} & \mathbf{0}_{s \times s} \end{pmatrix} , \\ &= \begin{pmatrix} \cos(\mathbf{N}) & \mathbf{0}_{s \times (r-s)} & -\sin(\mathbf{N}) \\ \mathbf{0}_{(r-s) \times s} & \mathbf{I}_{r-s} & \mathbf{0}_{(r-s) \times s} \\ \sin(\mathbf{N}) & \mathbf{0}_{s \times (r-s)} & \cos(\mathbf{N}) \end{pmatrix} . \end{aligned}$$

Again, extracting the first r columns when $r > s$ brings,

$$\exp \begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{N}_{[s,r]})^\top \\ \mathbf{N}_{[s,r]} & \mathbf{0}_{s \times s} \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} = \begin{pmatrix} \cos(\mathbf{N}) & \mathbf{0}_{s \times (r-s)} \\ \mathbf{0}_{(r-s) \times s} & \mathbf{I}_{r-s} \\ \sin(\mathbf{N}) & \mathbf{0}_{s \times (r-s)} \end{pmatrix} = \begin{pmatrix} \cos(\mathbf{N}_{[r]}) \\ \sin(\mathbf{N}_{[s,r]}) \end{pmatrix} , \quad (2.80)$$

which concludes the proof of Lemma 2.25-(a) by combining (2.79)-(2.80).

(b) Let $\mathbf{B} \in \text{St}_r(\mathbb{R}^d)$, $\mathbf{B}_\perp \in \text{St}_s(\mathbb{R}^d)$ and $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ with $\mathbf{B}^\top \mathbf{B}_\perp = \mathbf{0}_{r \times s}$, $\mathbf{B}_\perp^\top \mathbf{B}_\perp = \mathbf{I}_s$ and $\mathbf{D} = \mathbf{B}_\perp \mathbf{C}$, where $\mathbf{C} \in \mathbb{R}^{s \times r}$. We write the singular value decomposition $\mathbf{C}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{V}}^\top$ with $\tilde{\mathbf{U}} \in \text{O}_s(\mathbb{R})$, $\tilde{\mathbf{V}} \in \text{O}_r(\mathbb{R})$ and $\tilde{\Sigma}$ is a $s \times r$ non-negative diagonal matrix. Define the following matrices

$$\mathbf{H} = \begin{pmatrix} \tilde{\mathbf{V}} & \mathbf{0}_{r \times (d-r)} \\ \mathbf{0}_{(d-r) \times r} & \tilde{\mathbf{U}} \end{pmatrix},$$

$$\hat{\mathbf{C}} = \begin{pmatrix} \mathbf{0}_{r \times r} & -(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{C}^\top \\ \mathbf{C}(\mathbf{B}^\top \mathbf{B})^{-1/2} & \mathbf{0}_{s \times s} \end{pmatrix}, \quad \hat{\Sigma} = \begin{pmatrix} \mathbf{0}_{r \times r} & -\tilde{\Sigma}^\top \\ \tilde{\Sigma} & \mathbf{0}_{s \times s} \end{pmatrix}.$$

Remark that \mathbf{H} is a $d \times d$ orthogonal matrix, meaning that $\mathbf{H}^\top = \mathbf{H}^{-1}$, and with these notations, we have that $\mathbf{H} \hat{\Sigma} \mathbf{H}^\top = \hat{\mathbf{C}}$.

We start computing the right-hand term in (2.76) by rewriting the matrix exponential,

$$\exp(\hat{\mathbf{C}}) = \sum_{k=0}^{+\infty} \frac{\hat{\mathbf{C}}^k}{k!} = \sum_{k=0}^{+\infty} \frac{\mathbf{H} \hat{\Sigma}^k \mathbf{H}^{-1}}{k!} = \mathbf{H} \left(\sum_{k=0}^{+\infty} \frac{\hat{\Sigma}^k}{k!} \right) \mathbf{H}^{-1} = \mathbf{H} \exp(\hat{\Sigma}) \mathbf{H}^{-1}. \quad (2.81)$$

Next, we compute

$$\mathbf{H}^{-1} \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} = \begin{pmatrix} \tilde{\mathbf{V}}^\top \\ \mathbf{0}_{s \times r} \end{pmatrix} = \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} \tilde{\mathbf{V}}^\top. \quad (2.82)$$

We are now ready to compute the right-hand side of (2.76), using $\mathbf{H} \hat{\Sigma} \mathbf{H}^\top = \hat{\mathbf{C}}$, (2.81), (2.82) and (2.75):

$$\begin{aligned} & \left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \quad \mathbf{B}_\perp \right) \exp(\hat{\mathbf{C}}) \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} \\ &= \left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \quad \mathbf{B}_\perp \right) \mathbf{H} \exp(\hat{\Sigma}) \mathbf{H}^{-1} \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} \\ &= \left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \tilde{\mathbf{V}} \quad \mathbf{B}_\perp \tilde{\mathbf{U}} \right) \exp(\hat{\Sigma}) \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} \tilde{\mathbf{V}}^\top \\ &= \left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \tilde{\mathbf{V}} \quad \mathbf{B}_\perp \tilde{\mathbf{U}} \right) \begin{pmatrix} \cos(\tilde{\Sigma}_{[r]}) \\ \sin(\tilde{\Sigma}) \end{pmatrix} \tilde{\mathbf{V}}^\top. \end{aligned}$$

However, by definition of the Grassmann manifold as a quotient, for any $\mathbf{B}_1 \in \text{St}_r(\mathbb{R}^d)$ and $\mathbf{M} \in \text{GL}_r(\mathbb{R})$, $[\mathbf{B}_1 \mathbf{M}] = [\mathbf{B}_1]$. Therefore, we have

$$\begin{aligned} \left[\left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \quad \mathbf{B}_\perp \right) \exp(\hat{\mathbf{C}}) \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{s \times r} \end{pmatrix} \right] &= \left[\left(\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \tilde{\mathbf{V}} \quad \mathbf{B}_\perp \tilde{\mathbf{U}} \right) \begin{pmatrix} \cos(\tilde{\Sigma}_{[r]}) \\ \sin(\tilde{\Sigma}) \end{pmatrix} \right] \\ &= \left[\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \tilde{\mathbf{V}} \cos(\tilde{\Sigma}_{[r]}) + \mathbf{B}_\perp \tilde{\mathbf{U}} \mathbf{I}_{[s,r]} \sin(\tilde{\Sigma}_{[r]}) \right]. \end{aligned}$$

2.7. Complementary proofs for Section 2.3.1

Finally, the last term is equal to $\text{Exp}_{[\mathbf{B}]}(\mathbf{D})$ given in (2.62) since $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \{\mathbf{B}_\perp \tilde{\mathbf{U}}(\mathbf{I}_r)_{[s,r]}\} \tilde{\boldsymbol{\Sigma}}_{[r]} \tilde{\mathbf{V}}^\top$ is actually a compact singular value decomposition of $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}$. □

Computing the Cut Locus Consider the spectral norm of $\mathbf{M} \in \mathbb{R}^{d \times r}$, defined as

$$\|\mathbf{M}\| = \sup \{ \|\mathbf{M}v\|_2 : v \in \mathbb{R}^r, \|v\|_2 = 1 \} , \quad (2.83)$$

where $\|\cdot\|_2 : (v_i)_{i \in \{1, \dots, r\}} \in \mathbb{R}^r \mapsto \sqrt{\sum_{i=1}^r v_i^2}$. It is well known [Mey00, Equation (5.2.7)] that $\|\mathbf{M}\|$ is equal to the largest singular value of \mathbf{M} . For any $j, k \in \{1, \dots, d\}$, consider $\mathbf{e}_{j,k}$ the $d \times d$ matrix all of whose elements are zero, except the one at row j and column k , which is equal to 1.

Proposition 2.26. *Let $[\mathbf{B}]$ in $\text{Gr}_r(\mathbb{R}^d)$ and consider*

$$\mathcal{I}_{[\mathbf{B}]} = \left\{ \mathbf{D} \in \text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d) : \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\| < \pi/2 \right\} ,$$

where $\|\cdot\|$ is the spectral norm defined above.

(a) *The set $\mathcal{I}_{[\mathbf{B}]}$ is open in the tangent space $\text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ and the restriction of the Riemannian exponential map,*

$$\left(\text{Exp}_{[\mathbf{B}]} \right)_{|\mathcal{I}_{[\mathbf{B}]}} : \mathcal{I}_{[\mathbf{B}]} \rightarrow \text{Gr}_r(\mathbb{R}^d) ,$$

is a smooth diffeomorphism onto its image, which is exactly $\mathcal{U}_{\mathbf{B}}$.

(b) *For any $\mathbf{D} \in \text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, such that $\|\mathbf{D}\|_{[\mathbf{B}]} = 1$, the cutting time $c(\mathbf{D})$, defined in (2.159) satisfies*

$$c(\mathbf{D}) \leq \pi / \left(2 \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\| \right) .$$

(c) *The tangent injectivity domain $\text{TD}([\mathbf{B}])$, defined in Section 2.11.6, is equal to $\mathcal{I}_{[\mathbf{B}]}$. Therefore the tangent cut locus of $[\mathbf{B}]$ is the set:*

$$\left\{ \mathbf{D} \in \text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d) : \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\| = \pi/2 \right\} .$$

Proof. (a) Let us start by proving that $\mathcal{I}_{[\mathbf{B}]}$ is an open subset of $\text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$. Since the mapping

$$f : \begin{cases} \text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d) & \longrightarrow & \mathbb{R}_+ \\ \mathbf{D} & \longmapsto & \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\| \end{cases}$$

is a norm, it is continuous. Thus $\mathcal{I}_{[\mathbf{B}]} = f^{-1}([0, \pi/2))$ is open in $\text{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ as the inverse image of the open set $[0, \pi/2)$ (open in \mathbb{R}_+).

Now, take $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ and consider the singular value decomposition: $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, where $\mathbf{U} \in \mathbb{R}^{d \times r}$ satisfying $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$, $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$ is nonnegative and diagonal and $\mathbf{V} \in \text{O}_r(\mathbb{R})$ is a $r \times r$ -orthogonal matrix. We have, using that the spectral norm is equal to the largest singular value, $\mathbf{D} \in \mathcal{I}_{[\mathbf{B}]}$ if, and only if $\|\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top\| < \pi/2$, *i.e.*, $\|\mathbf{\Sigma}\|_\infty < \pi/2$, where for any matrix $\mathbf{A} \in \mathbb{R}^{r \times r}$, $\|\mathbf{A}\|_\infty = \max_{i,j \in \{1, \dots, r\}} |\mathbf{A}_{ij}|$ is the maximum norm. To prove that $\text{Exp}_{[\mathbf{B}]}$ is invertible on $\mathcal{I}_{[\mathbf{B}]}$, recall the expression of the exponential map given in (2.62) and take $\mathbf{D} \in \mathcal{I}_{[\mathbf{B}]}$, with the decomposition given above, $\mathbf{D} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top(\mathbf{B}^\top \mathbf{B})^{1/2}$. Using that $\cos(\mathbf{\Sigma})$ is invertible,

$$\begin{aligned} \text{Exp}_{[\mathbf{B}]}(\mathbf{D}) &= \left[\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \cos(\mathbf{\Sigma}) + \mathbf{U} \sin(\mathbf{\Sigma}) \right] \\ &= \left[\mathbf{B} + \mathbf{U} \tan(\mathbf{\Sigma}) \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2} \right] \\ &= \text{Ret}_{[\mathbf{B}]} \left(\mathbf{U} \tan(\mathbf{\Sigma}) \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2} \right). \end{aligned}$$

Let us prove that $\text{Ret}_{[\mathbf{B}]}$ and $\mathbf{D} \mapsto \mathbf{U} \tan(\mathbf{\Sigma}) \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2}$ are invertible. Similarly to the computation of the inverse of the parameterization φ_α for $\alpha \in \mathcal{A}$ in the proof of Proposition 2.22, define

$$\Psi_{\mathbf{B}} : \begin{cases} \mathcal{U}_{\mathbf{B}} & \longrightarrow \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d) \\ [\mathbf{Y}] & \longmapsto \sigma_{\mathbf{B}}([\mathbf{Y}]) - \mathbf{B} \end{cases},$$

where $\sigma_{\mathbf{B}}([\mathbf{Y}])$ is given by (2.69). It is easy to verify that $\Psi_{\mathbf{B}}$ is the inverse map of $\text{Ret}_{[\mathbf{B}]}$ on its image $\mathcal{U}_{\mathbf{B}}$. For the second map, consider $\mathbf{N} = \mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}$. Using Lemma 2.35 for \mathbf{N} and the mappings \tan and \arctan , we have that $\mathbf{U} \tan(\mathbf{\Sigma}) \mathbf{V}^\top$ and $\mathbf{U} \arctan(\mathbf{\Sigma}) \mathbf{V}^\top$ only depend on \mathbf{D} and \mathbf{B} , thus

$$\begin{cases} \mathcal{I}_{[\mathbf{B}]} & \longrightarrow \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d) \\ \mathbf{D} & \longmapsto \mathbf{U} \tan(\mathbf{\Sigma}) \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2} \end{cases}$$

is well-defined, and so is its inverse,

$$\begin{cases} \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d) & \longrightarrow \mathcal{I}_{[\mathbf{B}]} \\ \mathbf{D} & \longmapsto \mathbf{U} \arctan(\mathbf{\Sigma}) \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{1/2} \end{cases}.$$

We have proven that $\text{Exp}_{[\mathbf{B}]}$ is invertible on $\mathcal{I}_{[\mathbf{B}]}$ by composition. In particular, the image of $\text{Exp}_{[\mathbf{B}]}$ on $\mathcal{I}_{[\mathbf{B}]}$ is $\mathcal{U}_{\mathbf{B}} = \{[\mathbf{Y}] \in \text{Gr}_r(\mathbb{R}^d) : \det(\mathbf{B}^\top \mathbf{Y}) \neq 0\}$. Moreover, the Riemannian exponential is smooth [Lee19, Proposition 5.19-(c)]; therefore, to prove that it is a diffeomorphism, it is sufficient to prove that its derivative on $\mathcal{I}_{[\mathbf{B}]}$ is invertible.

We compute first the differential of the Riemannian exponential at a reference point of the Grassmann manifold, the subspace generated by the first r vectors of the canonical basis of \mathbb{R}^d . Then, using that the Grassmann manifold is a symmetrical manifold with isometry group $\text{O}_d(\mathbb{R})$, we can conclude about the differential of the Riemannian exponential at any point.

Denote by \mathbf{B}_0 the $d \times r$ -matrix whose column vectors are the first r vectors of the canonical basis of \mathbb{R}^d . We compute the exponential at $[\mathbf{B}_0]$ with tangent vector $\mathbf{D} \in$

2.7. Complementary proofs for Section 2.3.1

$T_{[\mathbf{B}_0]} \text{Gr}_r(\mathbb{R}^d)$ using the formula given in Lemma 2.25-(b), precisely (2.76). Combining this formula with the decomposition $\mathbf{D} = (\mathbf{B}_0)_\perp \mathbf{C}$ where $(\mathbf{B}_0)_\perp$ is the $d \times (d-r)$ -matrix composed by the last $d-r$ vectors of the canonical basis of \mathbb{R}^d and \mathbf{C} is a real $(d-r) \times r$ -matrix, meaning

$$\mathbf{D} = \begin{pmatrix} \mathbf{0}_{r \times r} \\ \mathbf{C} \end{pmatrix},$$

equation (2.76) simplifies here since $(\mathbf{B}_0 \quad (\mathbf{B}_0)_\perp) = \mathbf{I}_d$ and $\mathbf{B}_0^\top \mathbf{B}_0 = \mathbf{I}_r$,

$$\begin{aligned} \text{Exp}_{[\mathbf{B}_0]}(\mathbf{D}) &= \left[(\mathbf{B}_0 \quad (\mathbf{B}_0)_\perp) \exp \begin{pmatrix} \mathbf{0}_{r \times r} & -\mathbf{C}^\top \\ \mathbf{C} & \mathbf{0}_{(d-r) \times (d-r)} \end{pmatrix} \mathbf{B}_0 \right] \\ &= \left[\exp \begin{pmatrix} \mathbf{0}_{r \times r} & -\mathbf{C}^\top \\ \mathbf{C} & \mathbf{0}_{(d-r) \times (d-r)} \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ \mathbf{0}_{(d-r) \times r} \end{pmatrix} \right], \end{aligned} \quad (2.84)$$

where \exp is the matrix exponential. From now on, denote by $\widehat{\mathbf{C}}$ the $d \times d$ skew-symmetric matrix in (2.84):

$$\widehat{\mathbf{C}} = \begin{pmatrix} \mathbf{0}_{r \times r} & -\mathbf{C}^\top \\ \mathbf{C} & \mathbf{0}_{(d-r) \times (d-r)} \end{pmatrix}.$$

Note that $\widehat{\mathbf{C}}^\top = -\widehat{\mathbf{C}}$, therefore $\exp(\widehat{\mathbf{C}})$ is an orthogonal matrix. This allows us to rewrite (2.84) using (2.72),

$$\text{Exp}_{[\mathbf{B}_0]}(\mathbf{D}) = \exp(\widehat{\mathbf{C}}) \cdot [\mathbf{B}_0].$$

To differentiate without confusion of the indexes, denote by $E(\mathbf{D}) = \text{Exp}_{[\mathbf{B}_0]}(\mathbf{D})$. We are looking to compute its differential, $(DE)_{\mathbf{D}}$. For any $\xi = (\mathbf{B}_0)_\perp \mathbf{K} \in T_{[\mathbf{B}_0]} \text{Gr}_r(\mathbb{R}^d)$, denote again by $\widehat{\xi}$ the augmented skew-symmetric matrix

$$\widehat{\xi} = \begin{pmatrix} \mathbf{0}_{r \times r} & -\mathbf{K}^\top \\ \mathbf{K} & \mathbf{0}_{(d-r) \times (d-r)} \end{pmatrix}.$$

By definition of the differential,

$$(DE)_{\mathbf{D}}(\xi) = \left. \frac{d}{dt} \left(\exp(\widehat{\mathbf{C}} + t\widehat{\xi}) \cdot [\mathbf{B}_0] \right) \right|_{t=0}. \quad (2.85)$$

Consider $\mathbf{O} : [0, 1] \rightarrow \mathbb{R}^{d \times d}$ defined by $\mathbf{O}(t) = \exp(-\widehat{\mathbf{C}}) \exp(\widehat{\mathbf{C}} + t\widehat{\xi})$ for any $t \in [0, 1]$. Since $\widehat{\mathbf{C}}$ and $\widehat{\xi}$ are skew-symmetric matrices, so is $\widehat{\mathbf{C}} + t\widehat{\xi}$ for any $t \in [0, 1]$. Therefore $\mathbf{O}(t) \in O_d(\mathbb{R})$. Moreover since \exp is smooth, \mathbf{O} is a smooth curve on $O_d(\mathbb{R})$, with initial value $\mathbf{O}(0) = \exp(\mathbf{0}_{d \times d}) = \mathbf{I}_d$. Thus, we can rewrite (2.85) using (2.73),

$$\begin{aligned} (DE)_{\mathbf{D}}(\xi) &= \exp(\widehat{\mathbf{C}}) \cdot \left. \frac{d}{dt} \left(\exp(-\widehat{\mathbf{C}}) \exp(\widehat{\mathbf{C}} + t\widehat{\xi}) \cdot [\mathbf{B}_0] \right) \right|_{t=0} \\ &= \exp(\widehat{\mathbf{C}}) \cdot \left. \frac{d}{dt} (\mathbf{O}(t) \cdot \pi(\mathbf{B}_0)) \right|_{t=0}. \end{aligned}$$

By (2.72) and the chain rule, we have

$$\begin{aligned} \left. \frac{d}{dt} (\mathbf{O}(t) \cdot \pi(\mathbf{B}_0)) \right|_{t=0} &= \left. \frac{d}{dt} (\pi(\mathbf{O}(t)\mathbf{B}_0)) \right|_{t=0} \\ &= (D\pi)_{\mathbf{O}(0)\mathbf{B}_0} \left(\dot{\mathbf{O}}(0)\mathbf{B}_0 \right) = (D\pi)_{\mathbf{B}_0} \left(\dot{\mathbf{O}}(0)\mathbf{B}_0 \right). \end{aligned}$$

This computation gives,

$$(DE)_{\mathbf{D}}(\xi) = \exp(\widehat{\mathbf{C}}) \cdot (D\pi)_{\mathbf{B}_0} \left(\dot{\mathbf{O}}(0)\mathbf{B}_0 \right),$$

where

$$\dot{\mathbf{O}}(0) = \left. \frac{d}{dt} \left(\exp(-\widehat{\mathbf{C}}) \exp(\widehat{\mathbf{C}} + t\hat{\xi}) \right) \right|_{t=0}.$$

However, using [HH03, Theorem 5.4], we have

$$\dot{\mathbf{O}}(0) = \sum_{\ell=1}^{+\infty} (-1)^{\ell+1} \frac{\text{ad}_{\widehat{\mathbf{C}}}^{\ell-1}(\hat{\xi})}{\ell!}, \quad (2.86)$$

where $\text{ad}_{\widehat{\mathbf{C}}}$ is the linear map on $\mathbb{R}^{d \times d}$ defined by $\text{ad}_{\widehat{\mathbf{C}}}(\mathbf{M}) = \widehat{\mathbf{C}}\mathbf{M} - \mathbf{M}\widehat{\mathbf{C}}$, for any $\mathbf{M} \in \mathbb{R}^{d \times d}$. Using the block decomposition of the matrices $\widehat{\mathbf{C}}$ and $\hat{\xi}$ for the product $\text{ad}_{\widehat{\mathbf{C}}}(\hat{\xi})$, combined with a straightforward induction on $\ell \in \mathbb{N}^*$, proves that for any $\ell \in \mathbb{N}^*$, $\xi \in \mathbb{R}^{(d-r) \times r}$ and $\mathbf{C} \in \mathbb{R}^{(d-r) \times r}$, there exist real matrices $\mathbf{F}_\ell, \mathbf{G}_\ell, \mathbf{K}_\ell, \mathbf{L}_\ell$ such that

$$\text{ad}_{\widehat{\mathbf{C}}}^{2\ell}(\hat{\xi}) = \begin{pmatrix} \mathbf{0}_{r \times r} & \mathbf{K}_\ell \\ \mathbf{L}_\ell & \mathbf{0}_{(d-r) \times (d-r)} \end{pmatrix}, \quad \text{ad}_{\widehat{\mathbf{C}}}^{2\ell-1}(\hat{\xi}) = \begin{pmatrix} \mathbf{F}_\ell & \mathbf{0}_{r \times (d-r)} \\ \mathbf{0}_{(d-r) \times r} & \mathbf{G}_\ell \end{pmatrix}. \quad (2.87)$$

Moreover $\text{ad}_{\widehat{\mathbf{C}}}^0(\hat{\xi}) = \hat{\xi}$, meaning (2.87) holds with $\mathbf{L}_0 = \mathbf{K}$ and $\mathbf{K}_0 = -\mathbf{K}^\top$. Separating the odd and even terms in the sum in (2.86) gives

$$\begin{aligned} \dot{\mathbf{O}}(0)\mathbf{B}_0 &= - \sum_{\ell=1}^{+\infty} \frac{1}{(2\ell)!} \text{ad}_{\widehat{\mathbf{C}}}^{2\ell-1}(\hat{\xi}) \mathbf{B}_0 + \sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \text{ad}_{\widehat{\mathbf{C}}}^{2\ell}(\hat{\xi}) \mathbf{B}_0 \\ &= \xi + \sum_{\ell=1}^{+\infty} \frac{1}{(2\ell)!} \begin{pmatrix} -\mathbf{F}_\ell \\ \mathbf{L}_\ell \end{pmatrix}. \end{aligned}$$

Using Proposition 2.23 and (2.70),

$$\text{Ker}((D\pi)_{\mathbf{B}_0}) = \mathbf{V}_{\mathbf{B}_0} = \left\{ \begin{pmatrix} \mathbf{C} \\ \mathbf{0}_{(d-r) \times r} \end{pmatrix} : \mathbf{C} \in \mathbb{R}^{r \times r} \right\},$$

meaning that for any matrices \mathbf{F}, \mathbf{L} of sizes $r \times r$ and $(d-r) \times r$ respectively,

$$(D\pi)_{\mathbf{B}_0} \begin{pmatrix} \mathbf{F} \\ \mathbf{L} \end{pmatrix} = (D\pi)_{\mathbf{B}_0} \begin{pmatrix} \mathbf{0}_{r \times r} \\ \mathbf{L} \end{pmatrix}.$$

Therefore,

$$\begin{aligned}
 (\mathrm{D}\pi)_{\mathbf{B}_0}(\dot{\mathbf{O}}(0)\mathbf{B}_0) &= (\mathrm{D}\pi)_{\mathbf{B}_0} \left(\begin{array}{c} \mathbf{0}_{r \times r} \\ \sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \mathbf{L}_\ell \end{array} \right) \\
 &= (\mathrm{D}\pi)_{\mathbf{B}_0} \left(\sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \mathrm{ad}_{\widehat{\mathbf{C}}}^{2\ell}(\widehat{\xi}) \mathbf{B}_0 \right).
 \end{aligned} \tag{2.88}$$

Recalling the expression of the action of the orthogonal group on the tangent space, this summarizes as

$$(\mathrm{DE})_{\mathbf{D}}(\xi) = \exp(\widehat{\mathbf{C}}) \cdot (\mathrm{D}\pi)_{\mathbf{B}_0} \left(\sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \mathrm{ad}_{\widehat{\mathbf{C}}}^{2\ell}(\widehat{\xi}) \mathbf{B}_0 \right).$$

Consider $\mathrm{sn} : \mathbb{R}^{(d-r) \times r} \rightarrow \mathcal{L}(\mathbb{R}^{d \times d})$ the linear operator defined by

$$\mathrm{sn}(\mathbf{M}) = \sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \mathrm{ad}_{\widehat{\mathbf{M}}}^{2\ell}, \quad \text{where } \widehat{\mathbf{M}} = \begin{pmatrix} \mathbf{0}_{r \times r} & -\mathbf{M}^\top \\ \mathbf{M} & \mathbf{0}_{(d-r) \times (d-r)} \end{pmatrix}.$$

To prove that $(\mathrm{DE})_{\mathbf{D}}$ is invertible as long as $\mathbf{D} \in \mathcal{I}_{[\mathbf{B}_0]}$, we already have that the action of $\exp(\widehat{\mathbf{C}})$ on $\mathrm{T}_{[\mathbf{B}_0]}\mathrm{Gr}_r(\mathbb{R}^d)$ is invertible since $\exp(-\widehat{\mathbf{C}})$ is its inverse. Moreover, the consequences of Proposition 2.23 state that $(\mathrm{D}\pi)_{\mathbf{B}_0}$ is an isomorphism between $\mathbf{H}_{\mathbf{B}_0}$, given by (2.71), and $\mathrm{T}_{[\mathbf{B}_0]}\mathrm{Gr}_r(\mathbb{R}^d)$. Using the definition (2.71), we have that

$$\mathbf{H}_{\mathbf{B}_0} = \left\{ \begin{pmatrix} \mathbf{0}_{r \times r} \\ \mathbf{N} \end{pmatrix} : \mathbf{N} \in \mathbb{R}^{(d-r) \times r} \right\},$$

and, using (2.88), one can see that

$$\sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \mathrm{ad}_{\widehat{\mathbf{C}}}^{2\ell}(\widehat{\xi}) \mathbf{B}_0 \in \mathbf{H}_{\mathbf{B}_0}.$$

Therefore, to prove that $(\mathrm{DE})_{\mathbf{D}}$ is invertible, it is sufficient to prove that the linear operator $\mathrm{sn}(\mathbf{C}) : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d \times d}$ is an isomorphism. We do this in two steps. First, consider a singular value decomposition of \mathbf{C} , $\mathbf{C} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$, where $\mathbf{U} \in \mathrm{O}_{d-r}(\mathbb{R})$, $\boldsymbol{\Sigma} \in \mathbb{R}^{(d-r) \times r}$ is a diagonal matrix with nonnegative coefficients and $\mathbf{V} \in \mathrm{O}_r(\mathbb{R})$ and define $\mathbf{H} \in \mathbb{R}^{d \times d}$ the following matrix,

$$\mathbf{H} = \begin{pmatrix} \mathbf{V} & \mathbf{0}_{r \times (d-r)} \\ \mathbf{0}_{(d-r) \times r} & \mathbf{U} \end{pmatrix}.$$

Using that $\mathbf{C} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ and computations block-by-block, we have that $\mathbf{H}\widehat{\boldsymbol{\Sigma}}\mathbf{H}^\top = \widehat{\mathbf{C}}\mathbf{H}\widehat{\boldsymbol{\Sigma}}\mathbf{H}^{-1} = \widehat{\mathbf{C}}$, where

$$\widehat{\boldsymbol{\Sigma}} = \begin{pmatrix} \mathbf{0}_{r \times r} & -\boldsymbol{\Sigma}^\top \\ \boldsymbol{\Sigma} & \mathbf{0}_{(d-r) \times (d-r)} \end{pmatrix}.$$

Next, for any $\mathbf{N} \in \text{GL}_d(\mathbb{R})$, define the conjugation operator $\rho_{\mathbf{N}} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d \times d}$ by $\rho_{\mathbf{N}}(\mathbf{M}) = \mathbf{N}\mathbf{M}\mathbf{N}^{-1}$ for any $\mathbf{M} \in \mathbb{R}^{d \times d}$. The operator $\text{ad}_{\widehat{\mathbf{C}}}$ can be written as the composition $\rho_{\mathbf{H}} \circ \text{ad}_{\widehat{\Sigma}} \circ \rho_{\mathbf{H}}^{-1}$, using the singular value decomposition, indeed for any $\mathbf{M} \in \mathbb{R}^{d \times d}$,

$$\begin{aligned} \rho_{\mathbf{H}} \circ \text{ad}_{\widehat{\Sigma}} \circ \rho_{\mathbf{H}}^{-1}(\mathbf{M}) &= \rho_{\mathbf{H}} \left(\widehat{\Sigma} \mathbf{H}^{\top} \mathbf{M} \mathbf{H} - \mathbf{H}^{\top} \mathbf{M} \mathbf{H} \widehat{\Sigma} \right) = \mathbf{H} \widehat{\Sigma} \mathbf{H}^{\top} \mathbf{M} - \mathbf{M} \mathbf{H} \widehat{\Sigma} \mathbf{H}^{\top} \\ &= \widehat{\mathbf{C}} \mathbf{M} - \mathbf{M} \widehat{\mathbf{C}} = \text{ad}_{\widehat{\mathbf{C}}}(\mathbf{M}) . \end{aligned}$$

The consequence of this identity is that we can also factorize $\text{sn}(\mathbf{C})$ as follows,

$$\begin{aligned} \text{sn}(\mathbf{C}) &= \sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \left(\rho_{\mathbf{H}} \circ \text{ad}_{\widehat{\Sigma}} \circ \rho_{\mathbf{H}}^{-1} \right)^{2\ell} \\ &= \sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \rho_{\mathbf{H}} \circ \text{ad}_{\widehat{\Sigma}}^{2\ell} \circ \rho_{\mathbf{H}}^{-1} = \rho_{\mathbf{H}} \circ \text{sn}(\widehat{\Sigma}) \circ \rho_{\mathbf{H}}^{-1} . \end{aligned}$$

This concludes the first step of proving that $(DE)_{\mathbf{D}}$ is an isomorphism. Since $\rho_{\mathbf{H}}$ is clearly an isomorphism, we only need to prove now that $\text{sn}(\widehat{\Sigma})$ is an isomorphism.

The second step is proving this by finding a basis of eigenvectors for the operator $\text{ad}_{\widehat{\Sigma}}^2$ and for $\text{sn}(\widehat{\Sigma})$. Define $s = d - r$ and consider the following matrices for any $i, j \in \{1, \dots, r \wedge s\}$:

$$\begin{aligned} \mathbf{a}_{r,i,j} &= \mathbf{e}_{i,j} + \mathbf{e}_{r+i,r+j} , & \mathbf{b}_{r,i,j} &= \mathbf{e}_{r+i,j} - \mathbf{e}_{i,r+j} , \\ \mathbf{c}_{r,i,j} &= \mathbf{e}_{i,j} - \mathbf{e}_{r+i,r+j} , & \mathbf{d}_{r,i,j} &= \mathbf{e}_{r+i,j} + \mathbf{e}_{i,r+j} , \end{aligned} \quad (2.89)$$

where, for any $k, \ell \in \{1, \dots, d\}$, $\mathbf{e}_{k,\ell}$ is the $d \times d$ matrix whose coefficients are zero except for the coefficient (k, ℓ) which is equal to one.

Denote the diagonal elements of $\widehat{\Sigma}$ by $(\sigma_i)_{i \in \{1, \dots, r \wedge s\}}$, this allows us to write

$$\widehat{\Sigma} = \sum_{i=1}^{r \wedge s} \sigma_i \mathbf{b}_{r,i,i} . \quad (2.90)$$

Using (2.89)-(2.90) and that for any $i, j, k, \ell \in \{1, \dots, d\}$, $\mathbf{e}_{i,j} \mathbf{e}_{k,\ell} = \delta_{j,k} \mathbf{e}_{i,\ell}$, where $\delta_{j,k} = 1$ if $j = k$, and zero otherwise, we compute, for $i, j \in \{1, \dots, r \wedge s\}$,

$$\begin{aligned} \text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{i,j}) &= \sigma_i \mathbf{e}_{r+i,j} + \sigma_j \mathbf{e}_{i,r+j} , & \text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{r+i,r+j}) &= -\sigma_j \mathbf{e}_{r+i,j} - \sigma_i \mathbf{e}_{i,r+j} , \\ \text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{r+i,j}) &= -\sigma_i \mathbf{e}_{i,j} + \sigma_j \mathbf{e}_{r+i,r+j} , & \text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{i,r+j}) &= -\sigma_j \mathbf{e}_{i,j} + \sigma_i \mathbf{e}_{r+i,r+j} . \end{aligned}$$

Which, by adding or subtracting the equations together, by the definitions (2.89), gives

$$\begin{aligned} \text{ad}_{\widehat{\Sigma}}(\mathbf{a}_{r,i,j}) &= (\sigma_i - \sigma_j) \mathbf{b}_{r,i,j} , & \text{ad}_{\widehat{\Sigma}}(\mathbf{b}_{r,i,j}) &= (\sigma_j - \sigma_i) \mathbf{a}_{r,i,j} , \\ \text{ad}_{\widehat{\Sigma}}(\mathbf{c}_{r,i,j}) &= (\sigma_i + \sigma_j) \mathbf{d}_{r,i,j} , & \text{ad}_{\widehat{\Sigma}}(\mathbf{d}_{r,i,j}) &= -(\sigma_i + \sigma_j) \mathbf{c}_{r,i,j} . \end{aligned}$$

By applying twice these equations, we obtain that the matrices $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$ are eigenvectors for $\text{ad}_{\widehat{\Sigma}}^2$,

$$\begin{aligned} \text{ad}_{\widehat{\Sigma}}^2(\mathbf{a}_{r,i,j}) &= -(\sigma_i - \sigma_j)^2 \mathbf{a}_{r,i,j} , & \text{ad}_{\widehat{\Sigma}}^2(\mathbf{b}_{r,i,j}) &= -(\sigma_i - \sigma_j)^2 \mathbf{b}_{r,i,j} , \\ \text{ad}_{\widehat{\Sigma}}^2(\mathbf{c}_{r,i,j}) &= -(\sigma_i + \sigma_j)^2 \mathbf{c}_{r,i,j} , & \text{ad}_{\widehat{\Sigma}}^2(\mathbf{d}_{r,i,j}) &= -(\sigma_i + \sigma_j)^2 \mathbf{d}_{r,i,j} . \end{aligned}$$

2.7. Complementary proofs for Section 2.3.1

For any $i, j \in \{1, \dots, r \wedge s\}$ we have found four linearly independent eigenvectors, we can even see that any two different matrices defined in (2.89) are orthogonal for the natural euclidian scalar product on $\mathbb{R}^{d \times d}$, thus we have found in total $4(r \wedge s)^2$ linearly independant eigenvectors. If $r = s$, we have $r = s = d/2$, meaning we have found exactly d^2 eigenvectors, which concludes the diagonalization of $\text{ad}_{\hat{\Sigma}}^2$ in this case. To summarize, we look at the following table for the multiplicities of the eigenvalues of $\text{ad}_{\hat{\Sigma}}^2$.

Spectrum of $\text{ad}_{\hat{\Sigma}}^2$, case $r = s$			
Eigenvalue	Eigenvectors	Range for (i, j)	Multiplicity
$-(\sigma_i - \sigma_j)^2$	$\mathbf{a}_{r,i,j}, \mathbf{b}_{r,i,j}$	$i, j \in \{1, \dots, r\}$	2
$-(\sigma_i + \sigma_j)^2$	$\mathbf{c}_{r,i,j}, \mathbf{d}_{r,i,j}$	$i, j \in \{1, \dots, r\}$	2

(2.91)

We treat separately the cases $r < s$ and $r > s$ even though they are very similar. First, if $r < s$, then the matrices above only have nonnegative coefficients in position (i, j) for $i, j \in \{1, \dots, 2r\}$. So, using (2.89)-(2.90), we compute for $i \in \{1, \dots, s - r\}$ and $j \in \{1, \dots, r\}$,

$$\begin{aligned} \hat{\Sigma} \mathbf{e}_{2r+i,j} &= \mathbf{0}_{d \times d}, & \mathbf{e}_{2r+i,j} \hat{\Sigma} &= -\sigma_j \mathbf{e}_{2r+i,r+j}, \\ \hat{\Sigma} \mathbf{e}_{2r+i,r+j} &= \mathbf{0}_{d \times d}, & \mathbf{e}_{2r+i,r+j} \hat{\Sigma} &= \sigma_j \mathbf{e}_{2r+i,j}. \end{aligned}$$

Using these equations, we compute the image of these matrices by $\text{ad}_{\hat{\Sigma}}$ and $\text{ad}_{\hat{\Sigma}}^2$,

$$\begin{aligned} \text{ad}_{\hat{\Sigma}}(\mathbf{e}_{2r+i,j}) &= \sigma_j \mathbf{e}_{2r+i,r+j}, & \text{ad}_{\hat{\Sigma}}(\mathbf{e}_{2r+i,r+j}) &= -\sigma_j \mathbf{e}_{2r+i,j} \text{ and} \\ \text{ad}_{\hat{\Sigma}}^2(\mathbf{e}_{2r+i,j}) &= -\sigma_j^2 \mathbf{e}_{2r+i,j}, & \text{ad}_{\hat{\Sigma}}^2(\mathbf{e}_{2r+i,r+j}) &= -\sigma_j^2 \mathbf{e}_{2r+i,r+j}. \end{aligned}$$

This means that $\mathbf{e}_{2r+i,j}$ and $\mathbf{e}_{2r+i,r+j}$ are eigenvectors of $\text{ad}_{\hat{\Sigma}}^2$, with eigenvalue $-\sigma_j^2$. Similarly, for $i \in \{1, \dots, r\}$ and $j \in \{1, \dots, s - r\}$,

$$\begin{aligned} \hat{\Sigma} \mathbf{e}_{i,2r+j} &= \sigma_i \mathbf{e}_{r+i,2r+j}, & \mathbf{e}_{i,2r+j} \hat{\Sigma} &= \mathbf{0}_{d \times d}, \\ \hat{\Sigma} \mathbf{e}_{r+i,2r+j} &= -\sigma_i \mathbf{e}_{i,2r+j}, & \mathbf{e}_{r+i,2r+j} \hat{\Sigma} &= \mathbf{0}_{d \times d}. \end{aligned}$$

Using these equations, we compute the image of these matrices by $\text{ad}_{\hat{\Sigma}}$ and $\text{ad}_{\hat{\Sigma}}^2$,

$$\begin{aligned} \text{ad}_{\hat{\Sigma}}(\mathbf{e}_{i,2r+j}) &= \sigma_i \mathbf{e}_{r+i,2r+j}, & \text{ad}_{\hat{\Sigma}}(\mathbf{e}_{r+i,2r+j}) &= -\sigma_i \mathbf{e}_{i,2r+j} \text{ and} \\ \text{ad}_{\hat{\Sigma}}^2(\mathbf{e}_{i,2r+j}) &= -\sigma_i^2 \mathbf{e}_{i,2r+j}, & \text{ad}_{\hat{\Sigma}}^2(\mathbf{e}_{r+i,2r+j}) &= -\sigma_i^2 \mathbf{e}_{r+i,2r+j}. \end{aligned}$$

This means that $\mathbf{e}_{i,2r+j}$ and $\mathbf{e}_{r+i,2r+j}$ are eigenvectors of $\text{ad}_{\hat{\Sigma}}^2$, with eigenvalue $-\sigma_i^2$. Finally, for $i, j \in \{1, \dots, s - r\}$,

$$\hat{\Sigma} \mathbf{e}_{2r+i,2r+j} = \mathbf{e}_{2r+i,2r+j} \hat{\Sigma} = \mathbf{0}_{d \times d}.$$

This gives

$$\text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{2r+i,2r+j}) = \text{ad}_{\widehat{\Sigma}}^2(\mathbf{e}_{2r+i,2r+j}) = \mathbf{0}_{d \times d}.$$

When we count how many eigenvectors we have found with these computations, there are $r(s-r)$ matrices in $\{\mathbf{e}_{2r+i,j} : i \in \{1, \dots, s-r\}, j \in \{1, \dots, r\}\}$, in $\{\mathbf{e}_{2r+i,r+j} : i \in \{1, \dots, s-r\}, j \in \{1, \dots, r\}\}$, in $\{\mathbf{e}_{i,2r+j} : i \in \{1, \dots, r\}, j \in \{1, \dots, s-r\}\}$ and in $\{\mathbf{e}_{r+i,2r+j} : i \in \{1, \dots, r\}, j \in \{1, \dots, s-r\}\}$, totalling $4r(s-r)$ matrices. There are $(s-r)^2$ matrices in $\{\mathbf{e}_{2r+i,2r+j} : i \in \{1, \dots, s-r\}, j \in \{1, \dots, s-r\}\}$ and all these matrices form an orthogonal family, which gives that they are linearly independent.

To summarize, we look at the following table for the multiplicities of the eigenvalues of $\text{ad}_{\widehat{\Sigma}}^2$.

Spectrum of $\text{ad}_{\widehat{\Sigma}}^2$, case $r < s$			
Eigenvalue	Eigenvectors	Range for (i, j)	Multiplicity
$-(\sigma_i - \sigma_j)^2$	$\mathbf{a}_{r,i,j}, \mathbf{b}_{r,i,j}$	$i, j \in \{1, \dots, r\}$	2
$-(\sigma_i + \sigma_j)^2$	$\mathbf{c}_{r,i,j}, \mathbf{d}_{r,i,j}$	$i, j \in \{1, \dots, r\}$	2
$-\sigma_j^2$	$\mathbf{e}_{2r+i,j},$ $\mathbf{e}_{2r+i,r+j},$ $\mathbf{e}_{j,2r+i},$ $\mathbf{e}_{r+j,2r+i}$	$i \in \{1, \dots, s-r\},$ $j \in \{1, \dots, r\}$	$4(s-r)$
0	$\mathbf{e}_{2r+i,2r+j}$	$i, j \in \{1, \dots, s-r\}$	$(s-r)^2$

(2.92)

In total, we have found $4r^2 + 4r(s-r) + (s-r)^2 = (2r + (s-r))^2 = (s+r)^2 = d^2$ eigenvectors, which are all linearly independent since they are mutually orthogonal.

In the case $r > s$, the eigenvectors to compute are the following. For $i \in \{1, \dots, r-s\}, j \in \{1, \dots, s\}$, we have

$$\begin{aligned} \widehat{\Sigma} \mathbf{e}_{s+i,j} &= \mathbf{0}_{d \times d}, & \mathbf{e}_{s+i,j} \widehat{\Sigma} &= -\sigma_j \mathbf{e}_{s+i,r+j}, \\ \widehat{\Sigma} \mathbf{e}_{s+i,r+j} &= \mathbf{0}_{d \times d}, & \mathbf{e}_{s+i,r+j} \widehat{\Sigma} &= \sigma_j \mathbf{e}_{s+i,j}. \end{aligned}$$

Using these equations, we compute the image of these matrices by $\text{ad}_{\widehat{\Sigma}}$ and $\text{ad}_{\widehat{\Sigma}}^2$,

$$\begin{aligned} \text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{s+i,j}) &= \sigma_j \mathbf{e}_{s+i,r+j}, & \text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{s+i,r+j}) &= -\sigma_j \mathbf{e}_{s+i,j} \text{ and} \\ \text{ad}_{\widehat{\Sigma}}^2(\mathbf{e}_{s+i,j}) &= -\sigma_j^2 \mathbf{e}_{s+i,j}, & \text{ad}_{\widehat{\Sigma}}^2(\mathbf{e}_{s+i,r+j}) &= -\sigma_j^2 \mathbf{e}_{s+i,r+j}. \end{aligned}$$

This means that $\mathbf{e}_{s+i,j}$ and $\mathbf{e}_{s+i,r+j}$ are eigenvectors of $\text{ad}_{\widehat{\Sigma}}^2$, with eigenvalue $-\sigma_j^2$. Similarly, for $i \in \{1, \dots, s\}$ and $j \in \{1, \dots, r-s\}$,

$$\begin{aligned} \widehat{\Sigma} \mathbf{e}_{i,s+j} &= \sigma_i \mathbf{e}_{r+i,s+j}, & \mathbf{e}_{i,s+j} \widehat{\Sigma} &= \mathbf{0}_{d \times d}, \\ \widehat{\Sigma} \mathbf{e}_{r+i,s+j} &= -\sigma_i \mathbf{e}_{i,s+j}, & \mathbf{e}_{r+i,s+j} \widehat{\Sigma} &= \mathbf{0}_{d \times d}. \end{aligned}$$

Using these equations, we compute the image of these matrices by $\text{ad}_{\widehat{\Sigma}}$ and $\text{ad}_{\widehat{\Sigma}}^2$,

$$\begin{aligned} \text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{i,s+j}) &= \sigma_i \mathbf{e}_{r+i,s+j}, & \text{ad}_{\widehat{\Sigma}}(\mathbf{e}_{r+i,s+j}) &= -\sigma_i \mathbf{e}_{i,s+j} \text{ and} \\ \text{ad}_{\widehat{\Sigma}}^2(\mathbf{e}_{i,s+j}) &= -\sigma_i^2 \mathbf{e}_{i,s+j}, & \text{ad}_{\widehat{\Sigma}}^2(\mathbf{e}_{r+i,s+j}) &= -\sigma_i^2 \mathbf{e}_{r+i,s+j}. \end{aligned}$$

2.7. Complementary proofs for Section 2.3.1

This means that $\mathbf{e}_{i,s+j}$ and $\mathbf{e}_{r+i,s+j}$ are eigenvectors of $\text{ad}_{\widehat{\Sigma}}^2$, with eigenvalue $-\sigma_i^2$. Finally, for $i, j \in \{1, \dots, r-s\}$,

$$\widehat{\Sigma} \mathbf{e}_{s+i,s+j} = \mathbf{e}_{s+i,s+j} \widehat{\Sigma} = \mathbf{0}_{d \times d}.$$

This gives

$$\text{ad}_{\widehat{\Sigma}}^2(\mathbf{e}_{s+i,s+j}) = \text{ad}_{\widehat{\Sigma}}^2(\mathbf{e}_{s+i,s+j}) = \mathbf{0}_{d \times d}.$$

When we count how many eigenvectors we have found with these computations, there are $s(r-s)$ matrices in $\{\mathbf{e}_{s+i,j} : i \in \{1, \dots, r-s\}, j \in \{1, \dots, s\}\}$, in $\{\mathbf{e}_{s+i,r+j} : i \in \{1, \dots, r-s\}, j \in \{1, \dots, s\}\}$, in $\{\mathbf{e}_{i,s+j} : i \in \{1, \dots, s\}, j \in \{1, \dots, r-s\}\}$ and in $\{\mathbf{e}_{r+i,s+j} : i \in \{1, \dots, s\}, j \in \{1, \dots, r-s\}\}$, totalling $4s(r-s)$ matrices. There are $(r-s)^2$ matrices in $\{\mathbf{e}_{s+i,s+j} : i \in \{1, \dots, r-s\}, j \in \{1, \dots, r-s\}\}$ and all these matrices form an orthogonal family, which gives that they are linearly independant.

To summarize, we look at the following table for the multiplicities of the eigenvalues of $\text{ad}_{\widehat{\Sigma}}^2$.

Spectrum of $\text{ad}_{\widehat{\Sigma}}^2$, case $r > s$			
Eigenvalue	Eigenvectors	Range for (i, j)	Multiplicity
$-(\sigma_i - \sigma_j)^2$	$\mathbf{a}_{r,i,j}, \mathbf{b}_{r,i,j}$	$i, j \in \{1, \dots, s\}$	2
$-(\sigma_i + \sigma_j)^2$	$\mathbf{c}_{r,i,j}, \mathbf{d}_{r,i,j}$	$i, j \in \{1, \dots, s\}$	2
$-\sigma_j^2$	$\mathbf{e}_{s+i,j},$ $\mathbf{e}_{s+i,r+j},$ $\mathbf{e}_{j,s+i},$ $\mathbf{e}_{r+j,s+i}$	$i \in \{1, \dots, r-s\},$ $j \in \{1, \dots, s\}$	$4(r-s)$
0	$\mathbf{e}_{s+i,s+j}$	$i, j \in \{1, \dots, r-s\}$	$(r-s)^2$

(2.93)

In total, we have found $4s^2 + 4s(r-s) + (r-s)^2 = (2s + (r-s))^2 = (s+r)^2 = d^2$ eigenvectors, which are all linearly independant since they are mutually orthogonal.

Now, to compute the eigenvalues of $\text{sn}(\Sigma)$ from those of $\text{ad}_{\widehat{\Sigma}}^2$, for any eigenvalue $\lambda \in \mathbb{R}$ of $\text{ad}_{\widehat{\Sigma}}^2$ using tables (2.91)-(2.92)-(2.93), we have that λ is nonpositive, meaning we prefer to write $\sigma = -\lambda \in \mathbb{R}_+$. We denote by $\mathbf{v}_\sigma \in \mathbb{R}^{d \times d}$ an eigenvector of $\text{ad}_{\widehat{\Sigma}}^2$ of eigenvalue $-\sigma$. Then

$$\begin{aligned} \text{sn}(\Sigma)(\mathbf{v}_\sigma) &= \sum_{\ell=0}^{+\infty} \frac{1}{(2\ell+1)!} \left(\text{ad}_{\widehat{\Sigma}}^2\right)^\ell (\mathbf{v}_\sigma) = \sum_{\ell=0}^{+\infty} \frac{(-\sigma)^\ell}{(2\ell+1)!} \mathbf{v}_\sigma \\ &= \sum_{\ell=0}^{+\infty} \frac{(-1)^\ell (\sqrt{\sigma})^{2\ell}}{(2\ell+1)!} \mathbf{v}_\sigma = \frac{\sin(\sqrt{\sigma})}{\sqrt{\sigma}} \mathbf{v}_\sigma = \text{sinc}(\sqrt{\sigma}) \mathbf{v}_\sigma, \end{aligned} \quad (2.94)$$

where sinc is defined over \mathbb{R} by $\text{sinc}(x) = \sin(x)/x$ for $x \in \mathbb{R} \setminus \{0\}$ and $\text{sinc}(0) = 1$. By

using (2.94) on the results of (2.91)-(2.92)-(2.93), we obtain,

Spectrum of $\text{sn}(\Sigma)$		
Eigenvalue	Range for (i, j)	Multiplicity
$\text{sinc}(\sigma_i - \sigma_j)$	$i, j \in \{1, \dots, r \wedge s\}$	2
$\text{sinc}(\sigma_i + \sigma_j)$	$i, j \in \{1, \dots, r \wedge s\}$	2
$\text{sinc}(\sigma_j)$	$i \in \{1, \dots, d - 2(r \wedge s)\},$ $j \in \{1, \dots, r \wedge s\}$	$4[d - 2(r \wedge s)]$ (if $r \neq s$)
1	$i, j \in \{1, \dots, d - 2(r \wedge s)\}$	$[d - 2(r \wedge s)]^2$

We have removed the column for the eigenvectors, because it is the same as the one in previous tables (2.91)-(2.92)-(2.93). Now, to prove that $\text{sn}(\Sigma)$ is an isomorphism, we only need to check that all eigenvalues are nonzero. The condition $\mathbf{D} \in \mathcal{I}_{[\mathbf{B}_0]}$ gives that for any $i \in \{1, \dots, r \wedge s\}$, $\sigma_i \in [0, \pi/2)$, meaning that $\text{sinc}(\sigma_i) \in (2/\pi, 1]$, as well as that for any $i, j \in \{1, \dots, r \wedge s\}$, $\text{sinc}(\sigma_i - \sigma_j) \in (2/\pi, 1]$ and $\text{sinc}(\sigma_i + \sigma_j) \in (0, 1]$. This concludes that $\text{sn}(\Sigma)$ is an isomorphism. And by previously stated arguments, this concludes that $(DE)_{\mathbf{D}}$ is an isomorphism, therefore the derivative of $\text{Exp}_{[\mathbf{B}_0]}$ on $\mathcal{I}_{[\mathbf{B}_0]}$ is invertible.

Now, to prove that for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$, the derivative of $\text{Exp}_{[\mathbf{B}]}$ on $\mathcal{I}_{[\mathbf{B}]}$ is invertible, we use the following symmetry argument. Let $([\mathbf{B}], \mathbf{D}) \in \text{TGr}_r(\mathbb{R}^d)$, since the Grassmann manifold is symmetric, its isometry group $O_d(\mathbb{R})$ acts transitively on the manifold. Therefore, there exists $\mathbf{g}_0 \in O_d(\mathbb{R})$ such that $\mathbf{g}_0 \cdot [\mathbf{B}_0] = [\mathbf{B}]$. Now, differentiating (2.161) brings, for any $([\mathbf{B}_1], \mathbf{D}_1) \in \text{TGr}_r(\mathbb{R}^d)$, $\mathbf{g} \in O_d(\mathbb{R})$ and $\xi \in \text{T}_{[\mathbf{B}_1]}\text{Gr}_r(\mathbb{R}^d)$,

$$\mathbf{g} \cdot (\text{DExp}_{[\mathbf{B}_1]})_{\mathbf{D}_1}(\xi) = (\text{DExp}_{\mathbf{g} \cdot [\mathbf{B}_1]})_{\mathbf{g} \mathbf{D}_1}(\mathbf{g} \xi) . \quad (2.95)$$

Applying (2.95) to $[\mathbf{B}_0]$, \mathbf{g}_0 and $\mathbf{g}_0^{-1} \mathbf{D} \in \text{T}_{[\mathbf{B}_0]}\text{Gr}_r(\mathbb{R}^d)$

$$\mathbf{g}_0 \cdot (\text{DExp}_{[\mathbf{B}_0]})_{\mathbf{g}_0^{-1} \mathbf{D}}(\xi) = (\text{DExp}_{[\mathbf{B}]})_{\mathbf{D}}(\mathbf{g}_0 \xi) ,$$

which brings, for any $\xi' \in \text{T}_{[\mathbf{B}]}\text{Gr}_r(\mathbb{R}^d)$,

$$(\text{DExp}_{[\mathbf{B}]})_{\mathbf{D}}(\xi') = \mathbf{g}_0 \cdot (\text{DExp}_{[\mathbf{B}_0]})_{\mathbf{g}_0^{-1} \mathbf{D}}(\mathbf{g}_0^{-1} \xi') . \quad (2.96)$$

However, we have that $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}$ and $\mathbf{g}_0^{-1} \mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}$ have the same singular values, because if $\mathbf{U} \Sigma \mathbf{V}^\top$ is a singular value decomposition of $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}$, then $(\mathbf{g}_0^{-1} \mathbf{U}) \Sigma \mathbf{V}^\top$ is a singular value decomposition of $\mathbf{g}_0^{-1} \mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}$. Therefore, if $\mathbf{D} \in \mathcal{I}_{[\mathbf{B}]}$, then $\mathbf{g}_0^{-1} \mathbf{D}$ is in $\mathcal{I}_{[\mathbf{B}_0]}$.

By definition of a group action, the action of \mathbf{g}_0 on $\text{T}_{[\mathbf{B}_0]}\text{Gr}_r(\mathbb{R}^d)$ is an invertible mapping, moreover, we have just proven that $(\text{DExp}_{[\mathbf{B}_0]})_{\mathbf{g}_0^{-1} \mathbf{D}}$ is invertible. Using (2.96), we conclude that $(\text{DExp}_{[\mathbf{B}]})_{\mathbf{D}}$ is an isomorphism as a composition of invertible mappings. This concludes the proof that $\text{Exp}_{[\mathbf{B}]}$ is a diffeomorphism from $\mathcal{I}_{[\mathbf{B}]}$ to $\mathcal{U}_{\mathbf{B}}$.

(b) Let $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, such that $\|\mathbf{D}\|_{\mathbf{B}} = 1$. Let $t > \pi/(2\|\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}\|)$. Our goal is to prove that there exists $\mathbf{C} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ such that $\text{Exp}_{[\mathbf{B}]}(t\mathbf{D}) = \text{Exp}_{[\mathbf{B}]}(\mathbf{C})$ and $\|\mathbf{C}\|_{[\mathbf{B}]} < \|t\mathbf{D}\|_{[\mathbf{B}]} = t$. Once we have that, since for any $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, the length $\ell(\gamma_{\mathbf{D}})$ of the geodesic $\gamma_{\mathbf{D}} : s \in [0, 1] \mapsto \text{Exp}_{[\mathbf{B}]}(s\mathbf{D})$ is equal to $\|\mathbf{D}\|_{[\mathbf{B}]}$, see [Lee19, Corollary 5.6(b)], it comes that $\rho_{\Theta}([\mathbf{B}], \text{Exp}_{[\mathbf{B}]}(t\mathbf{D})) \leq \ell(\gamma_{\mathbf{C}}) = \|\mathbf{C}\|_{[\mathbf{B}]} < t$, and thus $\gamma_{\mathbf{D}}$ is not length-minimizing at t , meaning $t > c(\mathbf{D})$. Since we do this for any $t > \pi/(2\|\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}\|)$, we conclude that $c(\mathbf{D}) \leq \pi/(2\|\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}\|)$.

Now, take the singular value decomposition $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ with \mathbf{U} a $d \times r$ -matrix such that $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$, $\mathbf{\Sigma}$ a $r \times r$ -nonnegative diagonal matrix and \mathbf{V} a $r \times r$ orthogonal matrix. Denote by $\mathbf{\Sigma} = \text{diag}((\Sigma_i)_{i \in \{1, \dots, r\}})$, and, up to permuting the columns of \mathbf{U} , we can assume $(\Sigma_i)_{i \in \{1, \dots, r\}}$ nondecreasing. Since we have

$$\|t\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2}\| = t \|\mathbf{\Sigma}\|_{\infty} > \pi/2,$$

then $t\Sigma_1 > \pi/2$. Now, to obtain another preimage of $\text{Exp}_{[\mathbf{B}]}(t\mathbf{D})$, we want to exploit the periodicity of cos and sin, and therefore modify $\mathbf{\Sigma}$ into a matrix with smaller coefficients. Since $t\Sigma_1$ is the largest value of $t\Sigma$, and we have $\|t\mathbf{\Sigma}\|_{\infty} > \pi/2$, then $t\Sigma_1 > \pi/2$.

Consider the two following $r \times r$ -matrices, defined block-by-block

$$\mathbf{A} = \begin{pmatrix} \pi & \mathbf{0}_{1 \times (r-1)} \\ \mathbf{0}_{(r-1) \times 1} & \mathbf{0}_{(r-1) \times (r-1)} \end{pmatrix}, \quad \mathbf{S} = \begin{pmatrix} -1 & \mathbf{0}_{1 \times (r-1)} \\ \mathbf{0}_{(r-1) \times 1} & \mathbf{I}_{r-1} \end{pmatrix}.$$

Assume in the first place that $t\Sigma_1 > \pi$. Let $\tilde{\mathbf{\Sigma}} = t\mathbf{\Sigma} - \mathbf{A}$, which is still a nonnegative diagonal matrix, and consider $\mathbf{C} = \mathbf{U}\tilde{\mathbf{\Sigma}}\mathbf{V}^\top(\mathbf{B}^\top \mathbf{B})^{1/2}$. Then the SVD of $\mathbf{C}(\mathbf{B}^\top \mathbf{B})^{-1/2}$ is actually $\mathbf{U}\tilde{\mathbf{\Sigma}}\mathbf{V}^\top$ and, using that $\cos(x - \pi) = -\cos(x)$ and $\sin(x - \pi) = -\sin(x)$,

$$\begin{aligned} \text{Exp}_{[\mathbf{B}]}(\mathbf{C}) &= \left[\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \cos(t\mathbf{\Sigma} - \mathbf{A}) + \mathbf{U} \sin(t\mathbf{\Sigma} - \mathbf{A}) \right] \\ &= \left[\mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \cos(t\mathbf{\Sigma})\mathbf{S} + \mathbf{U} \sin(t\mathbf{\Sigma})\mathbf{S} \right] \\ &= \text{Exp}_{[\mathbf{B}]}(t\mathbf{D}), \end{aligned}$$

by using that \mathbf{S} is invertible. Using the SVD and that for any $m \times n$ matrix \mathbf{M} and any $n \times m$ matrix \mathbf{N} , $\text{Tr}(\mathbf{MN}) = \text{Tr}(\mathbf{NM})$, we can easily check that

$$\begin{aligned} \|\mathbf{D}\|_{\mathbf{B}}^2 &= \text{Tr} \left((\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{D}^\top \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right) \\ &= \text{Tr} \left(\mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^\top \right) \\ &= \text{Tr} \left(\mathbf{V}^\top \mathbf{V}\mathbf{\Sigma}^2 \right) \\ &= \text{Tr} \left(\mathbf{\Sigma}^2 \right) \\ &= \sum_{i=1}^r \Sigma_i^2. \end{aligned} \tag{2.97}$$

Therefore, for \mathbf{C} we obtain $\|\mathbf{C}\|_{[\mathbf{B}]}^2 = \|t\mathbf{D}\|_{[\mathbf{B}]}^2 - t^2\boldsymbol{\Sigma}_1^2 + (t\boldsymbol{\Sigma}_1 - \pi)^2 = \|t\mathbf{D}\|_{[\mathbf{B}]}^2 - 2t\pi\boldsymbol{\Sigma}_1 + \pi^2 < \|t\mathbf{D}\|_{[\mathbf{B}]}^2 - \pi^2 < \|t\mathbf{D}\|_{[\mathbf{B}]}^2$. This concludes this case of the proof.

Assume now $\pi/2 < t\boldsymbol{\Sigma}_1 \leq \pi$. Let $\tilde{\boldsymbol{\Sigma}} = t\mathbf{S}\boldsymbol{\Sigma} + \mathbf{A} = \text{diag}(\pi - t\boldsymbol{\Sigma}_1, t\boldsymbol{\Sigma}_2, \dots, t\boldsymbol{\Sigma}_r)$, which is still a nonnegative diagonal matrix, and consider $\tilde{\mathbf{V}} = \mathbf{V}\mathbf{S}$ as well as $\mathbf{C} = \mathbf{U}\tilde{\boldsymbol{\Sigma}}\tilde{\mathbf{V}}^\top (\mathbf{B}^\top \mathbf{B})^{1/2}$. Then an SVD of $\mathbf{C}(\mathbf{B}^\top \mathbf{B})^{-1/2}$ is actually $\mathbf{U}\tilde{\boldsymbol{\Sigma}}\tilde{\mathbf{V}}^\top$ since $\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} = \mathbf{S}\mathbf{V}^\top \mathbf{V}\mathbf{S} = \mathbf{S}\mathbf{S} = \mathbf{I}_r$. Now, using that $\cos(\pi - x) = -\cos(x)$ and $\sin(\pi - x) = \sin(x)$,

$$\begin{aligned} \text{Exp}_{[\mathbf{B}]}(\mathbf{C}) &= \left[\mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V}\mathbf{S} \cos(t\mathbf{S}\boldsymbol{\Sigma} + \mathbf{A}) + \mathbf{U} \sin(t\mathbf{S}\boldsymbol{\Sigma} + \mathbf{A}) \right] \\ &= \left[\mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V}\mathbf{S} \cos(t\boldsymbol{\Sigma}) \mathbf{S} + \mathbf{U} \sin(t\boldsymbol{\Sigma}) \right]. \end{aligned}$$

Using that \mathbf{S} and $\cos(t\boldsymbol{\Sigma})$ commute because they are diagonal, we have

$$\text{Exp}_{[\mathbf{B}]}(\mathbf{C}) = \left[\mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \cos(t\boldsymbol{\Sigma}) + \mathbf{U} \sin(t\boldsymbol{\Sigma}) \right] = \text{Exp}_{[\mathbf{B}]}(t\mathbf{D}).$$

Now, using the calculation done in (2.97), we have $\|\mathbf{C}\|_{[\mathbf{B}]}^2 = \|t\mathbf{D}\|_{[\mathbf{B}]}^2 - t^2\boldsymbol{\Sigma}_1^2 + (\pi - t\boldsymbol{\Sigma}_1)^2 = \|t\mathbf{D}\|_{[\mathbf{B}]}^2 + \pi^2 - 2t\pi\boldsymbol{\Sigma}_1$. However $t\boldsymbol{\Sigma}_1 > \pi/2$, thus $2t\pi\boldsymbol{\Sigma}_1 > \pi^2$. Therefore $\|\mathbf{C}\|_{[\mathbf{B}]}^2 < \|t\mathbf{D}\|_{[\mathbf{B}]}^2$. This concludes the proof under the assumption $t\boldsymbol{\Sigma}_1 > \pi/2$.

(c) Let us prove that $\mathcal{I}_{[\mathbf{B}]} = \text{TD}([\mathbf{B}])$, where the tangent injectivity domain $\text{TD}([\mathbf{B}]) = \{t\mathbf{C} : \mathbf{C} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d), \|\mathbf{C}\|_{[\mathbf{B}]} = 1, t < c(\mathbf{C})\}$, where the cut time $c(\cdot)$ is defined in (2.159). The inclusion $\text{TD}([\mathbf{B}]) \subset \mathcal{I}_{[\mathbf{B}]}$ is given by Proposition 2.26-(b), we prove the other one by contradiction.

Suppose there exists $\mathbf{W} \in \mathcal{I}_{[\mathbf{B}]} \setminus \text{TD}([\mathbf{B}])$. Setting $\mathbf{M} = \mathbf{W}/\|\mathbf{W}\|_{[\mathbf{B}]}$, since $\mathbf{W} \notin \text{TD}([\mathbf{B}])$, $\|\mathbf{W}\|_{[\mathbf{B}]} \geq c(\mathbf{M})$. However, $c(\mathbf{M})\mathbf{M}$ is not in $\text{TD}([\mathbf{B}])$ and, by definition of $\mathcal{I}_{[\mathbf{B}]}$, $\{t\mathbf{M} : t \leq \|\mathbf{W}\|_{[\mathbf{B}]}\} \subset \mathcal{I}_{[\mathbf{B}]}$ because $\mathbf{W} = \|\mathbf{W}\|_{[\mathbf{B}]} \mathbf{M} \in \mathcal{I}_{[\mathbf{B}]}$. Therefore $c(\mathbf{M})\mathbf{M} \in \mathcal{I}_{[\mathbf{B}]} \setminus \text{TD}([\mathbf{B}])$. We set $\mathbf{D} = c(\mathbf{M})\mathbf{M}$, which means that we have $c(\mathbf{M}) = \|\mathbf{D}\|_{[\mathbf{B}]}$ and by definition of the cut time given in (2.159), $\text{Exp}_{[\mathbf{B}]}(\mathbf{D}) \in \text{Cut}([\mathbf{B}])$.

Now, we can apply [Lee19, Proposition 10.32] to our point $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ and tangent vector $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$. Since $c(\mathbf{M}) = \|\mathbf{D}\|_{[\mathbf{B}]}$, we are in the case [Lee19, Proposition 10.32-(b)], meaning that the geodesic $\gamma : t \in [0, \|\mathbf{D}\|_{[\mathbf{B}]}] \mapsto \text{Exp}_{[\mathbf{B}]}(t\mathbf{D}) \in \text{Gr}_r(\mathbb{R}^d)$ is minimizing and that one or both of the following conditions are true:

(i) $\text{Exp}_{[\mathbf{B}]}(\mathbf{D})$ is conjugate to $[\mathbf{B}]$ along γ , *i.e.* \mathbf{D} is a critical point of $\text{Exp}_{[\mathbf{B}]}$, see [Lee19, Proposition 10.20].

(ii) There are at least two unit-speed minimizing geodesics from $[\mathbf{B}]$ to $\text{Exp}_{[\mathbf{B}]}(\mathbf{D})$.

To find the contradiction, we prove that both of these conditions are false. Firstly, (i) does not hold using Proposition 2.26-(a) and that \mathbf{D} belongs to $\mathcal{I}_{[\mathbf{B}]}$.

Now, since (i) does not hold, (ii) is true. Therefore, there exists $\mathbf{D}' \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ such that $\mathbf{D}' \neq \mathbf{D}$, that the geodesic $\gamma_{\mathbf{D}'} : t \in [0, 1] \mapsto \text{Exp}_{[\mathbf{B}]}(t\mathbf{D}')$ is minimizing on $[0, 1]$ and that $\text{Exp}_{[\mathbf{B}]}(\mathbf{D}) = \text{Exp}_{[\mathbf{B}]}(\mathbf{D}')$. Indeed, to claim $\mathbf{D}' \neq \mathbf{D}$, we use that having two different unit-speed geodesics means that $\mathbf{D}'/\|\mathbf{D}'\|_{[\mathbf{B}]} \neq \mathbf{D}/\|\mathbf{D}\|_{[\mathbf{B}]}$, however, since both have the

same image by $\text{Exp}_{[\mathbf{B}]}$ and both geodesics are minimizing, the distance $\rho_{\Theta}([\mathbf{B}], \text{Exp}_{[\mathbf{B}]}(\mathbf{D}))$ is equal to the length of both geodesics, namely $\|\mathbf{D}'\|_{[\mathbf{B}]}$ and $\|\mathbf{D}\|_{[\mathbf{B}]}$, which gives $\mathbf{D}' \neq \mathbf{D}$.

By injectivity of $\text{Exp}_{[\mathbf{B}]}$ on $\mathcal{I}_{[\mathbf{B}]}$ and that $\mathbf{D}' \neq \mathbf{D}$, we have $\mathbf{D}' \notin \mathcal{I}_{[\mathbf{B}]}$. Moreover, by definition of the cut time $c(\mathbf{D}'/\|\mathbf{D}'\|_{[\mathbf{B}]})$, and since $\gamma_{\mathbf{D}'}$ is minimizing on $[0, 1]$, we have that $\mathbf{D}' \in \text{TD}([\mathbf{B}]) \cup \text{TCut}([\mathbf{B}])$. Using Proposition 2.26-(b), we have $\text{TD}([\mathbf{B}]) \subset \mathcal{I}_{[\mathbf{B}]}$ which, combined with $\mathbf{D}' \notin \mathcal{I}_{[\mathbf{B}]}$, gives $\mathbf{D}' \in \text{TCut}([\mathbf{B}])$. We now look for a contradiction with respect to the injectivity of $\text{Exp}_{[\mathbf{B}]}$ on $\mathcal{I}_{[\mathbf{B}]}$.

Using Proposition 2.26-(a), we have that $\mathcal{I}_{[\mathbf{B}]}$ is open in $T_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$. Thus, there exists a radius $r_0 > 0$ such that the open ball $\mathbf{U} = \mathbf{B}(\mathbf{D}, r_0) \subset T_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ is contained inside $\mathcal{I}_{[\mathbf{B}]}$. Since $\mathbf{D} \neq \mathbf{D}'$, we have $\|\mathbf{D} - \mathbf{D}'\|_{[\mathbf{B}]} > 0$ and therefore, up to dividing r_0 by some positive number greater than one, we can assume that $\inf_{\mathbf{C} \in \mathbf{U}} \|\mathbf{C} - \mathbf{D}'\|_{[\mathbf{B}]} > 0$. Denoting the preimage of $\text{Exp}_{[\mathbf{B}]}(\mathbf{U})$ by the mapping $\text{Exp}_{[\mathbf{B}]}$, $\mathbf{O} = \text{Exp}_{[\mathbf{B}]}^{-1}\{\text{Exp}_{[\mathbf{B}]}(\mathbf{U})\}$, we have, since $\text{Exp}_{[\mathbf{B}]}$ is a diffeomorphism, that $\text{Exp}_{[\mathbf{B}]}(\mathbf{U})$ is open in $\text{Gr}_r(\mathbb{R}^d)$ and that \mathbf{O} is open in $T_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$. Moreover \mathbf{O} contains \mathbf{D} and \mathbf{D}' . Therefore, there exists $r_1 > 0$ such that the open ball $\mathbf{U}' = \mathbf{B}(\mathbf{D}', r_1) \subset T_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ satisfies $\mathbf{U}' \subset \mathbf{O}$. By the assumption made on the choice of r_0 , and up to division of r_1 by a large enough positive number, we can assume that $\mathbf{U} \cap \mathbf{U}' = \emptyset$.

We prove now that there exists $\mathbf{C}' \in \mathbf{U}' \cap \mathcal{I}_{[\mathbf{B}]}$. Since $\mathbf{D}' \in \text{TCut}([\mathbf{B}])$, then for any $t \in [0, 1)$, $t\mathbf{D}' \in \text{TD}([\mathbf{B}]) \subset \mathcal{I}_{[\mathbf{B}]}$. Therefore the sequence of tangent vectors $((1 - 2^{-n}/\|\mathbf{D}'\|_{[\mathbf{B}]})\mathbf{D}')_{n \in \mathbb{N}^*}$ has its values on $\mathcal{I}_{[\mathbf{B}]}$. Moreover, there exists $N \in \mathbb{N}^*$ such that for any $n \geq N$, $2^{-n} < r_1$, thus

$$\begin{aligned} \left\| \mathbf{D}' - (1 - 2^{-N}/\|\mathbf{D}'\|_{[\mathbf{B}]})\mathbf{D}' \right\|_{[\mathbf{B}]} &= \left| 1 - 1 + 2^{-N}/\|\mathbf{D}'\|_{[\mathbf{B}]} \right| \|\mathbf{D}'\|_{[\mathbf{B}]} \\ &= 2^{-N} \|\mathbf{D}'\|_{[\mathbf{B}]} / \|\mathbf{D}'\|_{[\mathbf{B}]} < r_1. \end{aligned}$$

This gives that $\mathbf{C}' = (1 - 2^{-N}/\|\mathbf{D}'\|_{[\mathbf{B}]})\mathbf{D}' \in \mathbf{U}'$. We have therefore found $\mathbf{C}' \in \mathbf{U}' \cap \mathcal{I}_{[\mathbf{B}]}$.

Finally, by definition of \mathbf{U}' , $\text{Exp}_{[\mathbf{B}]}(\mathbf{C}') \in \text{Exp}_{[\mathbf{B}]}(\mathbf{U})$, thus there exists $\mathbf{C} \in \mathbf{U}$ such that $\text{Exp}_{[\mathbf{B}]}(\mathbf{C}') = \text{Exp}_{[\mathbf{B}]}(\mathbf{C})$. Since $\mathbf{U} \cap \mathbf{U}' = \emptyset$, we have $\mathbf{C}' \neq \mathbf{C}$ and, since $\mathbf{U} \subset \mathcal{I}_{[\mathbf{B}]}$, we have $\mathbf{C} \in \mathcal{I}_{[\mathbf{B}]}$. This is contradictory with Proposition 2.26-(a), which states the injectivity of $\text{Exp}_{[\mathbf{B}]}$ on $\mathcal{I}_{[\mathbf{B}]}$. This concludes the proof of (c). \square

2.8 Proofs for Section 2.3.2

2.8.1 A technical lemma

We preface our proofs by a version of Lemma 2.3 adapted to the new scheme (2.16).

Lemma 2.27. *Assume A1-A2-R1. Consider the sequence $(\theta_n)_{n \in \mathbb{N}^*}$ satisfying the scheme (2.16). Setting, for any $k \in \mathbb{N}$, $\Delta M_k = \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k}$, we have for any $n \in \mathbb{N}^*$*

and $\varepsilon > 0$,

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1} (\underline{c}_1 - 2L\eta_{k+1} - \bar{c}^2\varepsilon) \|h(\theta_k)\|_{\theta_k}^2 &\leq V(\theta_0) - V(\theta_{n+1}) \\ &+ \sum_{k=0}^n \eta_{k+1} \Delta M_k + \underline{c}_2 \Gamma_{n+1} \\ &+ \sum_{k=0}^n \eta_{k+1} \|\text{grad } V(\theta_k)\|_{\theta_k} \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} \\ &+ 2L \sum_{k=0}^n \eta_{k+1}^2 (\|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2) \\ &+ \sum_{k=0}^n \eta_{k+1} \{(4\varepsilon)^{-1} + 2L\eta_{k+1}\} \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2, \end{aligned}$$

where e, Δ and are defined by (2.8) and (2.17) respectively.

Proof. For any $k \geq 0$ and $t \in [0, 1]$, consider

$$\gamma^{(k)}(t) = \text{Exp}_{\theta_k} \left(t\eta_{k+1} \left\{ H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\} \right).$$

Note that $\dot{\gamma}^{(k)}(0) = \eta_{k+1} \{ H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \}$ and that $\ell(\gamma^{(k)}) = \eta_{k+1} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}$. Then, by Lemma 2.1, (2.8) and using that for any $\theta \in \Theta$, $a, b, c, d \in \text{T}_\theta \Theta$, $\|a+b+c+d\|_\theta^2 \leq 4(\|a\|_\theta^2 + \|b\|_\theta^2 + \|c\|_\theta^2 + \|d\|_\theta^2)$, we get that for any $k \geq 0$,

$$\begin{aligned} &\left| V(\theta_{k+1}) - V(\theta_k) - \eta_{k+1} \left\langle \text{grad } V(\theta_k), H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\rangle_{\theta_k} \right| \\ &\leq (L/2)\ell(\gamma^{(k)})^2 = (L\eta_{k+1}^2/2) \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2, \\ &\leq 2L\eta_{k+1}^2 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2 \right\}. \end{aligned}$$

Therefore, we get that for any $k \in \mathbb{N}$,

$$\begin{aligned} &-\eta_{k+1} \langle \text{grad } V(\theta_k), h(\theta_k) \rangle_{\theta_k} \\ &\leq V(\theta_k) - V(\theta_{k+1}) \\ &+ \eta_{k+1} \left\langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\rangle_{\theta_k} \\ &+ 2L\eta_{k+1}^2 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2 \right\}. \end{aligned} \tag{2.98}$$

By **A2-(b)** and the Cauchy-Schwarz inequality, we obtain for any $k \in \mathbb{N}^*$ and $\varepsilon > 0$,

$$\langle \text{grad } V(\theta_k), b_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \leq \bar{c}^2\varepsilon \|h(\theta_k)\|_{\theta_k}^2 + (1/4\varepsilon) \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2.$$

Thus, using **A2-(a)**, the Cauchy-Schwarz inequality and plugging this in (2.98) gives,

$$\begin{aligned} \eta_{k+1} (\underline{c}_1 - 2L\eta_{k+1} - \bar{c}^2\varepsilon) \|h(\theta_k)\|_{\theta_k}^2 &\leq V(\theta_k) - V(\theta_{k+1}) \\ &+ \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} + \eta_{k+1} \underline{c}_2 \\ &+ \eta_{k+1} \|\text{grad } V(\theta_k)\|_{\theta_k} \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} + 2L\eta_{k+1}^2 \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2 \\ &+ [\eta_{k+1}/4\varepsilon + 2L\eta_{k+1}^2] \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2 + 2L\eta_{k+1}^2 \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2. \end{aligned}$$

Adding these inequalities from 0 to n and rearranging terms concludes the proof. \square

2.8.2 Proof of Theorem 2.11

If **MD2**(a) holds for $a \in \mathbb{N}^*$, then **MD2**(\tilde{a}) holds for any $\tilde{a} \in \{0, 1, \dots, a\}$ and $m_{(\tilde{a})}$ will then stand for a constant such that almost surely, $\mathbb{E}[\|e_{\theta_n}(X_{n+1})\|_{\tilde{\theta}_n}^{\tilde{a}} | \mathcal{F}_n] \leq m_{(\tilde{a})}$. Before giving the proof of Theorem 2.11, we specify the statement of this result in Theorem 2.28 below. In particular, we give an explicit expression of $A_{n+1}^{(R)}$. We define

$$\bar{\eta} = \begin{cases} \underline{c}_1 [4(2L + 3\mathcal{L}^{(1)}\bar{c}h_\infty)]^{-1} & \text{if } \mathbf{R2-MD2(4)} \text{ hold ,} \\ \underline{c}_1 / (8L) & \text{if } \mathbf{R3-MD2(6)} \text{ hold ,} \end{cases} \quad (2.99)$$

$$C_R^{(2)} = \begin{cases} 2\underline{c}_1^{-1}(2L + 3\bar{c}h_\infty\mathcal{L}^{(1)}(\tilde{\theta}))(b_\infty^2 + m_{(2)}) & \text{if } \mathbf{R2-MD2(4)} \text{ hold ,} \\ 4\underline{c}_1^{-1}L(b_\infty^2 + m_{(2)}) & \text{if } \mathbf{R3-MD2(6)} \text{ hold ,} \end{cases} \quad (2.100)$$

$$C_R^{(4)} = 54L\underline{c}_1^{-1}(\mathcal{L}^{(1)}(\tilde{\theta}))^2\{h_\infty^4 + b_\infty^4 + m_{(4)}\} , \quad (2.101)$$

$$C_R^{(3)} = 3^2\mathcal{L}^{(2)}(\tilde{\theta})\bar{c}h_\infty\{h_\infty^3 + b_\infty^3 + m_{(3)}\} , \quad (2.102)$$

$$C_R^{(6)} = 3^5\mathcal{L}^{(2)}(\tilde{\theta})\{b_\infty^6 + h_\infty^6 + m_{(6)}\} , \quad (2.103)$$

where $\tilde{\theta} \in \Theta$ is fixed and can be chosen arbitrary.

Theorem 2.28. *Assume **A1-A2-A3-A4**, **R1**, **MD1** hold. Suppose either **R2**, **MD2(4)** or **R3**, **MD2(6)**. Consider $(\theta_k)_{k \in \mathbb{N}}$ defined by (2.12). Then, if $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta}$, for $\bar{\eta}$ defined in (2.99), for any $n \in \mathbb{N}$*

$$\mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq 2(\underline{c}_1\Gamma_{n+1})^{-1} \{ \mathbb{E}[V(\theta_0)] + C_R^{(2)}\Gamma_{n+1}^{(2)} + A_{n+1}^{(R)} \} + 2(b_\infty\bar{c}h_\infty + \underline{c}_2) / \underline{c}_1 ,$$

$$\text{with } A_{n+1}^{(R)} = \begin{cases} C_R^{(4)}\Gamma_{n+1}^{(4)} & \text{if } \mathbf{R2, MD2(4)} \text{ hold} \\ C_R^{(3)}\Gamma_{n+1}^{(3)} + C_R^{(6)}\Gamma_{n+1}^{(6)} & \text{if } \mathbf{R3, MD2(6)} \text{ hold} \end{cases} , \quad (2.104)$$

where $C_R^{(2)}$ is given by (2.100), $C_R^{(4)}$ is defined in (2.101), $C_R^{(3)}$, $C_R^{(6)}$ are defined in (2.102)-(2.103) and $I_n \in \{0, \dots, n\}$ is a random variable independent of \mathcal{F}_n and distributed according to (2.5).

Proof of Theorem 2.28. In all the proof $\tilde{\theta} \in \Theta$ is a fixed element of Θ , which will be used as we will apply Lemma 2.7. Also for ease of notation, we simply denote $\mathcal{L}_\infty^{(1)} = \mathcal{L}^{(1)}(\tilde{\theta})$ and similarly $\mathcal{L}^{(2)} = \mathcal{L}^{(2)}(\tilde{\theta})$.

Using the assumptions in Theorem 2.11, the proof consists in bounding each term of Lemma 2.27 after taking the expectation. We first make a first estimate which holds if either **R2** or **R3** hold.

Using **MD1** and that θ_k is \mathcal{F}_k -measurable, we have for any $k \in \mathbb{N}$, $\mathbb{E}[\Delta M_k] = 0$. In addition, by **A2-(a)-A4** we have

$$\|\text{grad } V(\theta_k)\|_{\theta_k} \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} \leq \bar{c}h_\infty \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} .$$

Therefore, taking expectation in Lemma 2.27, we get using V is non-negative, **A 3** and **MD 2(4)** or **MD 2(6)**, for any $n \in \mathbb{N}^*$ and $\varepsilon > 0$,

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1} (\underline{c}_1 - 2L\eta_{k+1} - \bar{c}^2\varepsilon) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] + \underline{c}_2\Gamma_{n+1} \\ &\quad + \bar{c}h_\infty \sum_{k=0}^n \eta_{k+1} \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} + 2L(b_\infty^2 + m_{(2)})\Gamma_{n+1}^{(2)} \\ &\quad + \sum_{k=0}^n \eta_{k+1}^2 \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2 + [b_\infty^2/(4\varepsilon)]\Gamma_{n+1}. \end{aligned} \quad (2.105)$$

Therefore, it remains to bound the terms involving the retraction bias Δ , for which we distinguish the two different sets of conditions **R 2-MD 2(4)** or **R 3, MD 2(6)**.

If R 2-MD 2(4) hold. Using Lemma 2.7-(a), **A 3-A 4** and Hölder inequality shows that

$$\begin{aligned} \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} &\leq \mathcal{L}^{(1)}\eta_{k+1} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^2, \\ &\leq 3\mathcal{L}^{(1)}\eta_{k+1} \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \right\} \\ \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2 &\leq 9 \times 3[\mathcal{L}^{(1)}]^2\eta_{k+1}^2 \left\{ h_\infty^4 + b_\infty^4 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^4 \right\}. \end{aligned}$$

Therefore, plugging these estimates in (2.105), using **MD 2(4)**, we get that

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1} (\underline{c}_1 - 2L\eta_{k+1} - \bar{c}^2\varepsilon) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] + \underline{c}_2\Gamma_{n+1} \\ &\quad + 3\mathcal{L}^{(1)}\bar{c}h_\infty \sum_{k=0}^n \eta_{k+1}^2 \left\{ \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] + b_\infty^2 + m_{(2)} \right\} \\ &\quad + 2L(b_\infty^2 + m_{(2)})\Gamma_{n+1}^{(2)} + [b_\infty^2/(4\varepsilon)]\Gamma_{n+1} \\ &\quad + 27\mathcal{L}^{(1)}\{b_\infty^4 + h_\infty^4 + m_{(4)}\}\Gamma_{n+1}^{(4)}. \end{aligned}$$

Rearranging terms and taking $\varepsilon = \underline{c}_1/(4\bar{c}^2)$, we get

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1} (3\underline{c}_1/4 - (2L + 3\mathcal{L}^{(1)}\bar{c}h_\infty)\eta_{k+1}) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] \\ \leq \mathbb{E}[V(\theta_0)] + (\underline{c}_2 + b_\infty^2\bar{c}^2/\underline{c}_1)\Gamma_{n+1} \\ + (2L + 3\mathcal{L}^{(1)}\bar{c}h_\infty)(b_\infty^2 + m_{(2)})\Gamma_{n+1}^{(2)} + 27\mathcal{L}^{(1)}\{b_\infty^4 + h_\infty^4 + m_{(4)}\}\Gamma_{n+1}^{(4)}. \end{aligned}$$

Using $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta} \leq \underline{c}_1[4(2L + 3\mathcal{L}^{(1)}\bar{c}h_\infty)]^{-1}$ ensures that $\underline{c}_1/4 \geq (2L + 3\mathcal{L}^{(1)}\bar{c}h_\infty)\eta_{k+1}$ for any $k \in \mathbb{N}$ and the proof follows.

If R 3, MD 2(6) hold. Using Lemma 2.7-(b), **A 3-A 4** and Hölder inequality shows that

$$\begin{aligned} \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} &\leq \mathcal{L}^{(2)}\eta_{k+1}^2 \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^3, \\ &\leq 3^2\mathcal{L}^{(2)}\eta_{k+1}^2 \left\{ h_\infty^3 + b_\infty^3 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^3 \right\} \\ \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2 &\leq 3^5[\mathcal{L}^{(2)}]^2\eta_{k+1}^2 \left\{ h_\infty^6 + b_\infty^6 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^6 \right\}. \end{aligned}$$

Therefore, plugging these estimates in (2.105), using **MD 2(6)**, we get that

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1} (\underline{c}_1 - 2L\eta_{k+1} - \bar{c}^2\varepsilon) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] + \underline{c}_2\Gamma_{n+1} \\ &\quad + 3^2\mathcal{L}^{(2)}\bar{c}h_\infty \left\{ h_\infty^3 + b_\infty^3 + m_{(3)} \right\} \Gamma_{n+1}^{(3)} + 2L(b_\infty^2 + m_{(2)})\Gamma_{n+1}^{(2)} \\ &\quad + 3^5\mathcal{L}^{(2)}\{b_\infty^6 + h_\infty^6 + m_{(6)}\}\Gamma_{n+1}^{(6)} + [b_\infty^2/(4\varepsilon)]\Gamma_{n+1}. \end{aligned}$$

Rearranging terms and taking $\varepsilon = \underline{c}_1/(4\bar{c}^2)$, we get

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1} (3\underline{c}_1/4 - 2L\eta_{k+1}) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] + (\underline{c}_2 + b_\infty^2 \bar{c}^2 / \underline{c}_1) \Gamma_{n+1} \\ &\quad + 2L(b_\infty^2 + m_{(2)}) \Gamma_{n+1}^{(2)} \\ &\quad + 3^2 \mathcal{L}^{(2)} \bar{c} h_\infty \{h_\infty^3 + b_\infty^3 + m_{(3)}\} \Gamma_{n+1}^{(3)} \\ &\quad + 3^5 \mathcal{L}^{(2)} \{b_\infty^6 + h_\infty^6 + m_{(6)}\} \Gamma_{n+1}^{(6)}. \end{aligned}$$

Using that $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta} \leq \underline{c}_1/(8L)$ ensures that $\underline{c}_1/4 \geq 2L\eta_{k+1}$ for any $k \in \mathbb{N}$ which completes the proof. \square

2.8.3 Proof of Theorem 2.12

Under **MA4**(w), for any $p \in [1, 3]$, $C_w^{(p)}$ stands for a constant such that

$$\sup_{k \in \mathbb{N}} \mathbb{E}[w^p(X_{k+1})] \leq C_w^{(p)}.$$

In addition, in the sequel, $\tilde{\theta} \in \Theta$ is a fixed element of Θ , which will be used as we will apply Lemma 2.7. Also for ease of notation, we simply denote $\mathcal{L}_\infty^{(1)} = \mathcal{L}^{(1)}(\tilde{\theta})$ and similarly $\mathcal{L}^{(2)} = \mathcal{L}^{(2)}(\tilde{\theta})$.

Before giving the proof of Theorem 2.12, we specify the statement of this result in Theorem 2.29 below. In particular, we give an explicit expression of $B_{n+1}^{(R)}$. We define

$$\bar{\eta} = \begin{cases} \underline{c}_1/[4(2L + D_{\hat{\varepsilon}} + 3\mathcal{L}^{(1)}\bar{c}h_\infty)]^{-1} & \text{if } \mathbf{R2} \text{ holds,} \\ \underline{c}_1/[4(2L + D_{\hat{\varepsilon}})]^{-1} & \text{if } \mathbf{R3} \text{ holds,} \end{cases} \quad (2.106)$$

$$D_R^{(2)} = \begin{cases} C_{\hat{\varepsilon}}^0 + (b_\infty^2 + C_w e_\infty^2)(3\bar{c}h_\infty \mathcal{L}^{(1)} + 2L) & \text{if } \mathbf{R2} \text{ holds,} \\ C_{\hat{\varepsilon}}^0 + 2L(b_\infty^2 + C_w e_\infty^2) & \text{if } \mathbf{R3} \text{ holds,} \end{cases} \quad (2.107)$$

$$D_{R2}^{(3)} = 3\mathcal{L}^{(1)}[L\hat{e}_\infty C_w \{h_\infty^2 + b_\infty^2 + e_\infty^2 C_w^{(2)}\} + L_{\hat{\varepsilon}} \bar{c} h_\infty \{h_\infty^2 + b_\infty^2 + e_\infty^2 C_w\}], \quad (2.108)$$

$$D_R^{(4)} = 3^2 \mathcal{L}^{(2)}[L\hat{e}_\infty C_w \{h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3)}\} + L_{\hat{\varepsilon}} \bar{c} h_\infty \{h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3/2)}\}], \quad (2.109)$$

$$D_{R3}^{(3)} = 3^2 \mathcal{L}^{(2)} \bar{c} h_\infty \{h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3/2)}\}, \quad (2.110)$$

$$D_R^{(6)} = 23^5 L [\mathcal{L}^{(2)}]^2 \{h_\infty^6 + b_\infty^6 + e_\infty^6 C_w^{(3)}\}. \quad (2.111)$$

Theorem 2.29. *Assume **A1-A2-A3-A4**, **R1**, **MA1-MA2**(w)-**MA4**(w) hold for some measurable function $w : \mathcal{X} \rightarrow [1, +\infty)$. Suppose in addition that **R2** or **R3** holds. Assume that $(\eta_k)_{k \in \mathbb{N}^*}$ is a sequence of stepsizes and $a_1, a_2 \geq 0$ satisfying (2.10).*

Consider $(\theta_k)_{k \in \mathbb{N}}$ defined by (2.12). Then, if $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta}$, for $\bar{\eta}$ defined in (2.106), for any $n \in \mathbb{N}$,

$$\begin{aligned} \mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] &\leq 2(\underline{c}_1 \Gamma_{n+1})^{-1} \{ \mathbb{E}[V(\theta_0)] + C(\eta_1) \\ &\quad + D_R^{(2)} \Gamma_{n+1}^{(2)} + B_{n+1}^{(R)} \} + 2(b_\infty \bar{c} h_\infty + \underline{c}_2) / \underline{c}_1, \\ \text{with } B_{n+1}^{(R)} &= \begin{cases} D_{R2}^{(3)} \Gamma_{n+1}^{(3)} + D_R^{(4)} \Gamma_{n+1}^{(4)} & \text{if } \mathbf{R}2 \text{ holds,} \\ D_{R3}^{(3)} \Gamma_{n+1}^{(3)} + D_R^{(6)} \Gamma_{n+1}^{(6)} & \text{if } \mathbf{R}3 \text{ holds,} \end{cases} \end{aligned} \quad (2.112)$$

where $C(\eta_1), D_R^{(2)}$ are given by (2.39)-(2.107); $D_{R2}^{(3)}, D_R^{(4)}$ are defined in (2.108)-(2.109); $D_{R3}^{(3)}, D_R^{(6)}$ are defined in (2.110)-(2.111) and $I_n \in \{0, \dots, n\}$ is a random variable independent of \mathcal{F}_n and distributed according to (2.5).

We begin the proof by showing a similar lemma to Lemma 2.16.

Lemma 2.30. Assume **A1-A2-A3-A4-MA1-MA2(w)-MA4(w)** hold for some measurable function $w : \mathcal{X} \rightarrow [1, +\infty)$. Suppose in addition that **R2** or **R3** holds. Let $(\eta_k)_{k \in \mathbb{N}^*}$ be a sequence satisfying (2.10) and consider $(\theta_k)_{k \in \mathbb{N}}$ defined by (2.16). Set for any $k \in \mathbb{N}$, $\Delta M_k = \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k}$. It holds that for any $n \in \mathbb{N}$,

$$|\mathbb{E}[\sum_{k=0}^n \eta_{k+1} \Delta M_k]| \leq D_{\hat{e}} \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] + C_{\hat{e}}^0 \Gamma_{n+1}^{(2)} + C(\eta_1) + B_{R,0},$$

where $D_{\hat{e}}, C_{\hat{e}}^0, C(\eta_1)$ are given in (2.41)-(2.39) and

$$\begin{aligned} B_{R,0} &= \begin{cases} E_{\text{Ret}}^{(3)} \Gamma_{n+1}^{(3)} & \text{if } \mathbf{R}2 \text{ holds,} \\ E_{\text{Ret}}^{(4)} \Gamma_{n+1}^{(4)} & \text{if } \mathbf{R}3 \text{ holds,} \end{cases} \\ E_{\text{Ret}}^{(3)} = D_{R2}^{(3)} &= 3\mathcal{L}^{(1)}[L\hat{e}_\infty C_w \{h_\infty^2 + b_\infty^2 + e_\infty^2 C_w^{(2)}\} \\ &\quad + L\bar{c} h_\infty \{h_\infty^2 + b_\infty^2 + e_\infty^2 C_w\}], \\ E_{\text{Ret}}^{(4)} = D_{R3}^{(4)} &= 3^2 \mathcal{L}^{(2)}[L\hat{e}_\infty C_w \{h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3)}\} \\ &\quad + L\bar{c} h_\infty \{h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3/2)}\}]. \end{aligned}$$

Proof. The proof is an adaptation of Lemma 2.16 in which we need to deal with the retraction bias. Consider the measurable function $\hat{e} : \Theta \times \mathcal{X} \rightarrow \text{T}\Theta$ which satisfies **MA2(w)** and for any $k \in \mathbb{N}$. By (2.12), we can consider $\gamma^{(k+1)} : [0, 1] \rightarrow \Theta$, the geodesic between θ_k and θ_{k+1} defined by $\gamma^{(k+1)}(t) = \text{Exp}\{t\eta_{k+1}(H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}))\}$ for any $t \in [0, 1]$. Note that for any $k \in \mathbb{N}$,

$$\begin{aligned} \ell(\gamma^{(k+1)}) &= \eta_{k+1} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}, \\ &\leq \eta_{k+1} \{ \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k} + \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} \}. \end{aligned} \quad (2.113)$$

Using that the parallel transport associated with the Levi-Civita connection is a linear isometry [Lee19, Proposition 5.5] and $(\text{T}_{01}^Y)^{-1} = \text{T}_{10}^Y$ by uniqueness of parallel transport

[Lee19, Theorem 4.32], we obtain the following decomposition

$$\mathbb{E} \left[- \sum_{k=0}^n \eta_{k+1} \Delta M_k \right] = - \mathbb{E} \left[\sum_{i=1}^5 A_i \right], \quad (2.114)$$

where

$$\begin{aligned} A_1 &= \sum_{k=1}^n \eta_{k+1} \langle \text{grad } V(\theta_k), \hat{e}_{\theta_k}(X_{k+1}) - P_{\theta_k} \hat{e}_{\theta_k}(X_k) \rangle_{\theta_k}, \\ A_2 &= \sum_{k=1}^n \eta_{k+1} \langle \text{grad } V(\theta_k), P_{\theta_k} \hat{e}_{\theta_k}(X_k) - T_{01}^{\gamma(k)} P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \rangle_{\theta_k}, \\ A_3 &= \sum_{k=1}^n \eta_{k+1} \langle T_{10}^{\gamma(k)} \text{grad } V(\theta_k) - \text{grad } V(\theta_{k-1}), P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \rangle_{\theta_{k-1}}, \\ A_4 &= \sum_{k=1}^n (\eta_{k+1} - \eta_k) \langle \text{grad } V(\theta_{k-1}), P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \rangle_{\theta_{k-1}}, \\ A_5 &= \eta_1 \langle \text{grad } V(\theta_0), \hat{e}_{\theta_0}(X_1) \rangle_{\theta_0} - \eta_{n+1} \langle \text{grad } V(\theta_n), \hat{e}_{\theta_n}(X_{n+1}) \rangle_{\theta_n}. \end{aligned}$$

We now bound each term of this decomposition. As for the proof of (2.115) in the proof of Lemma 2.16, using using **MA2(w)-(a)**, we have

$$\mathbb{E}[A_1] = 0. \quad (2.115)$$

Similarly, using the same argument as (2.45) and (2.46) in the proof of Lemma 2.16, we get

$$\begin{aligned} \mathbb{E}[|A_4|] &\leq 2\bar{c}\hat{e}_\infty C_w \eta_1 + 2\bar{c}a_2 \sum_{k=1}^n \eta_k \mathbb{E} \left[\|h(\theta_{k-1})\|_{\theta_{k-1}}^2 \right], \\ \mathbb{E}[|A_5|] &\leq \bar{c}\hat{e}_\infty \left(2 + C_w \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \right). \end{aligned} \quad (2.116)$$

It remains to deal with A_2 and A_3 for which we distinguish the case where **R2** or **R3** holds.

In the case where R2 holds. Using Lemma 2.7-(a), **A3-A4** and Hölder inequality shows that

$$\begin{aligned} \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} &\leq \mathcal{L}^{(1)} \eta_{k+1} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^2, \\ &\leq 3\mathcal{L}^{(1)} \eta_{k+1} \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \right\}. \end{aligned}$$

Using the Cauchy-Schwarz inequality, (2.113), **A2-(b)**, the definition of e (2.8), **A3**,

MA2(w)-(a)-(c) and Jensen's inequality, we obtain

$$\begin{aligned}
 |A_3| &\leq L \sum_{k=1}^n \eta_{k+1} \ell(\gamma^{(k)}) \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \\
 &\leq L \sum_{k=1}^n \eta_k \eta_{k+1} \left\{ \|H_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right\} \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \\
 &\quad + L \sum_{k=1}^n \eta_k \eta_{k+1} \|\Delta_{\theta_{k-1}, \eta_k}(X_k)\|_{\theta_{k-1}} \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \\
 &\leq L \hat{e}_\infty \sum_{k=1}^n \eta_k \eta_{k+1} \left(e_\infty w^{1/2}(X_k) + b_\infty + \|h(\theta_{k-1})\|_{\theta_{k-1}} \right) P_{\theta_{k-1}} w^{1/2}(X_k) \\
 &\quad + 3\mathcal{L}^{(1)} L \hat{e}_\infty \sum_{k=1}^n \eta_k^2 \eta_{k+1} \left\{ \|h(\theta_{k-1})\|_{\theta_{k-1}}^2 + b_\infty^2 + e_\infty^2 w(X_k) \right\} P_{\theta_{k-1}} w^{1/2}(X_k).
 \end{aligned}$$

Taking the expectation, using **MA 2**(w)-(b), **A 4**, that $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (2.10) and the Cauchy-Schwarz inequality brings,

$$\begin{aligned}
 \mathbb{E}[|A_3|] &\leq L \hat{e}_\infty C_w \{ [e_\infty + b_\infty] \Gamma_{n+1}^{(2)} + \sum_{k=1}^n \eta_k^2 \mathbb{E}[\|h(\theta_{k-1})\|_{\theta_{k-1}}^2] \} \\
 &\quad + 3L \hat{e}_\infty \mathcal{L}^{(1)} C_w [h_\infty^2 + b_\infty^2 + e_\infty^2 C_w^{(2)}] \Gamma_{n+1}^{(3)}. \quad (2.117)
 \end{aligned}$$

It remains to treat A_2 depending on the additional two conditions we consider. We start by proving a general bound which hold. Using the Cauchy-Schwarz inequality, **MA2**(w)-(a)-(c)-(d), (2.113) and **A 2**(a), we get,

$$\begin{aligned}
 |A_2| &\leq A_{2,1} + A_{2,2} \\
 A_{2,1} &= L \hat{e} \sum_{k=1}^n \eta_{k+1} \eta_k w^{1/2}(X_k) \|\text{grad } V(\theta_k)\|_{\theta_k} \left\{ \|H_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right. \\
 &\quad \left. + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right\}, \\
 &\leq L \hat{e} \bar{c} \sum_{k=1}^n \eta_{k+1} \eta_k w^{1/2}(X_k) \|h(\theta_k)\|_{\theta_k} \left\{ e_\infty w^{1/2}(X_k) \right. \\
 &\quad \left. + \|h(\theta_{k-1})\|_{\theta_{k-1}} + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right\}, \\
 A_{2,2} &= L \hat{e} \sum_{k=1}^n \eta_{k+1} \eta_k w^{1/2}(X_k) \|\text{grad } V(\theta_k)\|_{\theta_k} \left\| \Delta_{\theta_{k-1}, \eta_k}(X_k) \right\|_{\theta_k} \\
 &\leq 3L \hat{e} \mathcal{L}^{(1)} \sum_{k=1}^n \eta_{k+1} \eta_k^2 \|\text{grad } V(\theta_k)\|_{\theta_k} \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + e_\infty^2 w(X_k) \right\}.
 \end{aligned} \tag{2.118}$$

Similarly to the proof of (2.48) in the proof of Lemma 2.16, we have

$$\mathbb{E}[|A_{1,2}|] \leq L \hat{e} \bar{c} \left[(e_\infty + b_\infty) \{ h_\infty^2 a_1 + 1 \} + h_\infty^2 \right] C_w \Gamma_{n+1}^{(2)}. \quad (2.119)$$

Using **A2-(a)**, **MA2(w)-(a)-(b)**, **A4**, and $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (2.10), we get

$$\mathbb{E} [|A_{2,2}|] \leq 3L\hat{c}\mathcal{L}^{(1)}h_\infty\{h_\infty^2 + b_\infty^2 + C_w e_\infty^2\}\Gamma_{n+1}^{(3)}. \quad (2.120)$$

Combining (2.115)-(2.116)-(2.117)-(2.118) (2.119) and (2.120) with (2.114) completes the proof.

If R3 holds. Using Lemma 2.7-(b), **A3-A4** and Hölder inequality shows that

$$\begin{aligned} \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} &\leq \mathcal{L}^{(2)}\eta_{k+1}^2 \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^3, \\ &\leq 3^2 \mathcal{L}^{(2)}\eta_{k+1}^2 \left\{ h_\infty^3 + b_\infty^3 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^3 \right\}. \end{aligned}$$

Using the Cauchy-Schwarz inequality, (2.113), **A2-(b)**, the definition of e (2.8), **A3**, **MA2(w)-(a)-(c)**, and Jensen's inequality, we obtain

$$\begin{aligned} |A_3| &\leq L \sum_{k=1}^n \eta_{k+1} \ell(\gamma^{(k)}) \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \\ &\leq L \sum_{k=1}^n \eta_k \eta_{k+1} \left\{ \|H_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right\} \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \\ &\quad + L \sum_{k=1}^n \eta_k \eta_{k+1} \|\Delta_{\theta_{k-1}, \eta_k}(X_k)\|_{\theta_{k-1}} \|P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \\ &\leq L\hat{e}_\infty \sum_{k=1}^n \eta_k \eta_{k+1} \left(e_\infty w^{1/2}(X_k) + b_\infty + \|h(\theta_{k-1})\|_{\theta_{k-1}} \right) P_{\theta_{k-1}} w^{1/2}(X_k) \\ &\quad + 3^2 \mathcal{L}^{(2)} L\hat{e}_\infty \sum_{k=1}^n \eta_k^3 \eta_{k+1} \left\{ h_\infty^3 + b_\infty^3 + e_\infty^3 w^{3/2}(X_k) \right\} P_{\theta_{k-1}} w^{1/2}(X_k). \end{aligned}$$

Taking the expectation, using **MA2(w)-(b)**, **A4**, **MA4**, that $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (2.10) and the Cauchy-Schwarz inequality brings,

$$\begin{aligned} \mathbb{E} [|A_3|] &\leq L\hat{e}_\infty C_w \{ [e_\infty + b_\infty] \Gamma_{n+1}^{(2)} + \sum_{k=1}^n \eta_k^2 \mathbb{E} [\|h(\theta_{k-1})\|_{\theta_{k-1}}^2] \} \\ &\quad + 3^2 L\hat{e}_\infty \mathcal{L}^{(2)} C_w [h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3)}] \Gamma_{n+1}^{(4)}. \quad (2.121) \end{aligned}$$

It remains to treat A_2 depending on the additional two conditions we consider. We start by proving a general bound which hold. Using the Cauchy-Schwarz inequality, **MA2(w)-**

(a)-(c)-(d), (2.113) and A2-(a), we get,

$$\begin{aligned}
 |A_2| &\leq A_{2,1} + A_{2,2} \\
 A_{2,1} &\leq L_{\hat{c}} \sum_{k=1}^n \eta_{k+1} \eta_k w^{1/2}(X_k) \|h(\theta_k)\|_{\theta_k} \{e_{\infty} w^{1/2}(X_k) + \|h(\theta_{k-1})\|_{\theta_{k-1}} \\
 &\quad + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}}\}, \\
 A_{2,2} &= L_{\hat{c}} \sum_{k=1}^n \eta_{k+1} \eta_k w^{1/2}(X_k) \|\text{grad } V(\theta_k)\|_{\theta_k} \|\Delta_{\theta_{k-1}, \eta_k}(X_k)\|_{\theta_k} \\
 &\leq 3^2 L_{\hat{c}} \mathcal{L}^{(2)} \sum_{k=1}^n \eta_{k+1} \eta_k^3 \|\text{grad } V(\theta_k)\|_{\theta_k} \{h_{\infty}^3 + b_{\infty}^3 + e_{\infty} w^{3/2}(X_k)\}.
 \end{aligned} \tag{2.122}$$

Similarly to the proof of (2.48) in the proof of Lemma 2.16, we have

$$\mathbb{E}[|A_{1,2}|] \leq L_{\hat{c}} \bar{c} \left[(e_{\infty} + b_{\infty}) \{h_{\infty}^2 a_1 + 1\} + h_{\infty}^2 \right] C_w \Gamma_{n+1}^{(2)}. \tag{2.123}$$

Using A2-(a), MA2(w)-(a)-(b), MA4, A4, and $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (2.10), we get

$$\mathbb{E}[|A_{2,2}|] \leq 3^2 L_{\hat{c}} \bar{c} \mathcal{L}^{(2)} h_{\infty} \{h_{\infty}^3 + b_{\infty}^3 + C_w^{(3/2)} e_{\infty}^3\} \Gamma_{n+1}^{(4)}. \tag{2.124}$$

Combining (2.115)-(2.116)-(2.121)-(2.122) (2.123) and (2.124) with (2.114) completes the proof. \square

Proof of Theorem 2.29. Taking the expectation in the bound provided by Lemma 2.27 and $\varepsilon = \underline{c}_1/(4\bar{c}^2)$, and using V is non-negative, A2-(a), A3, MA2(w)-(a)-(b), we get

$$\begin{aligned}
 \sum_{k=0}^n \eta_{k+1} (3\underline{c}_1/4 - 2L\eta_{k+1}) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] + \mathbb{E}[\sum_{k=0}^n \eta_{k+1} \Delta M_k] + \underline{c}_2 \Gamma_{n+1} \\
 &\quad + \bar{c} h_{\infty} \sum_{k=0}^n \eta_{k+1} \mathbb{E}[\|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}] \\
 &\quad + 2Le_{\infty}^2 C_w \Gamma_{n+1}^{(2)} + 2L \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E}[\|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2] \\
 &\quad + b_{\infty}^2 \{\bar{c}^2/\underline{c}_1 \Gamma_{n+1} + 2L\Gamma_{n+1}^{(2)}\}.
 \end{aligned} \tag{2.125}$$

A bound on $\mathbb{E}[\sum_{k=0}^n \eta_{k+1} \Delta M_k]$ is provided by Lemma 2.30. It remains therefore to deal with the two terms $\sum_{k=0}^n \eta_{k+1} \mathbb{E}[\|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}]$ and $\sum_{k=0}^n \eta_{k+1}^2 \mathbb{E}[\|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2]$. These two terms will be treated differently depending if R2 or R3 holds.

If R2 holds. Using Lemma 2.7-(a), A3, A4, MA2(w)-(a) and MA4(w), and Hölder inequality shows that

$$\begin{aligned}
 \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} &\leq \mathcal{L}^{(1)} \eta_{k+1} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^2, \\
 \mathbb{E}[\|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}] &\leq 3\mathcal{L}^{(1)} \eta_{k+1} \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_{\infty}^2 + e_{\infty}^2 C_w \right\} \\
 \mathbb{E}[\|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2] &\leq 9 \times 3[\mathcal{L}^{(1)}]^2 \eta_{k+1}^2 \left\{ h_{\infty}^4 + b_{\infty}^4 + e_{\infty}^4 C_w^{(2)} \right\}.
 \end{aligned}$$

Therefore, plugging these estimates in (2.125) and using Lemma 2.30, we get that

$$\begin{aligned}
 \sum_{k=0}^n \eta_{k+1} (3\underline{c}_1/4 - 2L\eta_{k+1}) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] + (\underline{c}_2 + b_\infty^2 \bar{c}^2 / \underline{c}_1) \Gamma_{n+1} \\
 &+ D_{\hat{\varepsilon}} \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] + C_{\hat{\varepsilon}}^0 \Gamma_{n+1}^{(2)} + C(\eta_1) + E_{\text{Ret}}^{(3)} \Gamma_{n+1}^{(3)} \\
 &+ 3\mathcal{L}^{(1)} \bar{c} h_\infty \sum_{k=0}^n \eta_{k+1}^2 \|h(\theta_k)\|_{\theta_k} \\
 &+ (b_\infty^2 + C_w e_\infty^2) (3\bar{c} h_\infty \mathcal{L}^{(1)} + 2L) \Gamma_{n+1}^{(2)} \\
 &+ 2 \times 3^3 L [\mathcal{L}^{(1)}]^2 \{h_\infty^4 + b_\infty^4 + e_\infty^4 C_w^{(2)}\} \Gamma_{n+1}^{(4)}.
 \end{aligned}$$

Rearranging terms and using $\sup_{k \in \mathbb{N}^*} \eta_k < \bar{\eta}$, with $\bar{\eta}$ satisfying (2.106), we get

$$\begin{aligned}
 (\underline{c}_1/2) \sum_{k=0}^n \eta_{k+1} \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] + (\underline{c}_2 + b_\infty^2 \bar{c}^2 / \underline{c}_1) \Gamma_{n+1} \\
 &+ [C_{\hat{\varepsilon}}^0 + (b_\infty^2 + C_w e_\infty^2) (3\bar{c} h_\infty \mathcal{L}^{(1)} + 2L)] \Gamma_{n+1}^{(2)} + C(\eta_1) \\
 &+ E_{\text{Ret}}^{(3)} \Gamma_{n+1}^{(3)} + 2 \times 3^3 L [\mathcal{L}^{(1)}]^2 \{h_\infty^4 + b_\infty^4 + e_\infty^4 C_w^{(2)}\} \Gamma_{n+1}^{(4)},
 \end{aligned}$$

and taking $\varepsilon = \underline{c}_1 / (4\bar{c}^2)$, we get

$$\begin{aligned}
 \sum_{k=0}^n \eta_{k+1} (3\underline{c}_1/4 - (2L + 3\mathcal{L}^{(1)} \bar{c} h_\infty) \eta_{k+1}) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] \\
 &+ (\underline{c}_2 + b_\infty^2 \bar{c}^2 / \underline{c}_1) \Gamma_{n+1} + (2L + 3\mathcal{L}^{(1)} \bar{c} h_\infty) (b_\infty^2 + m_{(2)}) \Gamma_{n+1}^{(2)} \\
 &+ 27\mathcal{L}^{(1)} \{b_\infty^4 + h_\infty^4 + m_{(4)}\} \Gamma_{n+1}^{(4)}.
 \end{aligned}$$

Using $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta} \leq \underline{c}_1 [4(2L + 3\mathcal{L}^{(1)} \bar{c} h_\infty)]^{-1}$ ensures that $\underline{c}_1/4 \geq (2L + 3\mathcal{L}^{(1)} \bar{c} h_\infty) \eta_{k+1}$ for any $k \in \mathbb{N}$ and the proof follows.

If R3 holds. Using Lemma 2.7-(b), A3, A4, MA2(w)-(a) and MA4(w), and Hölder inequality shows that

$$\begin{aligned}
 \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} &\leq \mathcal{L}^{(2)} \eta_{k+1}^2 \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^3, \\
 \mathbb{E}[\|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}] &\leq 3^2 \mathcal{L}^{(2)} \eta_{k+1}^2 \{h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3/2)}\}, \\
 \mathbb{E}[\|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2] &\leq 3^5 [\mathcal{L}^{(2)}]^2 \eta_{k+1}^2 \{h_\infty^6 + b_\infty^6 + e_\infty^6 C_w^{(3)}\}.
 \end{aligned}$$

Therefore, plugging these estimates in (2.125) and using Lemma 2.30, we get that

$$\begin{aligned}
 \sum_{k=0}^n \eta_{k+1} (3\underline{c}_1/4 - 2L\eta_{k+1}) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] &\leq \mathbb{E}[V(\theta_0)] + (\underline{c}_2 + b_\infty^2 \bar{c}^2 / \underline{c}_1) \Gamma_{n+1} \\
 &+ D_{\hat{\varepsilon}} \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] + C_{\hat{\varepsilon}}^0 \Gamma_{n+1}^{(2)} + C(\eta_1) + E_{\text{Ret}}^{(4)} \Gamma_{n+1}^{(4)} \\
 &+ 2L(b_\infty^2 + C_w e_\infty^2) \Gamma_{n+1}^{(2)} \\
 &+ 3^2 \mathcal{L}^{(2)} \bar{c} h_\infty \{h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3/2)}\} \Gamma_{n+1}^{(3)} \\
 &+ 2 \times 3^5 L [\mathcal{L}^{(2)}]^2 \{h_\infty^6 + b_\infty^6 + e_\infty^6 C_w^{(3)}\} \Gamma_{n+1}^{(6)}.
 \end{aligned}$$

Rearranging terms and using $\sup_{k \in \mathbb{N}^*} \eta_k < \bar{\eta}$, with $\bar{\eta}$ satisfying (2.106), we get

$$\begin{aligned} & (\underline{c}_1/2) \sum_{k=0}^n \eta_{k+1} \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] \\ & \leq \mathbb{E}[V(\theta_0)] + (\underline{c}_2 + b_\infty^2 \bar{c}^2 / \underline{c}_1) \Gamma_{n+1} + [C_\varepsilon^0 + 2L(b_\infty^2 + C_w e_\infty^2)] \Gamma_{n+1}^{(2)} + C(\eta_1) \\ & \quad + E_{\text{Ret}}^{(4)} \Gamma_{n+1}^{(4)} + 3^2 \mathcal{L}^{(2)} \bar{c} h_\infty \left\{ h_\infty^3 + b_\infty^3 + e_\infty^3 C_w^{(3/2)} \right\} \Gamma_{n+1}^{(3)} \\ & \quad + 2 \times 3^5 L [\mathcal{L}^{(2)}]^2 \{ h_\infty^6 + b_\infty^6 + e_\infty^6 C_w^{(3)} \} \Gamma_{n+1}^{(6)}. \end{aligned}$$

Therefore, plugging these estimates in (2.105), using MD2(6), we get that

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1} (\underline{c}_1 - 2L\eta_{k+1} - \bar{c}^2 \varepsilon) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] & \leq \mathbb{E}[V(\theta_0)] + \underline{c}_2 \Gamma_{n+1} \\ & \quad + 3^2 \mathcal{L}^{(2)} \bar{c} h_\infty \left\{ h_\infty^3 + b_\infty^3 + m_{(3)} \right\} \Gamma_{n+1}^{(3)} + 2L(b_\infty^2 + m_{(2)}) \Gamma_{n+1}^{(2)} \\ & \quad + 3^5 \mathcal{L}^{(2)} \{ b_\infty^6 + h_\infty^6 + m_{(6)} \} \Gamma_{n+1}^{(6)} + [b_\infty^2 / (4\varepsilon)] \Gamma_{n+1}. \end{aligned}$$

Rearranging terms and taking $\varepsilon = \underline{c}_1 / (4\bar{c}^2)$, we get

$$\begin{aligned} \sum_{k=0}^n \eta_{k+1} (3\underline{c}_1/4 - 2L\eta_{k+1}) \mathbb{E}[\|h(\theta_k)\|_{\theta_k}^2] & \leq \mathbb{E}[V(\theta_0)] + (\underline{c}_2 + b_\infty^2 \bar{c}^2 / \underline{c}_1) \Gamma_{n+1} \\ & \quad + 2L(b_\infty^2 + m_{(2)}) \Gamma_{n+1}^{(2)} + 3^2 \mathcal{L}^{(2)} \bar{c} h_\infty \left\{ h_\infty^3 + b_\infty^3 + m_{(3)} \right\} \Gamma_{n+1}^{(3)} \\ & \quad + 3^5 \mathcal{L}^{(2)} \{ b_\infty^6 + h_\infty^6 + m_{(6)} \} \Gamma_{n+1}^{(6)}. \end{aligned}$$

Using that $\sup_{k \in \mathbb{N}^*} \eta_k \leq \bar{\eta} \leq \underline{c}_1 / (8L)$ ensures that $\underline{c}_1/4 \geq 2L\eta_{k+1}$ for any $k \in \mathbb{N}$ which completes the proof. \square

2.9 Proof of Section 2.4

2.9.1 Proofs of Section 2.4.1

Lemma 2.31. *Consider the Grassmann manifold $\text{Gr}_r(\mathbb{R}^d)$ defined in Example 2.9 and the objective function f defined in (2.18). Then, for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$,*

(a) *the covariant derivative of f at $[\mathbf{B}]$ is given by*

$$\nabla f([\mathbf{B}]) : \mathbf{D} \mapsto \nabla_{\mathbf{D}} f([\mathbf{B}]) = -\text{Tr} \left(\mathbf{A} \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \right).$$

and the Riemannian gradient of f at $[\mathbf{B}]$ is given by

$$\text{grad } f([\mathbf{B}]) = -\text{Proj}_{[\mathbf{B}_\perp]} \mathbf{A} \mathbf{B},$$

where $\text{Proj}_{[\mathbf{B}_\perp]} = \mathbf{I}_d - \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top$ is the projection matrix on $[\mathbf{B}_\perp]$, the orthogonal complement of $[\mathbf{B}]$.

(b) for $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, the covariant Hessian of f at $[\mathbf{B}]$ is given by

$$\begin{aligned} \text{Hess } f([\mathbf{B}])\{(\mathbf{D}, \mathbf{D})\} &= -\text{Tr} \left(\mathbf{A}\mathbf{D} \left(\mathbf{B}^\top \mathbf{B} \right)^{-1} \mathbf{D}^\top \right) \\ &\quad + \text{Tr} \left(\mathbf{A}\mathbf{B} \left(\mathbf{B}^\top \mathbf{B} \right)^{-1} \mathbf{D}^\top \mathbf{D} \left(\mathbf{B}^\top \mathbf{B} \right)^{-1} \mathbf{B}^\top \right). \end{aligned} \quad (2.126)$$

Alternatively, one can also express the Hessian of f at $[\mathbf{B}]$ as follows,

$$\text{Hess } f([\mathbf{B}])\{(\mathbf{D}, \mathbf{D})\} = \left\langle \mathbf{D} \left(\mathbf{B}^\top \mathbf{B} \right)^{-1} \mathbf{B}^\top \mathbf{A}\mathbf{B} - \text{Proj}_{[\mathbf{B}_\perp]} \mathbf{A}\mathbf{D}, \mathbf{D} \right\rangle_{[\mathbf{B}]} . \quad (2.127)$$

Proof of Lemma 2.31. The main idea of the proof is to use Theorem 2.37 and to characterize the iterated total derivatives of f using its gradient and Hessian. For any $([\mathbf{B}], \mathbf{D}) \in \mathbb{T}\text{Gr}_r(\mathbb{R}^d)$, consider the geodesic $\gamma : [0, 1] \rightarrow \text{Gr}_r(\mathbb{R}^d)$ such that $\gamma(0) = [\mathbf{B}]$, $\dot{\gamma}(0) = \mathbf{D}$, which is given by $\gamma(t) = \text{Exp}_{[\mathbf{B}]}(t\mathbf{D})$. Then, using Theorem 2.37 for the function f , the curve γ and $n = 2$, we compute the Taylor expansion of the map $f \circ \gamma : [0, 1] \rightarrow \mathbb{R}$. Using Section 2.11.8 and (2.163), we get that $(f \circ \gamma)'(0) = \nabla_{\dot{\gamma}(0)} f(\gamma(0)) = \nabla_{\mathbf{D}} f([\mathbf{B}])$ and $(f \circ \gamma)''(0) = \text{Hess } f([\mathbf{B}])\{(\mathbf{D}, \mathbf{D})\}$. Introducing this into the Taylor expansion gives, for any $t \in [0, 1]$,

$$\begin{aligned} f \circ \gamma(t) &= f([\mathbf{B}]) + t \nabla_{\mathbf{D}} f([\mathbf{B}]) + (t^2/2) \text{Hess } f([\mathbf{B}])\{(\mathbf{D}, \mathbf{D})\} \\ &\quad + (1/2) \int_0^t (t-s)^2 (f \circ \gamma)^{(3)}(s) ds . \end{aligned} \quad (2.128)$$

We can bound the last term by Ct^3 for some $C > 0$ since $(f \circ \gamma)^{(3)}$ is bounded on the compact set $[0, 1]$. Our strategy now is to obtain (2.128) in a different way, by computing explicitly the second-order Taylor expansion of $f \circ \gamma(t)$ when $t \rightarrow 0$. Then we identify the first and second order terms to conclude. Recall that using (2.18) and (2.14), for any $([\mathbf{B}], \mathbf{D}) \in \mathbb{T}\text{Gr}_r(\mathbb{R}^d)$, setting $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ the compact singular value decomposition, we have

$$f \circ \gamma(t) = -(1/2) \text{Tr} \left(\mathbf{B}_t^\top \mathbf{A}\mathbf{B}_t \left(\mathbf{B}_t^\top \mathbf{B}_t \right)^{-1} \right), \quad (2.129)$$

where $\mathbf{B}_t = \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \cos(t\mathbf{\Sigma}) + \mathbf{U} \sin(t\mathbf{\Sigma}) \in \text{St}_r(\mathbb{R}^d)$ satisfies $[\mathbf{B}_t] = \gamma(t)$. First, expanding the trigonometric functions applied to the diagonal elements of the matrix $\mathbf{\Sigma}$ brings,

$$\sin(t\mathbf{\Sigma}) \underset{t \rightarrow 0}{=} t\mathbf{\Sigma} + o(t^2), \text{ and } \cos(t\mathbf{\Sigma}) \underset{t \rightarrow 0}{=} \mathbf{I}_r - t^2 \mathbf{\Sigma}^2 / 2 + o(t^2).$$

Thus, the expansion for \mathbf{B}_t is the following,

$$\mathbf{B}_t \underset{t \rightarrow 0}{=} \mathbf{B} \left(\mathbf{B}^\top \mathbf{B} \right)^{-1/2} \mathbf{V} + t \mathbf{U}\mathbf{\Sigma} - t^2 \mathbf{B} \left(\mathbf{B}^\top \mathbf{B} \right)^{-1/2} \mathbf{V} \mathbf{\Sigma}^2 / 2 + o(t^2).$$

For the product $\mathbf{B}^\top \mathbf{B}$, derive

$$\begin{aligned} \mathbf{B}_t^\top \mathbf{B}_t \underset{t \rightarrow 0}{=} & \mathbf{V}^\top \left(\mathbf{B}^\top \mathbf{B} \right)^{-1/2} \mathbf{B}^\top \mathbf{B} \left(\mathbf{B}^\top \mathbf{B} \right)^{-1/2} \mathbf{V} + 2t \operatorname{sym} \left(\mathbf{V}^\top \left(\mathbf{B}^\top \mathbf{B} \right)^{-1/2} \mathbf{B}^\top \mathbf{U} \boldsymbol{\Sigma} \right) \\ & - t^2 \operatorname{sym} \left(\mathbf{V}^\top \left(\mathbf{B}^\top \mathbf{B} \right)^{-1/2} \mathbf{B}^\top \mathbf{B} \left(\mathbf{B}^\top \mathbf{B} \right)^{-1/2} \mathbf{V} \boldsymbol{\Sigma}^2 \right) + t^2 \boldsymbol{\Sigma} \mathbf{U}^\top \mathbf{U} \boldsymbol{\Sigma} + o(t^2), \end{aligned}$$

where $\operatorname{sym}(\mathbf{M}) = (\mathbf{M} + \mathbf{M}^\top)/2$ is the symmetrization of the square matrix \mathbf{M} . Now, we simplify the first term $\mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{B}^\top \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} = \mathbf{I}_r$. In the second term, using $\mathbf{U} \boldsymbol{\Sigma} = \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V}$ and $\mathbf{B}^\top \mathbf{D} = \mathbf{0}_{r \times r}$ by definition of the tangent space $\mathbb{T}_{[\mathbf{B}]} \operatorname{Gr}_r(\mathbb{R}^d)$ (see the paragraph around (2.71) and Example 2.9), the term of order t vanishes. Moreover, the terms of order t^2 are $-\boldsymbol{\Sigma}^2$ and $\boldsymbol{\Sigma}^2$ respectively, bringing

$$\mathbf{B}_t^\top \mathbf{B}_t \underset{t \rightarrow 0}{=} \mathbf{I}_r + o(t^2), \text{ and therefore } \left(\mathbf{B}_t^\top \mathbf{B}_t \right)^{-1} \underset{t \rightarrow 0}{=} \mathbf{I}_r + o(t^2). \quad (2.130)$$

Finally, derive the Taylor expansion of $\mathbf{B}_t \mathbf{B}_t^\top$, using $\mathbf{V} \boldsymbol{\Sigma}^2 \mathbf{V}^\top = (\mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top)^\top \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top = (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{D}^\top \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2}$ and similarly $\mathbf{U} \boldsymbol{\Sigma}^2 \mathbf{U}^\top = \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top$,

$$\begin{aligned} \mathbf{B}_t \mathbf{B}_t^\top \underset{t \rightarrow 0}{=} & \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top + 2t \operatorname{sym} \left(\mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \right) + t^2 \mathbf{U} \boldsymbol{\Sigma}^2 \mathbf{U}^\top \\ & - t^2 \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \boldsymbol{\Sigma}^2 \mathbf{V}^\top (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{B}^\top + o(t^2), \\ \underset{t \rightarrow 0}{=} & \operatorname{Proj}_{[\mathbf{B}]} + 2t \operatorname{sym} \left(\mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \right) + t^2 \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \\ & - t^2 \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top + o(t^2). \end{aligned} \quad (2.131)$$

Now, recall that for any $d \times r$ matrices \mathbf{M}, \mathbf{N} , $\operatorname{Tr}(\mathbf{M}^\top \mathbf{N}) = \operatorname{Tr}(\mathbf{N} \mathbf{M}^\top)$, thus (2.129) can be written using (2.130),

$$\begin{aligned} f \circ \gamma(t) &= -(1/2) \operatorname{Tr} \left(\mathbf{A} \mathbf{B}_t (\mathbf{B}_t^\top \mathbf{B}_t)^{-1} \mathbf{B}_t^\top \right), \\ &\underset{t \rightarrow 0}{=} -(1/2) \operatorname{Tr} \left(\mathbf{A} \mathbf{B}_t \mathbf{I}_r \mathbf{B}_t^\top \right) + o(t^2). \end{aligned}$$

Now recall that for any $d \times d$ matrix \mathbf{M} and any $d \times d$ symmetric matrix \mathbf{N} , $\operatorname{Tr}(\mathbf{N} \mathbf{M}) = \operatorname{Tr}(\mathbf{N} \mathbf{M}^\top)$, therefore $\operatorname{Tr}(\mathbf{N} \operatorname{sym}(\mathbf{M})) = \operatorname{Tr}(\mathbf{N} \mathbf{M})$. Using this and (2.131),

$$\begin{aligned} f \circ \gamma(t) \underset{t \rightarrow 0}{=} & f([\mathbf{B}]) - t \operatorname{Tr} \left(\mathbf{A} \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \right) - (t^2/2) \operatorname{Tr} \left(\mathbf{A} \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \right) \\ & + (t^2/2) \operatorname{Tr} \left(\mathbf{A} \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \right) + o(t^2). \end{aligned}$$

We conclude by identifying the terms as mentioned at the beginning of this proof.

The alternative representation is obtained by checking the following facts. First, $\mathbf{B}^\top \operatorname{Proj}_{[\mathbf{B}_\perp]} \mathbf{A} \mathbf{D} =$

$\mathbf{0}_{r \times r}$, therefore $\text{Proj}_{[\mathbf{B}_\perp]} \mathbf{A} \mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$. Also, by definition of the metric, the fact that $\text{Proj}_{[\mathbf{B}_\perp]}^\top = \text{Proj}_{[\mathbf{B}_\perp]}$ and that $\text{Proj}_{[\mathbf{B}_\perp]} \mathbf{D} = \mathbf{D}$,

$$-\left\langle \text{Proj}_{[\mathbf{B}_\perp]} \mathbf{A} \mathbf{D}, \mathbf{D} \right\rangle_{[\mathbf{B}]} = -\text{Tr} \left(\left(\mathbf{B}^\top \mathbf{B} \right)^{-1} \mathbf{D}^\top \mathbf{A} \mathbf{D} \right).$$

Secondly, $\mathbf{B}^\top \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{A} \mathbf{B} = \mathbf{0}_{r \times r}$, thus $\mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{A} \mathbf{B} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$. Moreover,

$$\text{Tr} \left(\mathbf{A} \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \right) = \left\langle \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{A} \mathbf{B}, \mathbf{D} \right\rangle_{[\mathbf{B}]}.$$

Combining these two last equations into (2.126) gives (2.127). \square

Proposition 2.32. *Consider the setting of Section 2.4.1 and assume **HP 2**. Consider $e_{[\mathbf{B}]}(x) = \{\mathbf{I}_d - \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top\}(xx^\top - \mathbf{A})\mathbf{B}$, for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$, $x \in \mathbb{R}^d$. Then, for any $[\mathbf{B}_0], [\mathbf{B}_1] \in \text{Gr}_r(\mathbb{R}^d)$, $x \in \mathbb{X}$ and geodesic curve $\gamma : [0, 1] \rightarrow \text{Gr}_r(\mathbb{R}^d)$ between $[\mathbf{B}_0]$ and $[\mathbf{B}_1]$,*

$$\left\| e_{[\mathbf{B}_1]}(x) - \mathbf{T}_{01}^\gamma e_{[\mathbf{B}_0]}(x) \right\|_{[\mathbf{B}_1]} \leq C \ell(\gamma) (1 + \|x\|_2^2). \quad (2.132)$$

In particular, for any $x \in \mathbb{R}^d$, $[\mathbf{B}] \mapsto e_{[\mathbf{B}]}(x)$ satisfies (2.11).

Proof. Note that by [EAS98, Section 2.5.3], for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$, $x \in \mathbb{R}^d$, $e_{[\mathbf{B}]}(x) = -\text{grad } f^x([\mathbf{B}]) + \text{grad } f([\mathbf{B}])$ where grad is the Riemannian gradient on $\text{Gr}_r(\mathbb{R}^d)$, $f^x = -\text{Tr}(xx^\top \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top)/2$ and $f([\mathbf{B}]) = -\text{Tr}(\mathbf{A} \mathbf{B}(\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top)/2$. Therefore, to conclude that (2.132) holds, using Lemma 2.38, it is sufficient to show that $\text{Hess } f([\mathbf{B}]), \text{Hess } f^x([\mathbf{B}]) : \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d) \rightarrow \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$ have operator norm upper bounded uniformly for $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ by $C\{1 + \|x\|_2^2\}$ for some constant $C > 0$. In other words, using the Hessian representation given in Lemma 2.31, we want to prove that there exists $C > 0$ such that for any $x \in \mathbb{R}^d$, $([\mathbf{B}], \mathbf{D}) \in \mathbb{T} \text{Gr}_r(\mathbb{R}^d)$, $|\text{Hess } f([\mathbf{B}])\{(\mathbf{D}, \mathbf{D})\}| \leq C\{1 + \|x\|_2^2\} \|\mathbf{D}\|_{[\mathbf{B}]}^2$ and $|\text{Hess } f^x([\mathbf{B}])\{(\mathbf{D}, \mathbf{D})\}| \leq C\{1 + \|x\|_2^2\} \|\mathbf{D}\|_{[\mathbf{B}]}^2$.

First, since f does not depend on x and $\text{Gr}_r(\mathbb{R}^d)$ is compact, there exists $C_1 \geq 0$ such that for any $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$, the operator norm with respect to the Riemannian norm of $\text{Hess } f([\mathbf{B}])$ is bounded by C_1 . This means that for any $([\mathbf{B}], \mathbf{D}) \in \mathbb{T} \text{Gr}_r(\mathbb{R}^d)$, $|\text{Hess } f([\mathbf{B}])\{(\mathbf{D}, \mathbf{D})\}| \leq C_1 \|\mathbf{D}\|_{[\mathbf{B}]}^2$. Secondly, the derivative and the Hessian of f^x are given by Lemma 2.31 by changing \mathbf{A} into xx^\top . Therefore, by taking the expression of the Hessian of f^x given in (2.127) and using the triangle and Cauchy-Schwarz inequalities,

$$\begin{aligned} & |\text{Hess } f^x([\mathbf{B}])\{(\mathbf{D}, \mathbf{D})\}| \\ & \leq \left| \left\langle \text{Proj}_{[\mathbf{B}_\perp]} xx^\top \mathbf{D}, \mathbf{D} \right\rangle_{[\mathbf{B}]} \right| + \left| \left\langle \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top xx^\top \mathbf{B}, \mathbf{D} \right\rangle_{[\mathbf{B}]} \right| \\ & \leq \|\mathbf{D}\|_{[\mathbf{B}]} \left(\left\| \text{Proj}_{[\mathbf{B}_\perp]} xx^\top \mathbf{D} \right\|_{[\mathbf{B}]} + \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top xx^\top \mathbf{B} \right\|_{[\mathbf{B}]} \right). \end{aligned} \quad (2.133)$$

We use in the following computations the relationship between the Frobenius norm $\|\cdot\|_2$ and the Riemannian norm, namely for any $\mathbf{D} \in \mathbb{T}_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, using that $(\mathbf{B}^\top \mathbf{B})^{-1/2}$ is a symmetric matrix,

$$\|\mathbf{D}\|_{[\mathbf{B}]}^2 = \text{Tr} \left((\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{D}^\top \mathbf{D} \right) = \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\|_2^2. \quad (2.134)$$

Using (2.134), the fact that the Frobenius norm is sub-multiplicative, and the definition of the spectral norm given in (2.83),

$$\begin{aligned} \left\| \text{Proj}_{[\mathbf{B}_\perp]} x x^\top \mathbf{D} \right\|_{[\mathbf{B}]} &= \left\| \text{Proj}_{[\mathbf{B}_\perp]} x x^\top \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\|_2 \\ &\leq \left\| \text{Proj}_{[\mathbf{B}_\perp]} x x^\top \right\|_2 \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\|_2 \\ &\leq \left\| \text{Proj}_{[\mathbf{B}_\perp]} \right\| \left\| x x^\top \right\|_2 \|\mathbf{D}\|_{[\mathbf{B}]} \\ &\leq \|x\|_2^2 \|\mathbf{D}\|_{[\mathbf{B}]} , \end{aligned}$$

where we have used that the spectral norm is an operator norm, that the operator norm of a projector (here $\text{Proj}_{[\mathbf{B}_\perp]}$) is equal to one, and that $\|x x^\top\|_2 = \|x\|_2^2$. For the second term in (2.133), we proceed similarly,

$$\begin{aligned} \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top x x^\top \mathbf{B} \right\|_{[\mathbf{B}]} &= \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top x x^\top \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\|_2 \\ &\leq \left\| \mathbf{D} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\|_2 \left\| (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{B}^\top \right\|_2 \left\| x x^\top \right\|_2 \left\| \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\|_2 \\ &\leq \left\| \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\|_2^2 \|x\|_2^2 \|\mathbf{D}\|_{[\mathbf{B}]} , \end{aligned}$$

where we have used that for any $\mathbf{C} \in \mathbb{R}^{d \times r}$, $\|\mathbf{C}\|_2 = \|\mathbf{C}^\top\|_2$. Finally,

$$\left\| \mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \right\|_2^2 = \text{Tr} \left(\mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \right) = \text{Tr} (\mathbf{I}_r) = r .$$

In conclusion, we have that

$$|\text{Hess } f^x([\mathbf{B}])(\mathbf{D}, \mathbf{D})| \leq (r+1) \|x\|_2^2 \|\mathbf{D}\|_{[\mathbf{B}]}^2 ,$$

which completes the proof. \square

2.9.2 Proofs of Section 2.4.2

Recall that we consider the Huber-like dissimilarity measure, $\rho_{\text{H}} : \Theta \times \Theta \rightarrow \mathbb{R}_+$, given for any $\theta_0, \theta_1 \in \Theta$ by

$$\rho_{\text{H}}(\theta_0, \theta_1) = \delta^2 \left[1 + \{\rho_{\Theta}(\theta_0, \theta_1)/\delta\}^2 \right]^{1/2} - \delta^2 , \quad (2.135)$$

where $\delta > 0$.

We present lemmas to verify **A 2–A 3**, **MD 1** for the robust barycenter problem. In addition, we show that the function $f(\theta)$ of (2.21),

$$f(\theta) = \int_{\Theta} \rho_{\mathbb{H}}(\theta, x) \pi(dx) \text{ for } \theta \in \Theta ,$$

is strictly g-convex. Thus the the robust barycenter of a probability distribution π on Θ exists and is unique. Consequently, as the geodesic SA scheme (2.23) finds a stationary point of (2.21), the strict g-convexity of f guarantees that such stationary point is globally optimal and is unique.

Lemma 2.33. *Let Θ be a Hadamard manifold with sectional curvature bounded below by $-\kappa^2$, $x \in \Theta$ and $\delta > 0$. Define for any $\theta \in \Theta$,*

$$V_2(\theta) = \rho_{\Theta}^2(x, \theta) \text{ and } V_1^x(\theta) = \delta^2 \left[V_2(\theta)/\delta^2 + 1 \right]^{1/2} - \delta^2 .$$

Then, V_1^x is a smooth function and its Riemannian gradient is given for any $\theta \in \Theta$ by

$$\text{grad } V_1^x(\theta) = - \text{Exp}_{\theta}^{-1}(x) \Big/ \left[V_2(\theta)/\delta^2 + 1 \right]^{1/2} . \quad (2.136)$$

Moreover, for any $\theta \in \Theta$, $v \in \mathbb{T}_{\theta}\Theta \setminus \{0\}$, its Hessian satisfies

$$0 < \text{Hess } V_1^x(\theta)(v, v) \leq (1 + \delta\kappa) \|v\|_{\theta}^2 . \quad (2.137)$$

Proof. The proof relies heavily on the computation of $\text{grad } V_2^x$ and the Hessian comparison of V_2^x done in [Jos05, Theorem 5.6.1]. Indeed, [Jos05, Theorem 5.6.1] shows V_2 is smooth and that for any $\theta \in \Theta$,

$$\text{grad } V_2(\theta) = -2\text{Exp}_{\theta}^{-1}(x) .$$

Hence, (2.136) follows by composition since $V_1(\theta) = \varpi \circ V_2(\theta)$ where $\varpi : t \rightarrow \delta^2[t/\delta^2 + 1]^{1/2} - \delta^2$ for $t \in \mathbb{R}_+$.

For any $v \in \mathbb{T}_v\Theta$, recall that $\text{Hess } V_1(\theta)(v, v) = \langle \nabla_v \text{grad } V_1(\theta), v \rangle_{\theta}$, where ∇ is the Levi-Civita connection – see Section 2.11.8. The product rule for the covariant derivative [GHL04, p. 73] for a smooth function f and vector field Y on Θ gives $\nabla(fY) = \nabla f \otimes Y + f\nabla Y$. Applying this result to

$$f(\theta) = \left[V_2(\theta)/\delta^2 + 1 \right]^{-1/2} , \quad Y(\theta) = -\text{Exp}_{\theta}^{-1}(x) ,$$

and using $Y = \text{grad } V_2/2$ gives

$$\begin{aligned} \text{Hess } V_1(\theta) = & - \text{Exp}_{\theta}^{-1}(x) \otimes \text{Exp}_{\theta}^{-1}(x) \Big/ \left[\delta^2 \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{3/2} \right] \\ & + \text{Hess } V_2(\theta) \Big/ \left[2 \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{1/2} \right] . \end{aligned} \quad (2.138)$$

Let $\theta \in \Theta$ and $v \in \mathbb{T}_\theta \Theta \setminus \{0\}$. On the one hand, we have $\text{Exp}_\theta^{-1}(x) \otimes \text{Exp}_\theta^{-1}(x)(v, v) = \langle \text{Exp}_\theta^{-1}(x), v \rangle_\theta^2$. Therefore, using Cauchy-Schwarz inequality,

$$0 \leq \text{Exp}_\theta^{-1}(x) \otimes \text{Exp}_\theta^{-1}(x)(v, v) \leq \left\| \text{Exp}_\theta^{-1}(x) \right\|_\theta^2 \|v\|_\theta^2 = \rho_\Theta^2(\theta, x) \|v\|_\theta^2, \quad (2.139)$$

since $\|\text{Exp}_\theta^{-1}(x)\|_\theta = \rho_\Theta(\theta, x)$. On the other hand, [Jos05, Theorem 5.6.1] implies that

$$2 \|v\|_\theta^2 \leq \text{Hess } V_2(\theta)(v, v) \leq 2\kappa\rho_\Theta(\theta, x) \coth[\kappa\rho_\Theta(\theta, x)] \|v\|_\theta^2. \quad (2.140)$$

Now, combining (2.139)-(2.140) in (2.138), and using $t \coth(t) \leq 1 + t$ for $t \geq 0$, it follows that

$$\begin{aligned} \text{Hess } V_1(\theta)(v, v) &\leq [1 + \kappa\rho_\Theta(\theta, x)] / \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{1/2} \|v\|_\theta^2 \\ &\leq (1 + \kappa\delta) \|v\|_\theta^2, \end{aligned}$$

where we have used $1 + V_2(\theta)/\delta^2 \geq \max(1, \rho_\Theta^2(\theta, x)/\delta^2)$. Similarly, we obtain

$$\begin{aligned} \text{Hess } V_1(\theta)(v, v) &\geq - \left(\rho_\Theta^2(\theta, x) \|v\|_\theta^2 \right) / \left(\delta^2 \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{3/2} \right) + \|v\|_\theta^2 / \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{1/2} > 0, \end{aligned}$$

which concludes the proof. \square

Proof of Proposition 2.15. To show that **A 2** holds with the stated values of \underline{c} , \bar{c} and L , note that the scheme (2.23) can be written

$$\theta_{n+1} = \text{Exp}_{\theta_n} \left(-\eta_{n+1} \text{grad } V_1^{X_{n+1}}(\theta_n) \right),$$

where V_1^x and its gradient are given in Lemma 2.33. Then, both $\text{grad } f$ and $\text{Hess } f$ can be computed differentiating under the integral in (2.21). Using Lemma 2.33, we know that for any $x \in \Theta$, $V_1 : \theta \mapsto \delta^2(1 + \rho_\Theta^2(\theta, x)/\delta^2)^{1/2} - \delta^2$ is smooth. Using (2.136) and $\|\text{Exp}_\theta^{-1}(x)\|_\theta = \rho_\Theta(\theta, x)$, we have that for any $x, \theta \in \Theta$, $\|\text{grad } V_1^x(\theta)\|_\theta \leq \delta$ and, for any $x, \theta \in \Theta, v \in \mathbb{T}_\theta \Theta$,

$$|\text{Hess } V_1^x(\theta)(v, v)| \leq (1 + \kappa\delta) \|v\|_\theta^2.$$

Using Lebesgue's dominated convergence theorem, we have for any $\theta \in \Theta$,

$$\text{grad } f(\theta) = - \int_\Theta \text{Exp}_\theta^{-1}(x) / \left[1 + \{\rho_\Theta(\theta, x)/\delta\}^2 \right]^{1/2} \pi(\text{d}x), \quad (2.141)$$

as well as for any $v \in \mathbb{T}_\theta \Theta$,

$$\text{Hess } f(\theta)(v, v) = \int_\Theta \text{Hess } V_1^x(\theta)(v, v) \pi(\text{d}x).$$

Thus, using (2.137) for any $v \in \mathbb{T}_\theta \Theta \setminus \{0\}$,

$$0 < \text{Hess } f(\theta)(v, v) \leq (1 + \kappa\delta) \|v\|_\theta^2. \quad (2.142)$$

This last inequality proves that the operator norm of $\text{Hess } f$ is upper bounded by $1 + \delta\kappa$. Therefore, using Lemma 2.38, it follows that **A2-(b)** holds with $L = 1 + \delta\kappa$.

Combining that for any $x, \theta \in \Theta$, $e_\theta(x) = \text{grad } V_1^x(\theta) - \text{grad } f(\theta)$, $\|\text{grad } V_1^x(\theta)\|_\theta \leq \delta$ and (2.141), we have that for any $x \in \Theta$, $\sup_{\theta \in \Theta} \|e_\theta(x)\|_\theta \leq 2\delta$, which means that the first point in (2.11), inside **A5**, holds for $w = \mathbf{1}$. For the second point of (2.11), inside **A5**, we use that we have just proven that $\text{grad } f$ and $\text{grad } V_1^x$, by the same arguments, are $(1 + \delta\kappa)$ -Lipschitz functions, for any $x \in \Theta$. Therefore for any $x, \theta_0, \theta_1 \in \Theta$ and geodesic $\gamma : [0, 1] \rightarrow \Theta$ between θ_0 and θ_1 ,

$$\begin{aligned} \|e_{\theta_1}(x) - T_{01}^\gamma e_{\theta_0}(x)\|_{\theta_1} &\leq \|\text{grad } V_1^x(\theta_1) - T_{01}^\gamma \text{grad } V_1^x(\theta_0)\|_{\theta_1} \\ &\quad + \|\text{grad } f(\theta_1) - T_{01}^\gamma \text{grad } f(\theta_0)\|_{\theta_1} \\ &\leq 2(1 + \delta\kappa)\ell(\gamma), \end{aligned}$$

concluding that **A5(1)** holds.

It remains to prove that **MD1** is satisfied with $\sigma_0 = \delta^2$ and $\sigma_1 = 0$ if $(X_n)_{n \in \mathbb{N}^*}$ is an i.i.d. family with distribution π . Using the previous calculations, we have $\|H_{\theta_n}(X_{n+1})\|_{\theta_n}^2 = \|V_1^{X_{n+1}}(\theta_n)\|_{\theta_n} \leq \delta^2$. Thus, in the notation of (2.8),

$$\mathbb{E} \left[\|e_{\theta_n}(X_{n+1})\|_{\theta_n}^2 \middle| \mathcal{F}_n \right] = \mathbb{E} \left[\|H_{\theta_n}(X_{n+1}) - h(\theta_n)\|_{\theta_n}^2 \middle| \mathcal{F}_n \right] \leq \delta^2, \quad (2.143)$$

since the conditional variance is bounded by the mean square. Now, the required values of σ_0 and σ_1 can be read from (2.143). \square

Proof of Proposition 2.14. Assume there exists $\tau \in \Theta$ such that (2.22) holds. We prove that the function f takes finite values for any $\theta \in \Theta$, i.e. $f(\theta) < +\infty$. Indeed, note the following inequality, which holds for all $x \geq 0$,

$$\delta^2 \left[1 + \{x/\delta\}^2 \right]^{1/2} - \delta^2 \leq \delta x.$$

From (2.135) and (2.21), this inequality implies for any $\theta \in \Theta$,

$$f(\theta) \leq \delta \int_{\Theta} \rho_{\Theta}(\theta, x) \pi(dx),$$

and, furthermore, by the triangle inequality,

$$f(\theta) \leq \delta \int_{\Theta} (\rho_{\Theta}(\theta, \tau) + \rho_{\Theta}(\tau, x)) \pi(dx) = \delta \rho_{\Theta}(\theta, \tau) + \delta \int_{\Theta} \rho_{\Theta}(\tau, x) \pi(dx). \quad (2.144)$$

Then, it follows from (2.22) and (2.144) that $f(\theta) < +\infty$ for $\theta \in \Theta$.

Further, from (2.142) in the proof of Lemma 2.15, $\text{Hess } f(\theta) \succ 0$, so f has strictly positive-definite Riemannian Hessian and, using Theorem 2.37 and $\nabla^2 f = \text{Hess } f$, is therefore geodesically strictly convex. It remains to show that f has a unique global minimizer θ^* . Since f is geodesically strictly convex, continuous, and bounded from below (indeed,

f is positive), it is enough to show that f is coercive, in the sense that $f(\theta) \rightarrow +\infty$ when $\rho_{\Theta}(\theta, \tau) \rightarrow +\infty$. To do so, note the following inequality holds for all real x ,

$$\delta^2 \left[1 + \{x/\delta\}^2 \right]^{1/2} - \delta^2 \geq 2^{-1/2}\delta(x - \delta).$$

Then, using (2.135) and (2.21), this inequality implies

$$f(\theta) \geq 2^{-1/2}\delta \int_{\Theta} (\rho_{\Theta}(\theta, x) - \delta) \pi(dx),$$

or, by the triangle inequality,

$$f(\theta) \geq 2^{-1/2}\delta \int_{\Theta} (|\rho_{\Theta}(\theta, \tau) - \rho_{\Theta}(\tau, x)| - \delta) \pi(dx) \quad \text{for } \theta \in \Theta.$$

However, this directly yields

$$f(\theta) \geq 2^{-1/2}\delta \left(\rho_{\Theta}(\theta, \tau) - \int_{\Theta} (\rho_{\Theta}(\tau, x) + \delta) \pi(dx) \right).$$

Clearly, the right-hand side increases to $+\infty$ when $\rho_{\Theta}(\theta, \tau) \rightarrow +\infty$. Thus, f is indeed coercive.

Finally, to show that f has a unique global minimizer θ^* , let $f_* = \inf\{f(\theta) : \theta \in \Theta\}$ and $(\theta_n)_{n \in \mathbb{N}}$ a sequence of points in Θ such that $\lim f(\theta_n) = f_*$. Since f takes finite values, $f_* < +\infty$. Therefore, $(f(\theta_n))_{n \in \mathbb{N}}$ is bounded, and by coercivity of f , so is $(\rho_{\Theta}(\theta_n, \tau))_{n \in \mathbb{N}}$. Therefore, there exists some $R > 0$ such that $\rho_{\Theta}(\theta_n, \tau) < R$ for all $n \in \mathbb{N}$. Because Θ is a complete Riemannian manifold, the metric ball $B(\tau, R)$ has compact closure (a consequence of the Hopf-Rinow theorem [Lee19]). This implies that the sequence θ_n has a convergent subsequence, whose limit θ^* belongs to the closure of $B(\tau, R)$. By continuity of f , it is clear that $f(\theta^*) = \lim f(\theta_n) = f_*$, so θ^* is indeed a global minimizer of f . This global minimizer is unique because f is geodesically strictly convex. \square

2.9.3 Details and specification for Section 2.4.3

The entries of the matrix \mathbf{M} are chosen such that

- (a) for any $i \in \{1, \dots, 50\}$, $\mathbf{M}_{i,i}$ is a sample from the uniform distribution on $[0.1, 0.7]$;
- (b) for any $k \in \{1, \dots, 23\}$ and $i \in \{1, d - k\}$, $\mathbf{M}_{i,k} = U_{i,k}(0.1 + 0.8(i - 1)/27)$, where $U_{i,k}$ is a sample from the uniform distribution on $[0.05, 0.1]$;
- (c) $\|\mathbf{M}\| < 1$.

By a straightforward induction, we have since for any $k \in \mathbb{N}$, $X_{k+1} = \mathbf{M}X_k + \sigma\varepsilon_{k+1}$, where $\sigma > 0$,

$$X_k = \mathbf{M}^k X_0 + \sigma \sum_{i=0}^{k-1} \mathbf{M}^{k-1-i} \varepsilon_{i+1}.$$

Therefore, for any $k \in \mathbb{N}^*$, X_k follows a Gaussian distribution with mean $\mathbf{M}^k X_0$ and covariance matrix $\sigma^2 \sum_{i=0}^{k-1} \mathbf{M}^{k-1-i} [\mathbf{M}^{k-1-i}]^\top$. Denote by P the Markov kernel associated with $(X_k)_{k \in \mathbb{N}}$. Then, it is straightforward to verify that it is strongly aperiodic and irreducible with the Lebesgue measure as an irreducibility measure. In addition, the zero-mean Gaussian distribution with covariance matrix $\Sigma_\pi = \sigma^2 \sum_{i=0}^{\infty} \mathbf{M}^i [\mathbf{M}^i]^\top$ is the unique stationary distribution of P . Note that Σ_π is well defined since $\|\mathbf{M}\| < 1$. In addition, using that $\|\mathbf{M}\| < 1$ again, an easy computation shows that there exists $c > 0$ small enough such that setting $w(x) = \exp(c\|x\|_2^2)$, there exists $\lambda \in (0, 1)$ and $b_w \in \mathbb{R}_+$ such that $Pw(x) \leq \lambda w(x) + b_w$ for any $x \in \mathbb{R}^d$. Therefore **HP 2** holds for this example. In addition, Σ_π satisfies the discrete Riccati equation: $\Sigma_\pi = \mathbf{M}\Sigma_\pi\mathbf{M}^\top + \sigma^2\mathbf{I}_{50}$ which is solved using the function `dare` of the Python library `control`².

2.10 Useful results from matrix analysis

In this section, we introduce two linear algebra or matrix analysis auxiliary results. They are related to the nonuniqueness of the singular value decomposition.

Lemma 2.34. *Let $\mathbf{N} \in \mathbb{R}^{d \times r}$ be a rectangular real matrix. Consider two singular value decompositions of \mathbf{N} , i.e. let $\mathbf{U}, \tilde{\mathbf{U}} \in \mathbb{R}^{d \times r}$ such that $\mathbf{U}^\top \mathbf{U} = \tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} = \mathbf{I}_r$, let $\Sigma, \tilde{\Sigma} \in \mathbb{R}^{r \times r}$ be two nonnegative diagonal matrices and let $\mathbf{V}, \tilde{\mathbf{V}} \in \text{O}_r(\mathbb{R})$ be two orthogonal matrices such that $\mathbf{N} = \mathbf{U}\Sigma\mathbf{V}^\top = \tilde{\mathbf{U}}\tilde{\Sigma}\tilde{\mathbf{V}}^\top$.*

Denote $(u_i)_{i \in \{1, \dots, r\}}, (\tilde{u}_i)_{i \in \{1, \dots, r\}}, (v_i)_{i \in \{1, \dots, r\}}$ and $(\tilde{v}_i)_{i \in \{1, \dots, r\}}$ the columns of $\mathbf{U}, \tilde{\mathbf{U}}, \mathbf{V}$ and $\tilde{\mathbf{V}}$ respectively, and denote $(\lambda_i)_{i \in \{1, \dots, r\}}$ and $(\tilde{\lambda}_i)_{i \in \{1, \dots, r\}}$ the diagonal coefficients of Σ and $\tilde{\Sigma}$ respectively. Then, for any $i, j \in \{1, \dots, r\}$ such that $\lambda_i \neq \tilde{\lambda}_j$, we have $v_i^\top \tilde{v}_j = 0$ and $u_i^\top \tilde{u}_j = 0$.

Proof. Consider $\text{Sing}_{\mathbf{N}} = \{\lambda_i : i \in \{1, \dots, r\}\}$, by uniqueness of the singular values of \mathbf{N} , we have $\text{Sing}_{\mathbf{N}} = \{\tilde{\lambda}_i : i \in \{1, \dots, r\}\}$. Since \mathbf{V} and $\tilde{\mathbf{V}}$ are orthogonal matrices, their column vectors form orthonormal families. We have, by definition of the singular value decomposition, $\mathbf{N}^\top \mathbf{N} \mathbf{V} = \mathbf{V} \Sigma^2$. For any $i \in \{1, \dots, r\}$, we can reformulate this column-wise, meaning

$$\mathbf{N}^\top \mathbf{N} v_i = \lambda_i^2 v_i. \quad (2.145)$$

Let $\lambda \in \text{Sing}_{\mathbf{N}}$ and consider $I_\lambda = \{i \in \{1, \dots, r\} : \lambda_i = \lambda\}$. One of the consequences of the spectral theorem for the matrix $\mathbf{N}^\top \mathbf{N}$ is that

$$\mathbf{E}_\lambda = \text{Span}((v_i)_{i \in I_\lambda}). \quad (2.146)$$

where $\mathbf{E}_\lambda = \text{Ker}(\mathbf{N}^\top \mathbf{N} - \lambda^2 \mathbf{I}_r)$. However, (2.145) holds also for $(\tilde{v}_i)_{i \in \{1, \dots, r\}}$ since $\mathbf{N}^\top \mathbf{N} \tilde{\mathbf{V}} = \tilde{\mathbf{V}} \tilde{\Sigma}^2$. Therefore, by (2.146), for any $i \in \{1, \dots, r\}$,

$$\tilde{v}_i \in \text{Span}(\{v_j : j \in I_{\tilde{\lambda}_i}\}).$$

²<https://python-control.readthedocs.io/en/0.8.3/intro.html>

Another consequence of the spectral theorem for $\mathbf{N}^\top \mathbf{N}$ is that for $\lambda_1, \lambda_2 \in \text{Sing}_{\mathbf{N}}$ such that $\lambda_1 \neq \lambda_2$, we have that \mathbf{E}_{λ_1} and \mathbf{E}_{λ_2} are orthogonal. Using this and (2.146), we have that for any $i, j \in \{1, \dots, r\}$, if $\lambda_i \neq \tilde{\lambda}_j$ then $v_i^\top \tilde{v}_j = 0$.

To obtain the same conclusion with the vectors $(u_i)_{i \in \{1, \dots, r\}}$ and $(\tilde{u}_i)_{i \in \{1, \dots, r\}}$, first we compute,

$$\mathbf{N}\mathbf{N}^\top \mathbf{U} = \mathbf{U}\Sigma^2,$$

which is the parallel to (2.145), meaning for any $i \in \{1, \dots, r\}$, u_i is an eigenvector of $\mathbf{N}\mathbf{N}^\top$. Now, by the spectral theorem applied to $\mathbf{N}\mathbf{N}^\top$, denoting \mathbf{S} the eigenvalues of $\mathbf{N}\mathbf{N}^\top$, we have

$$\mathbb{R}^d = \bigoplus_{\lambda \in \mathbf{S}} \mathbf{E}'_{\lambda}, \quad (2.147)$$

where $\mathbf{E}'_{\lambda} = \ker(\mathbf{N}\mathbf{N}^\top - \lambda \mathbf{I}_d)$, and the components of the sum are mutually orthogonal. For any $\lambda \in \mathbf{S}$, we can complete the family $\{u_i : i \in \{1, \dots, r\}, u_i \in \mathbf{E}'_{\lambda}\}$ into an orthonormal basis of \mathbf{E}'_{λ} . By combining these families, there exist some vectors $(u_i)_{i \in \{r+1, \dots, d\}}$ in \mathbb{R}^d that complete $(u_i)_{i \in \{1, \dots, r\}}$ into an orthonormal basis of \mathbb{R}^d that respects the decomposition (2.147). This means that $(u_i)_{i \in \{1, \dots, d\}}$ is an orthonormal basis of eigenvectors of $\mathbf{N}\mathbf{N}^\top$. Therefore, for any $i \in \{r+1, \dots, d\}$, there exists $\lambda \in \mathbf{S}$ such that $\mathbf{N}\mathbf{N}^\top u_i = \lambda u_i$. However $\text{Span}((\lambda_j u_j)_{j \in \{1, \dots, r\}}) \subset \text{Im}(\mathbf{N}\mathbf{N}^\top)$ and, since $\text{Rank}(\mathbf{N}) = \text{Rank}(\mathbf{N}\mathbf{N}^\top)$ which coincides with the cardinality of $\{j \in \{1, \dots, r\} : \lambda_j \neq 0\}$, $\text{Span}((\lambda_j u_j)_{j \in \{1, \dots, r\}}) = \text{Im}(\mathbf{N}\mathbf{N}^\top)$ because the dimensions match. It follows that $\lambda = 0$, because $\mathbf{N}\mathbf{N}^\top u_i = \lambda u_i \in \text{Im}(\mathbf{N}\mathbf{N}^\top)$ and u_i is orthogonal to $\text{Im}(\mathbf{N}\mathbf{N}^\top)$ since it is orthogonal to the family $(u_j)_{j \in \{1, \dots, r\}}$. The rest of the proof follows from the same steps as the ones used for \mathbf{V} and $\tilde{\mathbf{V}}$. \square

Lemma 2.35. *Let $\mathbf{N} \in \mathbb{R}^{d \times r}$ be a rectangular real matrix. Consider $\mathbf{U}, \tilde{\mathbf{U}} \in \mathbb{R}^{d \times r}$ such that $\mathbf{U}^\top \mathbf{U} = \tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} = \mathbf{I}_r$, let $\Sigma, \tilde{\Sigma} \in \mathbb{R}^{r \times r}$ be nonnegative diagonal matrices and $\mathbf{V}, \tilde{\mathbf{V}} \in \text{O}_r(\mathbb{R})$ such that we have the two decompositions $\mathbf{N} = \mathbf{U}\Sigma\mathbf{V}^\top = \tilde{\mathbf{U}}\tilde{\Sigma}\tilde{\mathbf{V}}^\top$. For any $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ such that $f(0) = 0$, we have $\mathbf{U}f(\Sigma)\mathbf{V}^\top = \tilde{\mathbf{U}}f(\tilde{\Sigma})\tilde{\mathbf{V}}^\top$, where we apply f to a diagonal matrix by applying it to each diagonal coefficient individually. In particular, the quantity $\mathbf{U}f(\Sigma)\mathbf{V}^\top$ only depends on \mathbf{N} and f , therefore we can define the mapping $\psi : \mathbf{N} \mapsto \mathbf{U}f(\Sigma)\mathbf{V}^\top$ on $\mathbb{R}^{d \times r}$.*

Proof. We divide the proof into two cases. First, assume Σ and $\tilde{\Sigma}$ both have nonincreasing diagonal coefficients. By uniqueness of the singular values of the matrix \mathbf{N} and the assumption of nonincreasingness, we have $\Sigma = \tilde{\Sigma}$. Therefore $f(\Sigma) = f(\tilde{\Sigma})$. Denote $(v_i)_{i \in \{1, \dots, r\}}, (\tilde{v}_i)_{i \in \{1, \dots, r\}} \in (\mathbb{R}^r)^r$ the column vectors of \mathbf{V} and $\tilde{\mathbf{V}}$ respectively. Denote as well $(\lambda_i)_{i \in \{1, \dots, r\}}$ the diagonal coefficients of Σ . Define $\mathbf{N}_f = \mathbf{U}f(\Sigma)\mathbf{V}^\top$ and $\tilde{\mathbf{N}}_f = \tilde{\mathbf{U}}f(\Sigma)\tilde{\mathbf{V}}^\top$. We check now that $\mathbf{N}_f = \tilde{\mathbf{N}}_f$ by proving that, as linear maps from \mathbb{R}^r to \mathbb{R}^d , they coincide on the basis $(\tilde{v}_i)_{i \in \{1, \dots, r\}}$. Note that these definitions give the equations $\mathbf{N}\mathbf{V} = \mathbf{U}\Sigma$, $\mathbf{N}_f\mathbf{V} = \mathbf{U}f(\Sigma)$ and $\tilde{\mathbf{N}}_f\tilde{\mathbf{V}} = \tilde{\mathbf{U}}f(\Sigma)$, for which we use the column-wise interpretations. Let $i \in \{1, \dots, r\}$, using that $(v_j)_{j \in \{1, \dots, r\}}$ is an orthonormal basis of \mathbb{R}^r ,

and $\mathbf{N}_f \mathbf{V} = \mathbf{U}f(\boldsymbol{\Sigma})$,

$$\mathbf{N}_f \tilde{v}_i = \sum_{j=1}^r (\tilde{v}_i^\top v_j) \mathbf{N}_f v_j = \sum_{j=1}^r (\tilde{v}_i^\top v_j) f(\lambda_j) u_j .$$

Now, using that $f(0) = 0$, we have for any $x \in \mathbb{R}$, $f(x) = f(x) \mathbb{1}_{x \neq 0}$. Also, using that $\mathbf{N}\mathbf{V} = \mathbf{U}\boldsymbol{\Sigma}$, if $\lambda_j \neq 0$, $u_j = \mathbf{N}v_j/\lambda_j$. This brings

$$\mathbf{N}_f \tilde{v}_i = \sum_{j=1}^r (\tilde{v}_i^\top v_j) f(\lambda_j) \mathbb{1}_{\lambda_j \neq 0} u_j = \sum_{j=1}^r (\tilde{v}_i^\top v_j) (f(\lambda_j)/\lambda_j) \mathbb{1}_{\lambda_j \neq 0} \mathbf{N}v_j .$$

Now, we use that $\mathbf{N}\tilde{\mathbf{V}} = \tilde{\mathbf{U}}\boldsymbol{\Sigma}$ and Lemma 2.34, which gives us that $\tilde{v}_i^\top v_j = 0$ if $\lambda_j \neq \lambda_i$, and we write

$$\begin{aligned} \mathbf{N}_f \tilde{v}_i &= (f(\lambda_i)/\lambda_i) \mathbb{1}_{\lambda_i \neq 0} \sum_{j=1}^r (\tilde{v}_i^\top v_j) \mathbf{N}v_j = (f(\lambda_i)/\lambda_i) \mathbb{1}_{\lambda_i \neq 0} \mathbf{N} \left(\sum_{j=1}^r (\tilde{v}_i^\top v_j) v_j \right) \\ &= (f(\lambda_i)/\lambda_i) \mathbb{1}_{\lambda_i \neq 0} \mathbf{N} \tilde{v}_i = (f(\lambda_i)/\lambda_i) \mathbb{1}_{\lambda_i \neq 0} \lambda_i \tilde{u}_i = f(\lambda_i) \tilde{u}_i . \end{aligned}$$

We have proven $\mathbf{N}_f \tilde{\mathbf{V}} = \tilde{\mathbf{U}}f(\boldsymbol{\Sigma})$, which means that $\mathbf{N}_f = \tilde{\mathbf{N}}_f$ which concludes the proof of the first case.

Now, consider the general case. First, we prove that without loss of generality, we can assume $\mathbf{N} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ to be a singular value decomposition such that $\boldsymbol{\Sigma}$ has non-increasing diagonal coefficients. Indeed, one can always consider a new decomposition $\mathbf{N} = \mathbf{U}'\boldsymbol{\Sigma}'(\mathbf{V}')^\top$ such that $\boldsymbol{\Sigma}'$ has nonincreasing diagonal coefficients. Then, once we have proven $\mathbf{U}'f(\boldsymbol{\Sigma}')(\mathbf{V}')^\top = \mathbf{U}f(\boldsymbol{\Sigma})\mathbf{V}^\top$, we will have also proven $\mathbf{U}'f(\boldsymbol{\Sigma}')(\mathbf{V}')^\top = \tilde{\mathbf{U}}f(\tilde{\boldsymbol{\Sigma}})\tilde{\mathbf{V}}^\top$, and therefore $\tilde{\mathbf{U}}f(\boldsymbol{\Sigma})\tilde{\mathbf{V}}^\top = \mathbf{U}f(\boldsymbol{\Sigma})\mathbf{V}^\top$. Secondly, by uniqueness of the singular values of the matrix \mathbf{N} , there exists a permutation matrix $\mathbf{M} \in \mathbb{R}^{r \times r}$ such that $\mathbf{M}\tilde{\boldsymbol{\Sigma}}\mathbf{M}^\top = \boldsymbol{\Sigma}$. Recall that a permutation matrix is orthogonal, therefore $\mathbf{M}^\top\mathbf{M} = \mathbf{M}\mathbf{M}^\top = \mathbf{I}_r$. This brings,

$$\mathbf{N} = \tilde{\mathbf{U}} \left(\mathbf{M}^\top \mathbf{M} \right) \tilde{\boldsymbol{\Sigma}} \left(\mathbf{M}^\top \mathbf{M} \right) \tilde{\mathbf{V}}^\top = \left(\tilde{\mathbf{U}} \mathbf{M}^\top \right) \boldsymbol{\Sigma} \left(\tilde{\mathbf{V}} \mathbf{M}^\top \right)^\top . \quad (2.148)$$

We can now apply the first case in this proof to this decomposition and the one $\mathbf{N} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$. Therefore $(\tilde{\mathbf{U}}\mathbf{M}^\top)f(\boldsymbol{\Sigma})(\tilde{\mathbf{V}}\mathbf{M}^\top)^\top = \mathbf{U}f(\boldsymbol{\Sigma})\mathbf{V}^\top$. However, since f is applied on $\boldsymbol{\Sigma}$ coefficient-wise on its diagonal, $\mathbf{M}f(\tilde{\boldsymbol{\Sigma}})\mathbf{M}^\top = f(\boldsymbol{\Sigma})$. Performing the same computation as in (2.148) but with $f(\boldsymbol{\Sigma})$ brings $\mathbf{U}f(\boldsymbol{\Sigma})\mathbf{V}^\top = \tilde{\mathbf{U}}f(\tilde{\boldsymbol{\Sigma}})\tilde{\mathbf{V}}^\top$. \square

2.11 Preliminaries on Riemannian Geometry

This section reviews some basic concepts about Riemannian geometry. These concepts are essential to develop our main results on convergence of Riemannian SA.

2.11.1 Metric Tensor and Distance

A smooth manifold (at least C^2) Θ is equipped with a smooth symmetric and positive-definite metric tensor field $g \in \Gamma(T^2T^*\Theta)$, see [Lee19, Proposition 2.4]. To each $\theta \in \Theta$, this associates a scalar product g_θ on the tangent space $T_\theta\Theta$. When there is no confusion, we denote that [Lee19, Chapter 2, pages 11-12],

$$g_\theta(u, v) = \langle u, v \rangle_\theta, \quad u, v \in T_\theta\Theta$$

and the corresponding norm on $T_\theta\Theta$ is called the Riemannian norm, $\|u\|_\theta = \langle u, u \rangle_\theta^{1/2}$.

With the metric tensor, it is possible to define the notion of length of a curve. If $c : I \rightarrow \Theta$ is a differentiable curve, defined on some interval $I \subset \mathbb{R}$, with velocity \dot{c} , then its length is [Lee19, page 34]

$$\ell(c) = \int_I \|\dot{c}(t)\|_{c(t)} dt. \quad (2.149)$$

The length $\ell(c)$ is invariant by reparametrization: $\ell(c \circ \phi) = \ell(c)$ for any diffeomorphism $\phi : J \rightarrow I$, from an interval J onto I . Thus, without loss of generality, we consider only curves that are restricted to $c : [0, 1] \rightarrow \Theta$.

This can be used to turn Θ into a metric space. Indeed, if Θ is connected, the following is a well-defined distance function; satisfying the axioms of a metric space [Lee19, Theorem 2.55],

$$\rho_\Theta(\theta, \theta') = \inf \{ \ell(c) \mid c : [0, 1] \rightarrow \Theta; c(0) = \theta, c(1) = \theta' \}, \quad \theta, \theta' \in \Theta. \quad (2.150)$$

This is called the Riemannian distance induced by the metric tensor g .

The infimum in (2.150) is always attained for any $\theta, \theta' \in \Theta$, provided that the distance $\rho_\Theta(\cdot, \cdot)$ turns Θ into a complete metric space. This is a corollary of the Hopf-Rinow theorem, a fundamental theorem in Riemannian geometry [Lee19, Corollary 6.21].

2.11.2 Levi-Civita Connection

In the Riemannian manifold, the curve γ which attains the infimum in (2.150) is called a geodesic. Intuitively, a geodesic is a C^2 curve which has zero acceleration. This intuition can be formalized by introducing an affine connection ∇ [Lee19, page 89], compatible with the metric tensor g , called the Levi-Civita connection, or just Riemannian connection (to be precise, $\nabla = \nabla^g$, depends on the choice of g).

To each vector $u \in T_\theta\Theta$ and smooth vector field \mathfrak{X} on Θ , the connection ∇ associates a vector $\nabla_u \mathfrak{X}(\theta) \in T_\theta\Theta$. This vector is called the covariant derivative at θ of \mathfrak{X} in the direction of u . This is bilinear in u and \mathfrak{X} , and satisfies the product rule

$$\nabla_u(f\mathfrak{X}) = (uf)\mathfrak{X}(\theta) + f(\theta)\nabla_u\mathfrak{X}, \quad (2.151)$$

for any differentiable function $f : \Theta \rightarrow \mathbb{R}$, where uf denotes the derivative of f along u . Moreover, one has the following,

$$u \langle \mathfrak{X}, \mathfrak{Y} \rangle = \langle \nabla_u \mathfrak{X}, \mathfrak{Y} \rangle_\theta + \langle \mathfrak{X}, \nabla_u \mathfrak{Y} \rangle_\theta, \quad (2.152)$$

$$\nabla_{\mathfrak{X}}\mathfrak{Y} - \nabla_{\mathfrak{Y}}\mathfrak{X} = [\mathfrak{X}, \mathfrak{Y}] , \quad (2.153)$$

for any differentiable vector fields $\mathfrak{X}, \mathfrak{Y}$ on Θ , where $[\mathfrak{X}, \mathfrak{Y}]$ is the Lie bracket of the vector fields \mathfrak{X} and \mathfrak{Y} , itself a vector field. Here, (2.152) states that ∇ is compatible with the metric, and (2.153) states that ∇ is a connection with zero torsion.

The Levi-Civita connection [Lee19, Theorem 5.10] is defined as the unique affine connection ∇ which satisfies (2.152) and (2.153). Note that the uniqueness of this connection can be guaranteed by the Koszul's theorem, also known as the fundamental theorem of Riemannian geometry.

2.11.3 Geodesic Equation

It can be proved using (2.151) that $\nabla_u \mathfrak{X}$ depends only on the values of \mathfrak{X} along a curve tangent to the vector u [Lee19, Proposition 4.26]. This motivates the following definition. Consider $c : I \rightarrow \Theta$ as a smooth curve on Θ and \mathfrak{X} is an extendible vector field along c , this means that $\mathfrak{X} : I \rightarrow T\Theta$ satisfies $\mathfrak{X}(t) \in T_{c(t)}\Theta$ for any $t \in I$, see [Lee19, pages 100-101]. The covariant derivative of \mathfrak{X} along c is defined by

$$D_t \mathfrak{X} = \nabla_{\tilde{c}} \tilde{\mathfrak{X}} \circ c , \quad (2.154)$$

where $\tilde{\mathfrak{X}}$ is a vector field on Θ satisfying, for any $t \in I$, $\mathfrak{X}(t) = \tilde{\mathfrak{X}} \circ c(t)$. The reader should not confuse the index t in D_t , which is just a notation, with an actual real number $t \in I$.

A *geodesic* is thus a smooth curve $\gamma : I \rightarrow \Theta$, whose velocity $\dot{\gamma}$ is parallel along γ . If D_t is the covariant derivative along γ , then γ satisfies the geodesic equation [Lee19, page 103],

$$D_t \dot{\gamma}(t) = 0 , \quad t \in I . \quad (2.155)$$

The left-hand side of this equation is precisely the acceleration of the curve γ .

The geodesic equation is a non-linear ordinary differential equation of second order. For given initial conditions $\gamma(0) = \theta$ and $\dot{\gamma}(0) = u$, it has a unique solution $\gamma : (-\varepsilon, \varepsilon) \rightarrow \Theta$, for some $\varepsilon > 0$ [Lee19, Theorem 4.27]. If this solution can always be extended to a curve $\gamma : \mathbb{R} \rightarrow \Theta$, then Θ is called a complete Riemannian manifold. The Hopf-Rinow theorem states that this is equivalent to Θ being a complete metric space, with the distance function (2.150) [Lee19, Theorem 6.19].

2.11.4 Parallel Transport and Parallel Frames

Recall D_t the covariant derivative, associated with the Levi-Civita connection, along a curve $c : [0, 1] \rightarrow \Theta$, given in (2.154). Then, a vector field \mathfrak{X} is said to be parallel along c if it satisfies the parallel transport equation

$$D_t \mathfrak{X}(t) = 0 .$$

This is a first-order linear ordinary differential equation (ODE). Say $c(0) = \theta_0$, then, for a given initial condition $\mathfrak{X}(0) = u$, where $u \in T_{\theta_0}\Theta$, it follows that $u \mapsto \mathfrak{X}(t)$ is a linear mapping from $T_{\theta_0}\Theta$ to $T_{c(t)}\Theta$ [Lee19, Theorem 4.32]. This is denoted $T_{\theta_0}^c$, and by uniqueness of the solution to the ODE, $T_{\theta_0}^c$ is its linear inverse [Lee19, Equation (4.22)].

It is useful to derive an equivalent condition to (2.152), which holds for vector fields \mathfrak{X} and \mathfrak{Y} along c [Lee19, Proposition 5.5].

$$\frac{d}{dt} \langle \mathfrak{X}, \mathfrak{Y} \rangle_{c(t)} = \langle D_t \mathfrak{X}, \mathfrak{Y} \rangle_{c(t)} + \langle \mathfrak{X}, D_t \mathfrak{Y} \rangle_{c(t)} . \quad (2.156)$$

This equation yields that $t \mapsto \langle \mathfrak{X}, \mathfrak{Y} \rangle_{c(t)}$ is constant if \mathfrak{X} and \mathfrak{Y} are parallel vector fields along c . Thus T_{0t}^c preserves scalar products. In particular, if $(\mathbf{b}_i; i = 1, \dots, d)$ is an orthonormal basis of $T_{\theta_0} \Theta$, then the vector fields along c , defined by

$$e_i(t) = T_{0t}^c \mathbf{b}_i ,$$

form an orthonormal basis of the tangent space $T_{c(t)} \Theta$, for each $t \in I$. This is called a parallel orthonormal frame along c [Lee19, Equation (4.23)]. By linearity of T_{0t}^c , if $u \in T_{\theta_0} \Theta$ is written $u = \sum_{i=1}^d u^i \mathbf{b}_i$, then

$$T_{0t}^c u = \sum_{i=1}^d u^i e_i(t) . \quad (2.157)$$

In other words, parallel transport is obtained by simply propagating the coordinates u^i of the vector u along a parallel orthonormal frame.

2.11.5 Riemannian Exponential Map and Cut Locus

From now on, let us assume that Θ is a complete Riemannian manifold. A curve that attains the infimum in (2.150) is called a length-minimizing curve. While this curve is not always unique, it is always a geodesic [Lee19, Theorem 6.4]; in other words, it is a twice differentiable solution of the geodesic equation (2.155). On the other hand, it is important to keep in mind that a geodesic is not always a length-minimizing curve.

To give a concrete example, consider the geodesics of a sphere with its usual round metric [Lee19, Example 2.13] which are simply its great circles, i.e., intersections of the sphere with planes passing through the origin. Clearly, a portion of a great circle whose length is greater than π is not length-minimizing. Therefore, geodesics which start at some point θ on a sphere, are length-minimizing until they reach the opposite point $-\theta$. One says that the cut locus of the point θ on the sphere is the set $\{-\theta\}$.

Since Θ is complete, for $u \in T_{\theta} \Theta$, there exists a unique geodesic $\gamma_u : \mathbb{R} \rightarrow \Theta$ with $\gamma_u(0) = \theta$ and $\dot{\gamma}_u(0) = u$. Then, [Lee19, page 128] define the *exponential map* $\text{Exp} : T\Theta \rightarrow \Theta$ as:

$$\text{Exp}_{\theta}(u) = \gamma_u(1) , \quad (2.158)$$

which is a smooth map. For $u \in T_{\theta} \Theta$, with $\|u\|_{\theta} = 1$, let $c(u) > 0$ be the largest positive number t such that γ_u is length-minimizing when restricted to the interval $[0, t]$ [Lee19, page 307]:

$$c(u) = \sup \{t \geq 0 : \rho_{\Theta}(\theta, \gamma_u(t)) = t\} . \quad (2.159)$$

We call $c(u)$ the cut time of (θ, u) . The tangent cut locus of θ is then defined by

$$\text{TCut}(\theta) = \{c(u)u : u \in T_{\theta} \Theta ; \|u\|_{\theta} = 1\} .$$

Finally, the cut locus of θ is given by [Lee19, page 308]

$$\text{Cut}(\theta) = \text{Exp}_\theta \{ \text{TCut}(\theta) \} .$$

For example, if Θ is a unit sphere, then $\text{TCut}(\theta)$ is the set of tangent vectors $u \in T_\theta\Theta$ such that $\|u\|_\theta = \pi$. On the other hand, $\text{Cut}(\theta) = \{-\theta\}$, since $\|u\|_\theta = \pi$ implies $\text{Exp}_\theta(u) = -\theta$.

2.11.6 Injectivity Domain

The cut locus of a point in a complete Riemannian manifold gives valuable information regarding the topology of the manifold. Indeed, if Θ is a complete Riemannian manifold and θ is any point in Θ , then using [Lee19, Corollary 6.21], Θ can be decomposed into the disjoint union of two sets

$$\Theta = D(\theta) \cup \text{Cut}(\theta) , \tag{2.160}$$

where $D(\theta)$ is the injectivity domain,

$$D(\theta) = \text{Exp}_\theta \{ \text{TD}(\theta) \} , \text{ where } \text{TD}(\theta) = \{ tu : u \in T_\theta\Theta ; \|u\|_\theta = 1 ; 0 \leq t < c(u) \} ,$$

where $c(u)$ is given by (2.159). We observe:

Proposition 2.36. [Lee19, Theorem 10.34] *The Riemannian exponential map Exp_θ is a diffeomorphism of $\text{TD}(\theta)$ onto $D(\theta)$. Therefore, the inverse of the Riemannian exponential Exp_θ^{-1} is well-defined, and a diffeomorphism, on $D(\theta) = \Theta \setminus \text{Cut}(\theta)$.*

In fact, $\text{TD}(\theta)$ is an open, star-shaped subset of the tangent space $T_\theta\Theta$, so it has the topology of an open ball. Thus, (2.160) states that the topology of Θ is completely determined by $\text{Cut}(\theta)$.

2.11.7 Isometries and Homogeneous Spaces

An isometry g on the Riemannian manifold Θ is a diffeomorphism $g : \Theta \rightarrow \Theta$ which preserves the Riemannian metric. To express this, let $g \cdot \theta = g(\theta)$, and $g \cdot u = (Dg)_\theta(u)$ for each $\theta \in \Theta$ and $u \in T_\theta\Theta$. Here, $g(\theta) \in \Theta$ is simply the image of θ under the map g , and $(Dg)_\theta$ denotes the derivative of g at θ , so $(Dg)_\theta : T_\theta\Theta \rightarrow T_{g \cdot \theta}\Theta$. We say that g is an isometry if [Lee19, page 12]

$$\langle g \cdot u, g \cdot v \rangle_{g \cdot \theta} = \langle u, v \rangle_\theta \quad u, v \in T_\theta\Theta .$$

In other words, the linear map $u \mapsto g \cdot u$ from $T_\theta\Theta$ to $T_{g \cdot \theta}\Theta$ preserves scalar products. In particular, it also preserves norms, so $\|g \cdot u\|_{g \cdot \theta} = \|u\|_\theta$.

Isometries also preserve objects derived from the Riemannian metric such as distance, geodesics, among others. In particular, if $\gamma : I \rightarrow \Theta$ is a geodesic, and $g : \Theta \rightarrow \Theta$ is an isometry, then $\gamma' = g \circ \gamma : I \rightarrow \Theta$ is also a geodesic [Lee19, Corollary 5.14]. Now, if $\gamma(0) = \theta$ and $\dot{\gamma}(0) = u$, then $\gamma'(0) = g \cdot \theta$ and $\dot{\gamma}'(0) = g \cdot u$. From the definition of the Riemannian exponential (2.158), we get

$$g \cdot \text{Exp}_\theta(u) = \text{Exp}_{g \cdot \theta}(g \cdot u) . \tag{2.161}$$

The set \mathbf{G} of all isometries of a Riemannian manifold Θ forms a group under composition. A deep theorem, called Myers-Steenrod theorem, states that \mathbf{G} can always be given the structure of a Lie group, such that for each $\theta \in \Theta$, the group action $g \mapsto g \cdot \theta$ is a differentiable map [GHL04, page 66].

One calls Θ a Riemannian homogeneous space if its group of isometries \mathbf{G} acts transitively. Transitive action means that for any $\theta, \theta' \in \Theta$ there exists $g \in \mathbf{G}$ such that $g \cdot \theta = \theta'$. When Θ is a Riemannian homogeneous space, knowing the metric of Θ at just one point, $o \in \Theta$, is enough to know this metric anywhere [GHL04, page 67].

2.11.8 Riemannian Gradient, Hessian

The metric tensor $\langle \cdot, \cdot \rangle$ and Levi-Civita connection ∇ , on the Riemannian manifold Θ , can be used to generalize classical objects from analysis, like the gradient and Hessian of a C^2 function $V : \Theta \rightarrow \mathbb{R}$, as we introduce next.

The Riemannian gradient of V is a vector field $\text{grad } V$ on Θ , uniquely defined by the property [Lee19, Equation 2.14]

$$\langle \text{grad } V, u \rangle_\theta = (DV)_\theta(u), \quad u \in T_\theta\Theta, \quad (2.162)$$

where $(DV)_\theta : T_\theta\Theta \rightarrow \mathbb{R}$ is the differential of the function V at θ . As V is a real-valued function, it is useful to know that differentials, directional derivatives and covariant derivatives coincide $(DV)_\theta(u) = uV(\theta) = \nabla_u V(\theta)$. This definition makes it clear that the Riemannian gradient $\text{grad } V$ depends on the choice of metric on the manifold Θ , and does not arise from the manifold structure of Θ , in itself.

The Riemannian Hessian of V , denoted $\text{Hess } V$ is defined using the Levi-Civita connection. Precisely, it is the covariant derivative of the gradient $\text{grad } V$. For $\theta \in \Theta$, this gives the Hessian $\text{Hess } V(\theta) : T_\theta\Theta \rightarrow T_\theta\Theta$,

$$\text{Hess } V(\theta)u = \nabla_u \text{grad } V(\theta). \quad (2.163)$$

The Riemannian definition of the Hessian coincides with the covariant Hessian obtained from the Levi-Civita connection [Lee19, Example 4.22]. The covariant characterization gives, for any vector fields $\mathfrak{X}, \mathfrak{Y}$ on Θ ,

$$\langle \text{Hess } V \mathfrak{X}, \mathfrak{Y} \rangle = \nabla_{\mathfrak{X}} \nabla_{\mathfrak{Y}} V - \nabla_{\nabla_{\mathfrak{X}} \mathfrak{Y}} V = \mathfrak{X}(\mathfrak{Y} V) - (\nabla_{\mathfrak{X}} \mathfrak{Y}) V,$$

where the last equality comes from the remark regarding directional derivatives. One can see that (2.153) yields that the linear operator $\text{Hess } V(\theta) : T_\theta\Theta \rightarrow T_\theta\Theta$ is self-adjoint with respect to the Riemannian scalar product $\langle \cdot, \cdot \rangle_\theta$. Therefore, and not without a slight abuse of notation, we will also call Hess the resulting symmetric bilinear form

$$\text{Hess } V(\theta)(u, v) = \langle \text{Hess } V(\theta)u, v \rangle_\theta, \quad u, v \in T_\theta\Theta.$$

2.11.9 Taylor formula along a geodesic

To state the Taylor expansion theorem, we define the total covariant derivative, which allows us to write higher order derivatives as iterated total covariant derivatives. Let Θ be a smooth manifold and ∇ be a connection on $T\Theta$. Let $k, l \in \mathbb{N}$, then ∇ defines a connection in each tensor bundle $T^{(k,l)}T\Theta$, which we also denote by ∇ . See [Lee19, Proposition 4.15] for the properties of such connections and their proof. For every (k, l) -tensor field $F \in \Gamma(T^{(k,l)}T\Theta)$, the map ∇F defined over $\Gamma(T^{(0,1)}T\Theta) \times \dots \times \Gamma(T^{(0,1)}T\Theta) \times \Gamma(T^{(1,0)}T\Theta) \times \dots \times \Gamma(T^{(0,1)}T\Theta)$, where $\Gamma(T^{(0,1)}T\Theta)$ appears k times and $\Gamma(T^{(1,0)}T\Theta)$ appears $l + 1$ times, taking values in $C^\infty(\Theta)$, the smooth functions on Θ , given by

$$\nabla F : \omega^1, \dots, \omega^k, \mathfrak{Y}_1, \dots, \mathfrak{Y}_l, \mathfrak{Y} \mapsto (\nabla_{\mathfrak{Y}} F)(\omega^1, \dots, \omega^k, \mathfrak{Y}_1, \dots, \mathfrak{Y}_l),$$

defines a smooth $(k, l + 1)$ -tensor field on Θ called the total covariant derivative of F . See [Lee19, Proposition 4.17] for the proof.

Given a smooth curve $\gamma : I \rightarrow \Theta$ defined on an interval I , and any $k, l \in \mathbb{N}$, a (k, l) -tensor field along γ is a continuous map $F : I \rightarrow T^{(k,l)}T\Theta$, such that $F(t) \in T^{(k,l)}(T_{\gamma(t)}\Theta)$ for any $t \in I$, where $T^{(k,l)}T\Theta$ is the bundle of (k, l) -tensors on Θ , see e.g. [Lee19, Appendix B]. A vector field \mathfrak{Y} along γ is a $(1, 0)$ -tensor field, in which case for any $t \in I$, $\mathfrak{Y}(t)$ is just a tangent vector in $T_{\gamma(t)}\Theta$. We say that a tensor field F along γ is extendible if there exists a tensor field \tilde{F} defined on a neighborhood of $\gamma(I)$ such that $F = \tilde{F} \circ \gamma$.

We let $\mathcal{X}^{k,l}(\gamma)$ denote the set of smooth (k, l) -tensor fields along γ , and $\mathcal{X}(\gamma) = \mathcal{X}^{1,0}(\gamma)$ denote the set of smooth vector fields along γ . In particular, $\mathcal{X}^{0,0}(\gamma)$ is the set of smooth functions $g : I \rightarrow \mathbb{R}$ such that for any $t \in I$, $g(t) = (\gamma(t), f(t))$ for some smooth function $f : I \rightarrow \mathbb{R}$ and therefore can be identified with the set of smooth functions $f : I \rightarrow \mathbb{R}$. In the sequel, we adopt if no confusion is possible this identification. We extend to tensor fields along γ the following definition of the trace on tensors. For any (k, l) -tensor T , we denote by $\text{Tr}_{\square, \Delta}(T)$ the $(k - 1, l - 1)$ -tensor with component of index $(i_1, \dots, i_{k-1}, j_1, \dots, j_{l-1})$, given by $\sum_{m=1}^d T_{i_1, \dots, i_{k-1}, m, i_m, \dots, i_{k-1}}^{j_1, \dots, j_{l-1}, m, j_m, \dots, j_{l-1}}$. In particular, for any $\omega \in \mathcal{X}^{0,1}(\gamma)$, $\mathfrak{Y} \in \mathcal{X}(\gamma)$,

$$\text{Tr}_{(1,1)}(\omega \otimes \mathfrak{Y}) = \omega(\mathfrak{Y}).$$

Also, for any $F \in \mathcal{X}^{k,l}(\gamma)$, any $\omega^1, \dots, \omega^{k_0} \in \mathcal{X}^{0,1}(\gamma)$ and $\mathfrak{Y}_1, \dots, \mathfrak{Y}_{l_0} \in \mathcal{X}(\gamma)$, with $k_0 \leq k, l_0 \leq l$, denote by $[F : \omega^1 \otimes \dots \otimes \omega^{k_0} \otimes \mathfrak{Y}_1 \otimes \dots \otimes \mathfrak{Y}_{l_0}]$, the $(k - k_0, l - l_0)$ smooth tensor field along γ defined by the induction:

$$[F : \omega^{\otimes 1:i}] = \text{Tr}_{(1,l+1)}([F : \omega^{\otimes 1:(i-1)}] \otimes \omega^i) \quad (2.164)$$

$$[F : \omega^{\otimes 1:k_0} \otimes \mathfrak{Y}_{\otimes 1:j}] = \text{Tr}_{(k-k_0+1,1)}([F : \omega^{\otimes 1:k_0} \otimes \mathfrak{Y}_{\otimes 1:(j-1)}] \otimes \mathfrak{Y}_j), \quad (2.165)$$

setting $\omega^{\otimes 1:i} = \omega^1 \otimes \dots \otimes \omega^i$ and $\mathfrak{Y}_{\otimes 1:j} = \mathfrak{Y}_1 \otimes \dots \otimes \mathfrak{Y}_j$. Note that for any $\omega^{k-k_0+1}, \dots, \omega^k \in \mathcal{X}^{0,1}(\gamma)$ and $\mathfrak{Y}_{l-l_0+1}, \dots, \mathfrak{Y}_l \in \mathcal{X}(\gamma)$,

$$[F : \omega^{\otimes 1:k_0} \otimes \mathfrak{Y}_{\otimes 1:l_0}](\omega^{k-k_0+1}, \dots, \omega^k, \mathfrak{Y}_{l-l_0+1}, \dots, \mathfrak{Y}_l) = F(\omega^1, \dots, \omega^k, \mathfrak{Y}_1, \dots, \mathfrak{Y}_l). \quad (2.166)$$

Theorem 2.37. *Let Θ be a smooth manifold and ∇ be a connection on $T\Theta$. Let $\gamma : [0, 1] \rightarrow \Theta$ be a geodesic and $f : \Theta \rightarrow \mathbb{R}$ be a smooth function. Then, for any $t \in [0, 1], n \in \mathbb{N}$,*

$$\begin{aligned} f(\gamma(t)) &= \sum_{k_0=0}^n (t^{k_0}/k_0!) \nabla^{k_0} f_{\gamma(0)}(\dot{\gamma}(0), \dots, \dot{\gamma}(0)) \\ &\quad + \int_0^t [(t-s)^n/n!] \nabla^{n+1} f_{\gamma(s)}(\dot{\gamma}(s), \dots, \dot{\gamma}(s)) \, ds, \end{aligned} \quad (2.167)$$

where the $(0, k_0)$ -tensor field $\nabla^{k_0} f$ is the total derivative of order k_0 of the $(0, 0)$ -tensor field f .

For a definition of the total covariant derivative, see [Lee19, Proposition 4.17]. Also, in (3.97), remark that even though $\dot{\gamma}$ is only a vector field along γ , and not a vector field, the value of a tensor field $\nabla_{\mathfrak{X}} T$, where T is a smooth tensor field, evaluated at $\theta \in \Theta$ only depends on $\mathfrak{X}(\theta)$ and on values of T along a smooth curve $c : [0, 1] \rightarrow \Theta$ satisfying $c(0) = \theta$ and $\dot{c}(0) = \mathfrak{X}(\theta)$; see [Lee19, Proposition 4.26]. Therefore, an easy induction on k_0 proves that the expression $\nabla^{k_0} f_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t))$ in Theorem 2.37 is well defined for any $k_0 \in \mathbb{N}, t \in [0, 1]$.

Proof. Consider $\mathcal{V} : [0, 1] \rightarrow \Theta$ the smooth $(0, 0)$ -tensor field along γ defined by

$$\mathcal{V} = f \circ \gamma.$$

Then we check by induction on $n \in \mathbb{N}^*$ that \mathcal{V} is n -times differentiable with derivative of order n given for any $t \in [0, 1]$ by $\mathcal{V}^{(n)}(t) = D_t^n \mathcal{V}(t) = \nabla^n f_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t))$, where D_t is the covariant derivative operator along γ with respect to the connection ∇ , defined in [Dur+21, Proposition S14].

First, the case $n = 1$ is a direct application of [Lee19, Theorem 4.24] since f is an extension of \mathcal{V} . Assume now that the property holds for $n \in \mathbb{N}^*$. Then, for any $t_0 \in [0, 1]$, by smoothness of $\mathcal{V}^{(n)}$, [Lee19, Theorem 4.24] and the induction hypothesis,

$$\mathcal{V}^{(n+1)}(t_0) = D_t^{n+1} \mathcal{V}(t_0).$$

We now show that for any $t \in [0, 1]$, $D_t^{n+1} \mathcal{V}(t) = \nabla^{n+1} f_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t))$. Using [Dur+21, Lemma S15] on the smooth $(0, n)$ -tensor field along γ $F = (\nabla^n f) \circ \gamma$, we have

$$D_t [F : \dot{\gamma} \otimes \dots \otimes \dot{\gamma}] = [D_t F : \dot{\gamma} \otimes \dots \otimes \dot{\gamma}],$$

since $D_t \dot{\gamma} = 0$ because γ is a geodesic. Also, by (3.88), $[D_t F : \dot{\gamma} \otimes \dots \otimes \dot{\gamma}] = D_t F(\dot{\gamma}, \dots, \dot{\gamma})$. Finally, as $\nabla^n f$ is an extension of F , using the induction hypothesis and the definition of the total derivative give for any $t \in [0, 1]$,

$$\begin{aligned} D_t^{n+1} \mathcal{V}(t) &= D_t F(\dot{\gamma}, \dots, \dot{\gamma})(t) = \nabla_{\dot{\gamma}(t)}(\nabla^n f)_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t)) \\ &= (\nabla^{n+1} f)_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t)), \end{aligned}$$

concluding the induction.

Finally, (3.97) is simply a consequence of Taylor's formula with integral remainder of the real valued function \mathcal{V} . \square

If $\gamma : [0, 1] \rightarrow \Theta$ is a geodesic such that $\gamma(0) = \theta_0$ and $\gamma(1) = \theta_1$; which we simply call a geodesic between θ_0 and θ_1 , we have that the gradient defined by (2.162) also satisfies $\langle \text{grad } f_{\gamma(t)}, \dot{\gamma}(t) \rangle_{\gamma(t)} = D_t(f \circ \gamma)(t)$. Similarly, the Hessian defined by (2.163) satisfies, using [Dur+21, Proposition S14-(S72)-(iv)] and that $D_t \dot{\gamma} = 0$, we get for any $t \in [0, 1]$, $D_t^2(f \circ \gamma)(t) = \text{Hess } f_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t))$. These two arguments combined can be used to reformulate the result of Theorem 2.37 as follows

$$f(\theta_1) - f(\theta_0) = \langle \text{grad } f(\theta_0), \dot{\gamma}(0) \rangle_{\theta_0} + \int_0^1 (1-s) \text{Hess } f_{\gamma(s)}(\dot{\gamma}(s), \dot{\gamma}(s)) ds .$$

2.11.10 Bounded Hessian Implies Lipschitz Gradient

We can now state a result that can be very useful when the Riemannian Hessian is bounded.

Lemma 2.38. *If V has a continuous Riemannian Hessian $\text{Hess } V(\theta) : T_\theta \Theta \rightarrow T_\theta \Theta$, with operator norm upper bounded uniformly for $\theta \in \Theta$ by $C \geq 0$; then the Riemannian gradient $\text{grad } V$ satisfies the Lipschitz condition A 2-(b) with Lipschitz constant $L = C$.*

Proof. Let $(e_i; i = 1, \dots, d)$ be a parallel orthonormal frame along γ . Define $\nabla V^i : [0, 1] \rightarrow T_\theta \Theta$ by

$$\nabla V^i(t) = \langle \text{grad } V(\gamma(t)), e_i(t) \rangle_{\gamma(t)} , \quad t \in [0, 1] . \quad (2.168)$$

Also, note from (2.157), applied to $c = \gamma$ and $u = \text{grad } V(\theta_0)$, that

$$T_{\theta_1}^{\gamma} \text{grad } V(\theta_0) = \sum_{i=1}^d \nabla V^i(0) e_i(1) .$$

Then, since $(e_i(1); i = 1, \dots, d)$ is an orthonormal basis of $T_{\theta_1} \Theta$,

$$\|\text{grad } V(\theta_1) - T_{\theta_1}^{\gamma} \text{grad } V(\theta_0)\|_{\theta_1}^2 = \sum_{i=1}^d (\nabla V^i(1) - \nabla V^i(0))^2 \quad (2.169)$$

But, by applying (2.156) to $c = \gamma$ with $\mathfrak{X} = \text{grad } V \circ \gamma$ and $\mathfrak{Y} = e_i$, it follows from (2.154), (2.163) and (2.168),

$$\frac{d}{dt} \nabla V^i(t) = \langle \text{Hess } V(\gamma(t)) \dot{\gamma}(t), e_i(t) \rangle_{\gamma(t)} + \langle \text{grad } V(\gamma(t)), D_t e_i(t) \rangle_{\gamma(t)} . \quad (2.170)$$

Then, since each $e_i(t)$ is parallel along γ , $D_t e_i(t) = 0$. Plugging (2.170) into (2.169) yields, by the mean-value theorem, followed by Jensen's inequality,

$$\begin{aligned} \|\text{grad } V(\theta_1) - T_{\theta_1}^{\gamma} \text{grad } V(\theta_0)\|_{\theta_1} &= \left(\sum_{i=1}^d \left[\int_0^1 \langle \text{Hess } V(\gamma(t)) \dot{\gamma}(t), e_i(t) \rangle_{\gamma(t)} dt \right]^2 \right)^{1/2} , \\ &\leq \int_0^1 \left(\sum_{i=1}^d \left[\langle \text{Hess } V(\gamma(t)) \dot{\gamma}(t), e_i(t) \rangle_{\gamma(t)} \right]^2 \right)^{1/2} dt , \\ &= \int_0^1 \|\text{Hess } V(\gamma(t)) \dot{\gamma}(t)\|_{\gamma(t)} dt . \end{aligned}$$

The last equality comes from the fact that $(e_i(t); i = 1, \dots, d)$ is an orthonormal basis of $T_{\gamma(t)}\Theta$. Finally, using the definition of the operator norm $\|\cdot\|_{op, \gamma(t)}$ on $T_{\gamma(t)}\Theta$, the bound in the assumption and (2.149),

$$\|\text{grad } V(\theta_1) - T_{01}^\gamma \text{grad } V(\theta_0)\|_{\theta_1} \leq \int_0^1 \|\text{Hess } V(\gamma(t))\|_{op, \gamma(t)} \|\dot{\gamma}(t)\|_{\gamma(t)} dt \leq C\ell(\gamma).$$

This concludes the proof. □

CHAPTER 3

On Riemannian Stochastic Approximation Schemes with Fixed Step-Size

ALAIN DURMUS,¹ PABLO JIMÉNEZ¹, ÉRIC MOULINES¹, SALEM SAID².

This paper studies fixed step-size stochastic approximation (SA) schemes, including stochastic gradient schemes, in a Riemannian framework. It is motivated by several applications, where geodesics can be computed explicitly, and their use accelerates crude Euclidean methods. A fixed step-size scheme defines a family of time-homogeneous Markov chains, parametrized by the step-size. Here, using this formulation, non-asymptotic performance bounds are derived, under Lyapunov conditions. Then, for any step-size, the corresponding Markov chain is proved to admit a unique stationary distribution, and to be geometrically ergodic. This result gives rise to a family of stationary distributions indexed by the step-size, which is further shown to converge to a Dirac measure, concentrated at the solution of the problem at hand, as the step-size goes to 0. Finally, the asymptotic rate of this convergence is established, through an asymptotic expansion of the bias, and a central limit theorem.

Contents

3.1	Introduction	114
3.2	Constant stepsize analysis for a constrained scheme	117
3.2.1	Main Results	117
3.2.2	Two Examples of Lyapunov Functions	120
3.3	Asymptotic expansion and law in the unconstrained case	122
3.3.1	Asymptotic Expansion as $\eta \rightarrow 0$	122
3.3.2	A Central Limit Theorem on $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$	123

¹CMAP, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France (alain.durmus@polytechnique.edu, pablo.jimenez-moreno@polytechnique.edu, eric.moulines@polytechnique.edu).

²Laboratoire Jean Kuntzmann, CNRS, Université Grenoble Alpes (salem.said@univ-grenoble-alpes.fr).

3.4	Application to SGD	124
3.5	Numerical Experiments	127
3.6	Assumptions	130
3.7	Supplementary notation	131
3.8	Proofs of Section 3.2	132
3.8.1	Proof of Theorem 3.1	132
3.8.2	An alternative to Theorem 3.1-(b)	133
3.8.3	Proof of Theorem 3.2	134
3.8.4	Proof of Theorem 3.3	136
3.8.5	Proof of Proposition 3.4	136
3.8.6	Proof of Proposition 3.5	137
3.9	Proofs of Section 3.3	138
3.9.1	Proof of Theorem 3.6	140
3.9.2	Proof of Theorem 3.7	141
3.10	Proofs for Section 3.4	149
3.10.1	Proof of Lemma 3.8	149
3.10.2	Proof of Lemma 3.10	150
3.10.3	Proof of Proposition 3.11	151
3.10.4	Proof of Proposition 3.12	151
3.10.5	Proof of Theorem 3.13	152
3.11	Background on Markov chain theory and Riemannian geometry	156
3.11.1	Markov chain notions	157
3.11.2	Useful results from Riemannian geometry	157

3.1 Introduction

This paper deals with the study of fixed step-size Stochastic Approximation (SA) algorithms ([KY03; PJ92; RM51]), defined on a Riemannian manifold Θ with metric g . Specifically, consider the problem

$$\begin{aligned} &\text{find } \theta \in \Theta \text{ satisfying } h(\theta) = 0, \\ &\text{for a vector field } h : \Theta \rightarrow T\Theta, \end{aligned} \tag{3.1}$$

where $T\Theta$ denotes the tangent bundle of Θ , and h is only accessible through an oracle returning noisy estimates. The setting where $h = -\text{grad} f$ is of particular interest for minimizing a smooth function $f : \Theta \rightarrow \mathbb{R}$. In the Euclidean setting, Stochastic Gradient Descent (SGD) and its variants are now common methods for solving this problem ([Bot10; BB08a]). However, it should be stressed that (3.1) encompasses several other applications in stochastic optimization, reinforcement learning or maximum likelihood estimation, such as online Expectation Maximization algorithms ([CM09]), policy gradient ([BB01]) or Q-learning ([JJS93]). Minimization over a Riemannian manifold or its general formulation

(3.1) arises in many applications: Principal Component Analysis ([EAS98]), dictionary recovery ([SQW17]), matrix completion ([BA11]), smooth semidefinite programs ([BVB16]), tensor factorization ([Ish+11]), and Riemannian barycenter estimation ([Arn+12; SM19]). This has motivated the development of a comprehensive framework for stochastic optimization problems on Riemannian manifolds. One of the first contributions in this field is [Bon13], which derives asymptotic convergence results for SA on Riemannian manifolds. Non-asymptotic results are obtained by [ZS16] for a geodesically convex function f . This study has been followed and completed by [SKM19; ZRS16] which introduce and analyze a Riemannian counterpart of the Stochastic Variance Reduced Gradient (SVRG) algorithm. Since then, many existing methods or results from the Euclidean case have been considered in a Riemannian setting. For example, [KL17] suggest a Riemannian stochastic primal-dual algorithm and most recently [Tri+18] study an averaged version of Riemannian SGD.

In this paper, we are interested in the study of fixed step-size SA methods of the form

$$\begin{aligned} \theta_{n+1} &= \text{proj}_{\mathcal{S}} [\text{Exp}_{\theta_n} \{\eta H_{\theta_n}(X_{n+1})\}] , \\ \text{where } H_{\theta_n}(X_{n+1}) &= h(\theta_n) + e_{\theta_n}(X_{n+1}) . \end{aligned} \tag{3.2}$$

In (3.2), $\eta > 0$ is a step-size, $(X_n)_{n \in \mathbb{N}^*}$ is an $(\mathcal{F}_n)_{n \in \mathbb{N}}$ -adapted process, defined on a filtered probability space, with values in a measurable space $(\mathbf{X}, \mathcal{X})$, and $e : \Theta \times \mathbf{X} \rightarrow \text{T}\Theta$ is a measurable function, such that $\theta \mapsto e_{\theta}(x)$ is a vector field over Θ , for any $x \in \mathbf{X}$. In addition, $\text{Exp}_{\theta} : \text{T}_{\theta}\Theta \rightarrow \Theta$ is the Riemannian exponential mapping and $\text{proj}_{\mathcal{S}} : \Theta \rightarrow \mathcal{S}$ is a projection-like operator onto a subset $\mathcal{S} \subset \Theta$. This recursion is a natural extension of Euclidean SA, akin to the Robbins-Monroe algorithm, in a Riemannian setting.

In the Euclidean setting, the study of fixed step-size SA, and in particular SGD, has recently attracted much attention, see e.g. [Bac20; BM11; DDB17; MBB18; VBS19]. Indeed, first of all, the step-size η is the only parameter to tune, in contrast to the case where a decreasing sequence of step-sizes is used in (3.2). Furthermore, the forgetting of the initial condition is exponentially fast ([NB01; NWS14]).

We aim to show, in a general Riemannian framework, that the use of (3.2) provides a good solution for (3.1). To this end, we establish non-asymptotic and asymptotic properties of $(\theta_n)_{n \in \mathbb{N}}$, in the limit $\eta \rightarrow 0$. Our contributions can be summarized as follows.

- (1) We derive non-asymptotic bounds, for the convergence of $(\theta_n)_{n \in \mathbb{N}}$ to approximate solutions of (3.1), under general Lyapunov assumptions and mild assumptions on the manifold Θ and the subset \mathcal{S} .
- (2) Under additional regularity conditions, we show that $(\theta_n)_{n \in \mathbb{N}}$, as a Markov chain, admits a unique stationary distribution μ^{η} and is geometrically ergodic, *i.e.* converges to μ^{η} exponentially fast.
- (3) We study the limiting behavior of the family $(\mu^{\eta})_{\eta > 0}$ as $\eta \rightarrow 0$. In particular, we show that if (3.1) admits a unique solution θ^* and other suitable conditions hold, this family converges to the Dirac measure at θ^* . In addition, we asymptotically quantify this convergence, through a central limit theorem. Precisely, we prove that after a $\eta^{-1/2}$ -rescaling, this family of stationary distributions converges weakly to a normal distribution as $\eta \rightarrow 0$. These results illustrate the exponential forgetting of initial condition of the

scheme and that, at stationarity, the iterates $(\theta_n)_{n \in \mathbb{N}}$ stay in a $\mathcal{O}(\eta^{1/2})$ -neighborhood of θ^* . In addition, they can be understood as generalizations to Riemannian spaces of [Pfl86, Theorem 1] and [DDB17, Theorem 4].

(4) We apply our results to SGD. In particular, we establish the first non-asymptotic convergence bounds for strongly geodesically convex functions, without boundedness assumptions on the manifold Θ .

(5) Finally, we introduce and prove the convergence of an SGD scheme to compute the Riemannian barycenter, also known as the Karcher mean, of distributions on Hadamard manifolds. To the authors' knowledge, our contribution on this topic is one of the few without boundedness assumptions on the distribution.

For ease of reading, all the assumptions are gathered in the supplement Section 3.6. In the derivation of our results, we use crucially the fact that $(\theta_n)_{n \in \mathbb{N}}$ defines a Markov chain in Θ , under mild conditions. This interpretation has been successfully used in several papers dealing with the convergence of SA or SGD in Euclidean spaces; see e.g. [BMP90; FP99; KH81; Pfl86].

We consider a more general setting and milder conditions in comparison with most other studies in the field. Indeed, most papers do not consider the general SA framework, but only the case $h = -\text{grad } f$, dealing with SGD and its variants. To the authors' knowledge, only [Bon13; Dur+20] tackle the general SA problem (3.1). Our main contribution, compared to these two works, is to deal with the fixed step-size setting. Besides, our study considers general geodesically complete Riemannian manifolds which encompass Hadamard spaces, which have been the primary focus for [Tri+18; ZS16; ZRS16].

Furthermore, a majority of the previous studies on SGD in a Riemannian space (see e.g. [Ali+20; HG20; Tri+18; ZS16; ZRS16]), are purely local in nature, because of the assumption that $(\theta_n)_{n \in \mathbb{N}}$ stays almost surely in a (fixed and deterministic) compact and geodesically convex subset of Θ . For example, note that all the convergence results derived in [ZS16] depend on the diameter of the compact in which $(\theta_n)_{n \in \mathbb{N}}$ is assumed to stay. This assumption rarely holds in practice, and is quite difficult to verify in theory. It strongly limits the applicability of many results in the literature over the past few years. On the contrary, our results do not suffer from this problem, and can all be applied either on a compact or non-compact Riemannian manifold. As a result, we consider a new SA method to estimate the Karcher mean of a distribution π on Θ , see [Arn+12; IP18; Le04; ZS16], for which we derive non-asymptotic convergence bounds without boundedness conditions on the support of π .

Notations For any $\theta \in \Theta$ and $v, w \in T_\theta \Theta$, denote by $\mathfrak{g}_\theta(v, w) = \langle v, w \rangle_\theta$ and its corresponding norm by $\mathfrak{g}_\theta(v, v) = \|v\|_\theta^2$. $\rho_\Theta : \Theta \times \Theta \rightarrow \mathbb{R}_+$ denotes the distance associated with the Riemannian metric \mathfrak{g} . For any $\theta_0 \in \Theta, r > 0$, set $B(\theta_0, r) = \{\theta_1 \in \Theta : \rho_\Theta(\theta_0, \theta_1) < r\}$, the open ball centered at θ_0 with radius r . Similarly, we define closed balls in Θ by $\bar{B}(\theta_0, r) = \{\theta_1 \in \Theta : \rho_\Theta(\theta_0, \theta_1) \leq r\}$.

For a smooth function $g : \Theta \rightarrow \mathbb{R}$, we denote by $\text{grad } g$ its Riemannian gradient (see [Lee19, p. 27]) and by $\text{Hess } g$ its Riemannian, or covariant, Hessian (see [Lee19, Example 4.22]). For a curve $\gamma : I \rightarrow \Theta, T_{t_0, t_1}^\gamma : T_{\gamma(t_0)} \Theta \rightarrow T_{\gamma(t_1)} \Theta$ stands for the parallel transport map associated to the Levi-Civita connection along γ from $\gamma(t_0)$ to $\gamma(t_1)$ (see [Lee19,

Equation 4.22]). Moreover, for any $\theta \in \Theta$, under the assumption that Θ is complete, consider the Riemannian exponential map $\text{Exp}_\theta : T_\theta\Theta \rightarrow \Theta$, see [Lee19, Proposition 5.19]. This map projects a vector from the tangent space $T_\theta\Theta$ onto the manifold Θ , following a geodesic curve.

3.2 Constant stepsize analysis for a constrained scheme

3.2.1 Main Results

In this section, we study the Stochastic Approximation scheme (3.2), which is constrained on a subset $S \subset \Theta$. The following assumption on the manifold Θ and S is considered all along this paper and allows us to rigorously define proj_S .

A6. *Assume one of the following conditions.*

(i) Θ is a Hadamard manifold, i.e. a complete, simply connected Riemannian manifold with non-positive sectional curvature. In addition, S is a closed geodesically convex subset of Θ with non-empty interior.

(ii) Θ is a complete, connected Riemannian manifold and $S = \Theta$.

Note that under **A6**, the exponential map $\text{Exp} : T\Theta \rightarrow \Theta$ is well-defined, see [Lee19, Theorem 6.19]. Under **A6-(i)**, [Stu03, Proposition 2.6] shows that there exists $\text{proj}_S : \Theta \rightarrow S$ which is the Riemannian counterpart of the Euclidean projection onto a closed convex subset. More precisely, proj_S is the unique mapping from Θ to S such that for any $\theta \in \Theta$, $\rho_\Theta(\text{proj}_S(\theta), \theta) = \inf_{\theta' \in S} \rho_\Theta(\theta', \theta)$. Under **A6-(ii)**, we simply set $\text{proj}_S = \text{Id}$.

Recall that the recursion (3.2) only uses a noisy estimate H_θ of the mean field $h(\theta)$, for any $\theta \in \Theta$. We assume the following conditions on the noise to ensure convergence.

MD3. *The sequence $(X_n)_{n \in \mathbb{N}^*}$ is independent and identically distributed (i.i.d.). In addition, for any $\theta \in \Theta$, $\mathbb{E}[e_\theta(X_1)] = 0$ and there exist $\sigma_0^2, \sigma_1^2 > 0$ such that for any $\theta \in S$, $\mathbb{E}[\|e_\theta(X_1)\|_\theta^2] \leq \sigma_0^2 + \sigma_1^2 \|h(\theta)\|_\theta^2$.*

MD3 is referred to as the martingale difference setting which implies that $(\theta_n)_{n \in \mathbb{N}}$ is a time-homogeneous $(\mathcal{F}_n)_{n \in \mathbb{N}}$ -Markov chain, for which we denote by Q_η its corresponding Markov kernel. The Euclidean counterpart of this assumption consists in replacing the Riemannian norm with the Euclidean one, which is usual in Stochastic Approximation [see Duf97].

MD4. (i) \mathbb{P} -almost surely, the vector field $\theta \mapsto e_\theta(X_1)$ is continuous on Θ .

(ii) For any $\theta \in \Theta$, Leb_θ and the distribution of $e_\theta(X_1)$ are mutually absolutely continuous, where Leb_θ stands for the Lebesgue measure on $T_\theta\Theta$.

MD4 ensures topological and aperiodicity properties of the Markov chain under consideration. This condition is used in the study of the limiting behaviour of $(\theta_n)_{n \in \mathbb{N}}$. Note that the condition **MD4-(ii)** is automatically satisfied adding some Gaussian noise, i.e., when $e_\theta(X_i)$ is replaced by $e_\theta(X_i) + p_\theta(Z_i)$ where for any $\theta \in \Theta$, p_θ is any invertible linear application from \mathbb{R}^d to $T_\theta\Theta$ and $(Z_i)_{i \in \mathbb{N}^*}$ is a sequence of i.i.d. d -dimensional Gaussian

random variables with zero-mean and covariance matrix identity. The same assumption is considered in the Euclidean case in [Pfl86, Assumption A(v)-A(vi)].

To ensure recurrence of $(\theta_n)_{n \in \mathbb{N}}$, we assume the existence of a Lyapunov function $V : \Theta \rightarrow \mathbb{R}_+$ for the mean vector field h .

A6. (i) For any $\theta \in \Theta$, $V \circ \text{proj}_{\mathcal{S}}(\theta) \leq V(\theta)$.

(ii) V is continuously differentiable on Θ and its Riemannian gradient $\text{grad} V$ is geodesically L -Lipschitz, i.e., there exists $L \geq 0$ such that for any $\theta_0, \theta_1 \in \Theta$, and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ such that $\gamma(0) = \theta_0$ and $\gamma(1) = \theta_1$,

$$\|\text{grad} V(\theta_1) - T_{01}^{\gamma} \text{grad} V(\theta_0)\|_{\theta_1} \leq L\ell(\gamma), \quad (3.3)$$

where $\ell(\gamma) = \|\dot{\gamma}(0)\|_{\theta_0}$ is the length of the geodesic.

(iii) V is proper on \mathcal{S} , i.e., for any $M \geq 0$, there exists a compact set $\mathcal{K} \subset \mathcal{S}$ such that for any $\theta \in \mathcal{S} \setminus \mathcal{K}$, $V(\theta) > M$.

A7. There exist $C_1 \geq 0$ and $C_2 > 0$ such that for any $\theta \in \mathcal{S}$, $\|h(\theta)\|_{\theta}^2 + C_2 \langle \text{grad} V(\theta), h(\theta) \rangle_{\theta} \leq C_1$.

In addition, to quantify the convergence of $(\theta_n)_{n \in \mathbb{N}}$ in a neighborhood of a solution of (3.1), we consider the following condition for some compact set $\mathcal{K}^* \subset \mathcal{S}$.

A8 (\mathcal{K}^*). There exists $\lambda > 0$ such that for any $\theta \in \mathcal{S}$, $\langle \text{grad} V(\theta), h(\theta) \rangle_{\theta} \leq -\lambda V(\theta) \mathbb{1}_{\mathcal{S} \setminus \mathcal{K}^*}(\theta)$.

Note that under **A8**(\emptyset), if $h(\theta) = 0$, then $V(\theta) = 0$ since V is a nonnegative function.

It is relevant to recognize that **A6**, **A7** and **A8** boil down to standard stability and recurrence conditions; see e.g. [BMP90; Duf97]. In the Euclidean case when we assume the uniqueness of a solution θ^* , a common choice for V is $\theta \mapsto \|\theta - \theta^*\|^2$ because **A6** is always satisfied. In this context, **A7** and **A8** would read $\|h(\theta)\| \leq A + B\|\theta - \theta^*\|$ and $\inf_{\theta \in \mathcal{S} \setminus \mathcal{K}} \langle \theta - \theta^*, h(\theta^*) \rangle / \|\theta - \theta^*\|^2 \geq \lambda$ respectively. They are therefore the Riemannian counterparts of [Pfl86, Assumption A(i)-A(ii)]. However, the square distance is no longer a suitable candidate in non-compact Riemannian settings, and therefore selecting a Lyapunov function adapted to the manifold Θ and the geometry of the mean field h is all the more important. Note that **A6**-(iii) is automatically satisfied if \mathcal{S} is compact. In addition, in most cases \mathcal{K}^* and V are chosen such that $\mathcal{K}^* = \emptyset$ or $\{\theta \in \mathcal{S} : \|h(\theta)\|_{\theta} \leq \varepsilon\}$ for some $\varepsilon \geq 0$, $-C_2 \langle h(\theta), \text{grad} V(\theta) \rangle_{\theta} \geq \|h(\theta)\|_{\theta}^2$ for some $C_2 > 0$ and any $\theta \in \Theta$, and therefore **A7** is satisfied with $C_1 = 0$.

The use of Lyapunov functions is really common and widespread to analyze stochastic approximation schemes, see [Duf97; KH81; KY03]. However, compared to the Euclidean setting, the square distance cannot be used in many situations because it does not satisfy **A6**-(ii). This brought us to consider a different Lyapunov function and therefore develop an adapted framework for the Riemannian case; see Section 3.2.2 hereafter for more details.

We start with our first result which is established along with all the other statements of this section in the supplement Section 3.8.

Theorem 3.1. Assume **A6**, **MD3**, **A6**-(i)-(ii), **A7**.

3.2. Constant stepsize analysis for a constrained scheme

(a) Suppose in addition that for any $\theta \in \mathcal{S}$, $\langle \text{grad } V(\theta), h(\theta) \rangle_\theta \leq 0$. Then, for any $\eta \in (0, \bar{\eta}]$, $\theta_0 \in \mathcal{S}$, and $n \in \mathbb{N}^*$,

$$n^{-1} \sum_{k=0}^{n-1} \mathbb{E} [-\langle \text{grad } V(\theta_k), h(\theta_k) \rangle_{\theta_k}] \leq 2V(\theta_0)/(n\eta) + \eta b, \quad (3.4)$$

where $(\theta_n)_{n \in \mathbb{N}}$ is defined by (3.2) starting from θ_0 , $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1}$, $b = 2L\{\sigma_0^2 + C_1(1 + \sigma_1^2)\}$.

Suppose in addition that **A8**(\mathcal{K}^*) holds for some compact set $\mathcal{K}^* \subset \mathcal{S}$.

(b) Then for any $\eta \in (0, \bar{\eta}]$, $\theta_0 \in \mathcal{S}$, and $n \in \mathbb{N}^*$,

$$n^{-1} \sum_{k=0}^{n-1} \mathbb{E} [\mathbb{1}_{\mathcal{S} \setminus \mathcal{K}^*}(\theta_k) V(\theta_k)] \leq V(\theta_0)/(a n \eta) + \eta b/(2a), \quad (3.5)$$

where $a = \lambda/2$.

(c) Define $\|V\|_{\mathcal{K}^*} = \sup\{V(\theta) : \theta \in \mathcal{K}^*\}$ if $\mathcal{K}^* \neq \emptyset$ and $\|V\|_{\mathcal{K}^*} = 0$ otherwise. Then for any $\eta \in (0, \bar{\eta}]$, $\theta_0 \in \mathcal{S}$, and any $n \in \mathbb{N}^*$,

$$\mathbb{E}[V(\theta_n)] \leq \{1 - \eta a\}^n V(\theta_0) + \|V\|_{\mathcal{K}^*} + \eta b/(2a). \quad (3.6)$$

Note that Theorem 3.1 gives, in the case $\mathcal{K}^* = \emptyset$, non-asymptotic bounds of order η on $n^{-1} \sum_{k=0}^{n-1} \mathbb{E} [-\langle \text{grad } V(\theta_k), h(\theta_k) \rangle_{\theta_k}]$, $n^{-1} \sum_{k=0}^{n-1} \mathbb{E}[V(\theta_k)]$ and $\mathbb{E}[V(\theta_n)]$ as $n \rightarrow +\infty$. In addition, the forgetting of the initial condition in (3.4) and (3.5) is linear with respect to n , contrary to (3.6) where it is exponential. A statement similar to Theorem 3.1-(b) holds only assuming **A6**-(i)-(ii) and replacing **A8**(\mathcal{K}^*) by the condition that there exists $\lambda > 0$ such that for any $\theta \in \mathcal{S}$, $\langle \text{grad } V(\theta), h(\theta) \rangle_\theta \leq -\lambda \|h(\theta)\|_\theta^2 \mathbb{1}_{\mathcal{S} \setminus \mathcal{K}^*}(\theta)$. This result is postponed to the supplement Theorem 3.15-Section 3.8.2. Theorem 3.1-(a) is a generalization of [HS19, Lemma 7] for SGD under a general Lyapunov condition and milder assumptions. Theorem 3.1-(a) is also stated with very similar assumptions in [Dur+20, Theorem 2] – except that $\langle \text{grad } V(\theta), h(\theta) \rangle_\theta \leq 0$ in our setting is replaced with $\|\text{grad } V(\theta)\|_\theta \leq \bar{c} \|h(\theta)\|_\theta$ – where the result is not restricted to constant step-size settings. However, Theorem 3.1-(c) can only be obtained for constant step-size schemes, and the following results use it as a stepping stone. We show in Section 3.4, how this generalization can be applied to SGD to obtain better convergence guarantees. Finally, in the same Section, we show that Theorem 3.1-(b)-(c) can be used to derive non-asymptotic convergence bounds for SGD applied to a geodesically strongly convex function, without any boundedness assumptions on Θ .

The study of the asymptotic behavior of $(\theta_n)_{n \in \mathbb{N}}$ is the second step towards understanding the quality of the approximation to the solution of (3.1). We now show, under suitable assumptions and for $\eta \leq \bar{\eta}$ given in Theorem 3.1, first, that the chain is ergodic and admits a unique invariant distribution, and second, that this measure converges weakly to the Dirac measure at some point θ^* , as the stepsize of the scheme goes to zero. In other words, the family of stationary distributions $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$ concentrates around θ^* as $\eta \rightarrow 0$. Possible approximations of θ^* are therefore derived from sampling from μ^η or taking its Riemannian barycenter, for a small enough η . If the sequence $(\theta_n)_{n \in \mathbb{N}}$ is ergodic,

then as $n \rightarrow +\infty$ the marginal distributions of this Markov chain converge to μ^η and can be used in turn as proxy to solve (3.1). A remaining question is to provide an estimate of the approximation error as a function of the step-size η . This is tackled in Section 3.3.

Theorem 3.2. *Assume **A6**, **MD3**, **MD4**, **A6**, **A7** and **A8**(\mathbf{K}^*) for some compact set $\mathbf{K}^* \subset \mathbf{S}$. Let $\eta \in (0, \bar{\eta}]$ where $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1}$. Then, $(\theta_n)_{n \in \mathbb{N}}$ defined by (3.2) admits a unique stationary distribution μ^η and is Harris-recurrent. In addition, there exist $\rho \in [0, 1)$ and $C \geq 0$ such that for any $\theta_0 \in \mathbf{S}$ and $k \in \mathbb{N}$, $|\mathbb{E}[g(\theta_n)] - \int_{\Theta} g(\theta) d\mu^\eta(\theta)| \leq C\rho^n(1 + V(\theta_0))$, for any measurable function $g : \Theta \rightarrow \mathbb{R}$ satisfying $\sup_{\theta \in \Theta} \{|g|/V\} \leq 1$.*

Taking $n \rightarrow +\infty$ in Theorem 3.1-(c), we obtain by Theorem 3.2 that

$$|\int_{\Theta} g(\theta) d\mu^\eta(\theta)| \leq \|V\|_{\mathbf{K}^*} + \eta b/(2a), \quad (3.7)$$

for any measurable function $g : \Theta \rightarrow \mathbb{R}$ satisfying $\sup_{\theta \in \Theta} \{|g|/V\} \leq 1$. In the case $\|V\|_{\mathbf{K}^*} = 0$ (then $V(\theta) = 0$ for any $\theta \in \mathbf{K}^*$), we get $\int_{\Theta} V(\theta) d\mu^\eta(\theta) \leq \eta b/(2a)$. Therefore, this result indicates that the family $\{\mu^\eta : \eta \in (0, \bar{\eta}]\}$ concentrates in a $\mathcal{O}(\eta)$ -neighborhood of \mathbf{K}^* as $\eta \rightarrow 0$. In particular, if V admits a unique zero θ^* which corresponds in many applications to a solution of (3.1), then we can expect that $\{\mu^\eta : \eta \in (0, \bar{\eta}]\}$ converges in distribution to δ_{θ^*} , the Dirac measure at θ^* , as $\eta \rightarrow 0$. The specific additional conditions to obtain such a result are the following.

A9. *There exists $\theta^* \in \mathbf{S}$ such that for any $r > 0$, **A8**($\bar{\mathbf{B}}(\theta^*, r)$) holds and that there exists $c_r > 0$ satisfying for any $\theta \in \mathbf{S} \setminus \bar{\mathbf{B}}(\theta^*, r)$, $c_r \leq V(\theta)$.*

Note that assuming **A9** is weaker than assuming **A8**($\{\theta^*\}$) since in the first case the constant $\lambda > 0$ in **A8**($\bar{\mathbf{B}}(\theta^*, r)$) may depend on r . [Pfl86] considers a similar assumption in the Euclidean case [see Pfl86, Assumption A(i)].

As announced previously, we obtain the convergence in distribution of $\{\mu^\eta : \eta \in (0, \bar{\eta}]\}$.

Theorem 3.3. *Assume **A6**, **MD3**, **MD4**, **A6** and **A7** and let $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1}$.*

(a) *In addition suppose **A8**(\mathbf{K}^*) holds for some compact set $\mathbf{K}^* \subset \mathbf{S}$ and that there exists $c > 0$ such that for any $\theta \in \mathbf{S} \setminus \mathbf{K}^*$, $c \leq V(\theta)$. Then $\lim_{\eta \rightarrow 0} \mu^\eta\{\mathbf{K}^*\} = 1$, where μ^η is the stationary distribution of Q_η for $\eta \in (0, \bar{\eta}]$.*

(b) *In addition suppose **A9** holds. Then $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$ converges weakly to δ_{θ^*} , as $\eta \rightarrow 0$.*

3.2.2 Two Examples of Lyapunov Functions

Having stated the main results of this section, we give two examples of Lyapunov functions V under the following setting for Θ .

A7. *Θ is a Hadamard manifold. In addition, there exists $\kappa > 0$ such that the sectional curvature of Θ is bounded below by $-\kappa^2$.*

3.2. Constant stepsize analysis for a constrained scheme

A classical choice of Lyapunov function on Euclidean spaces is $\theta \mapsto \rho_{\Theta}^2(\theta, \theta^*)$, being both strongly convex and Lipschitz-gradient. However, this function does not satisfy **A6(ii)** as soon as Θ has non-zero curvature and is non-compact. In an effort to show the capital impact of curvature and in order to obtain a valid Lyapunov function satisfying the conditions **A6** and **A8**($\bar{B}(\theta^*, r)$) for $r > 0$, we now introduce the necessary assumptions and consider a truncated version of $\theta \mapsto \rho_{\Theta}^2(\theta, \theta^*)$.

Let $\mathbf{H} = \{\text{Exp}_{\theta}(tH_{\theta}(x)) : \theta \in \mathbf{S}, x \in \mathbf{X}, t \in [0, \eta]\}$ be the set of all points reached from geodesics $\gamma : [0, 1] \rightarrow \Theta$ of the form $\gamma(0) \in \mathbf{S}$ and $\dot{\gamma}(0) = \eta H_{\gamma(0)}(x)$, for any $x \in \mathbf{X}$. We assume in our next result that the closure of \mathbf{H} is compact which is implied for example in the case where \mathbf{S} is compact and $(\theta, x) \mapsto H_{\theta}(x)$ is bounded on $\mathbf{S} \times \mathbf{X}$.

Proposition 3.4. *Assume **A7** and that the closure $\bar{\mathbf{H}}$ of \mathbf{H} is compact, denote $D_{\mathbf{H}} = \text{diam}(\bar{\mathbf{H}})$. Consider a smooth function $\chi_{\mathbf{H}} : \Theta \rightarrow [0, 1]$ with compact support satisfying $\chi_{\mathbf{H}}(\theta) = 1$ for any $\theta \in \bar{\mathbf{H}}$ and for any $\theta \in \Theta$ such that $\inf_{\theta' \in \bar{\mathbf{H}}} \rho_{\Theta}(\theta', \theta) \geq 1$, it holds $\chi_{\mathbf{H}}(\theta) = 0$. Consider now $V_2 : \Theta \rightarrow \mathbb{R}_+$ defined for any $\theta \in \Theta$ by*

$$V_2(\theta) = \chi_{\mathbf{H}}(\theta) \rho_{\Theta}^2(\theta^*, \theta) + (1 - \chi_{\mathbf{H}}(\theta)) D_{\mathbf{H}}^2.$$

Then, **A6(i)-(ii)** holds with $V \leftarrow V_2$ and $L \leftarrow C_{\chi}(D_{\mathbf{H}} + 1)(1 + \kappa \coth(\kappa D_{\mathbf{H}}))$ where $C_{\chi} \geq 0$ is a constant only depending on $\chi_{\mathbf{H}}$. Suppose in addition that there exist $r > 0, \lambda_{\rho} > 0$ such that for any $\theta \in \mathbf{S}$,

$$-\left\langle \text{Exp}_{\theta}^{-1}(\theta^*), h(\theta) \right\rangle_{\theta} \leq -\lambda_{\rho} \rho_{\Theta}^2(\theta^*, \theta) \mathbb{1}_{\mathbf{S} \setminus \bar{B}(\theta^*, r)}(\theta). \quad (3.8)$$

Then **A8**($\bar{B}(\theta^*, r)$) holds with $\lambda \leftarrow \lambda_{\rho}$.

Note that under the setting of Proposition 3.4, $V_2(\theta) \geq c$ for any $\theta \in \mathbf{S} \setminus \bar{B}(\theta^*, r)$ by definition, since it is continuous. Clearly, **A6(iii)** does not hold for V_2 if \mathbf{S} is non-compact, since V_2 is constant outside of the support of $\chi_{\mathbf{H}}$. For this reason, and to weaken the assumptions of Proposition 3.4, we introduce a ‘‘Huberized’’ version of the distance to θ^* .

Proposition 3.5. *Assume **A7**. Let $\delta > 0$ and consider $V_1 : \Theta \rightarrow \mathbb{R}_+$ defined for any $\theta \in \Theta$ by*

$$V_1(\theta) = \delta^2 \left\{ (\rho_{\Theta}(\theta^*, \theta)/\delta)^2 + 1 \right\}^{1/2} - \delta^2. \quad (3.9)$$

Then, **A6** holds with $V \leftarrow V_1$ and $L \leftarrow 1 + \kappa\delta$. Suppose in addition that there exist $r > 0, \lambda_{\rho} > 0$ such that for any $\theta \in \mathbf{S}$, (3.8) holds. Then, **A8**($\bar{B}(\theta^*, r)$) holds and $\lambda \leftarrow \lambda_{\rho}$.

The proof is an extension of [Dur+20, Lemma 17] and relies heavily on this result. This Lyapunov function is still constructed upon the distance function, but as Proposition 3.5 shows, it is better suited for non-positive curvature spaces. Note that under the setting of Proposition 3.5, $V_1(\theta) \geq c$ for any $\theta \in \mathbf{S} \setminus \bar{B}(\theta^*, r)$ by definition since it is continuous.

It is worth mentioning that if either Proposition 3.4 or Proposition 3.5 can be applied, in order to use Theorem 3.1 and Theorem 3.2 (resp. Theorem 3.3-(b)) the only condition to verify (relative to the Lyapunov function) is **A7** (resp. are **A7** and **A9**).

3.3 Asymptotic expansion and law in the unconstrained case

The purpose of this section is to quantify the convergence derived in Theorem 3.3-(b). First, we establish an asymptotic expansion for the bias $\int_{\Theta} g(\theta) d\mu^{\eta}(\theta) - g(\theta^*)$ with respect to the step size η for g belonging to a certain class of smooth functions from Θ to \mathbb{R} . Our result can be applied to SGD ($h = -\text{grad } f$) and implies then an asymptotic expansion of $\int_{\Theta} \|\text{grad } f(\theta)\|_{\theta}^2 d\mu^{\eta}(\theta)$. Secondly, we establish that the convergence derived in Theorem 3.3-(b) occurs at a rate $\eta^{1/2}$, through a central limit theorem for $(\mu^{\eta})_{\eta \in (0, \bar{\eta}]}$. These two results can be understood as a bias-variance decomposition in which both terms are of order η and are therefore weak counterparts of [DDB17, Proposition 3, Theorem 5], [Pfl86, Theorem 1] in a Riemannian setting. The related proofs are postponed to the supplement Section 3.9.

3.3.1 Asymptotic Expansion as $\eta \rightarrow 0$

Here, we assume that A6-(ii) holds, Θ is compact and the conditions of Theorem 3.3-(b) hold. In addition, define the covariance tensor field Σ on Θ , for any $\theta \in \Theta$ by,

$$\Sigma(\theta) = \mathbb{E} [e_{\theta}(X_1) \otimes e_{\theta}(X_1)] . \quad (3.10)$$

Under appropriate conditions, letting $n \rightarrow +\infty$ in Theorem 3.1-(b), and Theorem 3.15 in the supplement, show respectively that $\int_{\Theta} V(\theta) d\mu^{\eta}(\theta)$ and $\int_{\Theta} \|h(\theta)\|_{\theta}^2 d\mu^{\eta}(\theta)$ are bounded by a term of order η . We specify this result in the case where $h = -\text{grad } f$ for a smooth objective function $f : \Theta \rightarrow \mathbb{R}$. More precisely, we establish in what follows a weak asymptotic expansion for $\int_{\Theta} \|\text{grad } f(\theta)\|_{\theta}^2 d\mu^{\eta}(\theta)$, as $\eta \rightarrow 0$ based on the following result for which we assume:

MD5. Σ is a continuous $(2, 0)$ -tensor field on Θ .

Denote the contraction of a covariant 2-tensor F with a contravariant 2-tensor G on Θ by $[F : G]$; see Section 3.11.2 (3.86)-(3.87) in the supplementary for more details. For two matrices A, B , $[A : B]$ just corresponds to $\text{Tr}(AB^{\top})$, where \top denotes the transpose.

Theorem 3.6. *Assume A6-(ii), h is continuous, $h(\theta^*) = 0$ and Θ is compact. Assume also MD3, MD4, MD5, A6, A7 and A9. Let $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1}$. Then for any $\eta \in (0, \bar{\eta}]$ and smooth function $g : \Theta \rightarrow \mathbb{R}$, we have*

$$- \int_{\Theta} \langle \text{grad } g(\theta), h(\theta) \rangle_{\theta} d\mu^{\eta}(\theta) = \frac{\eta}{2} [\text{Hess } g : \Sigma] (\theta^*) + \mathcal{R}_{g,\eta} ,$$

where $\lim_{\eta \rightarrow 0} \{|\mathcal{R}_{g,\eta}|/\eta\} = 0$.

Applying this result to SGD, i.e. $h = -\text{grad } f$ and $g = f$, we obtain that $\int_{\Theta} \|\text{grad } f(\theta)\|_{\theta}^2 d\mu^{\eta}(\theta) = (\eta/2) [\text{Hess } f : \Sigma] (\theta^*) + \mathcal{R}_{f,\eta}$ with $\lim_{\eta \rightarrow 0} \{|\mathcal{R}_{f,\eta}|/\eta\} = 0$.

3.3.2 A Central Limit Theorem on $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$

Now, we assume both **A6-(i)** and **A6-(ii)**, meaning $S = \Theta$ and Θ is a Hadamard manifold. Note that under this setting $\text{Exp}_\theta^{-1} : \Theta \rightarrow T_\theta \Theta$ is a well defined diffeomorphism for any $\theta \in \Theta$ by [Lee19, Proposition 12.9]. In addition, we assume that the other conditions of Theorem **3.3-(b)** hold. Following the approach of [Pfl86] in Euclidean SA, to find the asymptotic rate of convergence of the family $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$ defined in Section **3.2**, we establish a central limit theorem in $T_{\theta^*} \Theta$, for the family of pushforward measures $(\bar{\nu}^\eta)_{\eta \in (0, \bar{\eta}]}$ defined for any $A \in \mathcal{B}(T_{\theta^*} \Theta)$ by

$$\bar{\nu}^\eta(A) = \mu^\eta \left(\text{Exp}_{\theta^*}(\eta^{1/2} A) \right). \quad (3.11)$$

It is shown in Section **3.9** that for any $\eta \in (0, \bar{\eta}]$, $\bar{\nu}^\eta$ is the stationary distribution of the rescaled and projected Markov chain $(\bar{U}_n)_{n \in \mathbb{N}}$ defined for any $n \in \mathbb{N}$ by $\bar{U}_n = \eta^{-1/2} \text{Exp}_{\theta^*}^{-1}(\theta_n)$. Therefore, since under **A6-(i)-(ii)**, for any $u \in T_{\theta^*} \Theta$, $\rho_\Theta(\theta^*, \text{Exp}_{\theta^*}(u)) = \|u\|_{\theta^*}$ by [Lee19, Corollary 6.12, Proposition 12.9], showing a central limit theorem for the family $(\bar{\nu}^\eta)_{\eta \in (0, \bar{\eta}]}$ as $\eta \rightarrow 0$ shows that asymptotically $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$ concentrates in regions of diameter $\mathcal{O}(\eta^{1/2})$ around θ^* for the Riemannian distance.

We consider the following assumptions.

MD 6. *There exist $\varepsilon_e > 0$, $\tilde{\sigma}_0^2, \tilde{\sigma}_1^2 \geq 0$ such that for any $\theta \in \Theta$, $\mathbb{E}[\|e_\theta(X_1)\|_\theta^{2+\varepsilon_e}] \leq \tilde{\sigma}_0^2 + \tilde{\sigma}_1^2 V(\theta)$.*

A 10. *There exist a linear mapping $\mathbf{A} : T_{\theta^*} \Theta \rightarrow T_{\theta^*} \Theta$ and a map $\mathcal{H} : \Theta \rightarrow T_{\theta^*} \Theta$, such that for any $\theta \in \Theta$,*

$$h(\theta) = T_{01}^\gamma \left(\mathbf{A} \text{Exp}_{\theta^*}^{-1}(\theta) + \mathcal{H}(\theta) \right), \quad (3.12)$$

where θ^* is defined in **A9**, T_{01}^γ denotes parallel transport along the geodesic $\gamma : [0, 1] \rightarrow \Theta$ with $\gamma(0) = \theta^*$ and $\gamma(1) = \theta$, and $\lim_{\theta \rightarrow \theta^*} \{\|\mathcal{H}(\theta)\|_{\theta^*} / \rho_\Theta(\theta^*, \theta)\} = 0$. In addition, the eigenvalues of the matrix \mathbf{A} all have strictly negative real parts. Finally, there exists $C_3 > 0$ such that for any $\theta \in \Theta$, $\|h(\theta)\|_\theta \leq C_3 \rho_\Theta(\theta^*, \theta)$.

MD 6 is a strengthened version of **MD 3**. We show in Theorem **3.29** that (3.12) holds in the case h is twice continuously differentiable on Θ with $\mathbf{A} = \nabla h(\theta^*)$. Note that when $\Theta = \mathbb{R}^d$ is equipped with its Euclidean metric, (3.12) translates into $h(\theta) = \mathbf{A}(\theta - \theta^*) + \mathcal{H}(\theta)$ with $\lim_{\theta \rightarrow \theta^*} \{\|\mathcal{H}(\theta)\| / \|\theta - \theta^*\|\} = 0$, which is exactly the assumption in [Pfl86, Assumption B(ii)]. For ease of notation, we also denote by \mathbf{A} and $\Sigma(\theta^*)$ the matrices associated with these two linear applications in some orthonormal basis of $T_{\theta^*} \Theta$. **A10** guarantees the existence and uniqueness of the solution $\mathbf{V} \in \mathbb{R}^{d \times d}$ of the Lyapunov equation $\mathbf{A}\mathbf{V} + \mathbf{V}\mathbf{A}^\top = \Sigma(\theta^*)$, see [HJ94, Theorem 2.2.1].

We also assume that V can be compared to a function of the distance on Θ which leads to the strengthening of **A9**.

A 11. *There exists θ^* such that **A8**($\{\theta^*\}$) holds and there exists $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for any $\theta \in \Theta$, $V(\theta) \geq \phi(\rho_\Theta(\theta^*, \theta))$ and for any $r > 0$, $\inf_{[r, +\infty)} \phi > 0$. In addition, there exists $\bar{a} > 0$, such that $\lim_{r \rightarrow +\infty} \sup_{a \leq \bar{a}} a / \phi(a^{1/2} r) = 0$.*

Note that the assumption on the growth rate of the Lyapunov function is verified when $V = V_1$, considered in Proposition 3.5. In this case, we can take $\phi(r) = \delta^2[1+(r/\delta)^2]^{1/2} - \delta^2$. [Pfl86] only considers the Euclidean case with $V : \theta \mapsto \|\theta - \theta^*\|^2$, in which case **A11** boils down to $\langle \theta - \theta^*, h(\theta) \rangle \geq \|\theta - \theta^*\|^2 \lambda$. In addition, for this particular choice of Lyapunov function, it is equivalent to **A8**(\emptyset), as the other conditions result from the existence and uniqueness of the solution θ^* and the Euclidean properties of the squared distance function [see Pfl86, Assumption A(i)-B(iii)]. However, for our study, we have to generalize this assumption to any well-suited Lyapunov function, since, as we have already discussed, taking $V : \theta \mapsto \|\theta - \theta^*\|^2$ is no longer an option in the Riemannian case.

Theorem 3.7. *Assume **A6**-(i)-(ii), **MD3**, **MD4**, **MD5**, **MD6**, **A6**, **A7**, **A10** and **A11** hold. Suppose in addition that $h(\theta^*) = 0$, h is continuous and let $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1} \wedge (4C_3)^{-1}$. Then, the family of distributions $(\bar{\nu}^\eta)_{\eta \in (0, \bar{\eta}]}$, defined by (3.11), converges weakly to $N(0, \mathbf{V})$ as $\eta \rightarrow 0$ on $T_{\theta^*}\Theta$, where \mathbf{V} is the unique solution to the Lyapunov equation $\mathbf{A}\mathbf{V} + \mathbf{V}\mathbf{A}^\top = \Sigma(\theta^*)$.*

Even though $N(0, \mathbf{V})$ is a distribution on \mathbb{R}^d , we identify \mathbb{R}^d with $T_{\theta^*}\Theta$ using the same orthonormal basis as before. As mentioned in Section 3.2, Theorem 3.7 complements Theorem 3.3 because it proves that the asymptotic rate of convergence of $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$ to δ_{θ^*} is $\eta^{1/2}$, since $(\bar{U}_n)_{n \in \mathbb{N}}$ is rescaled by this factor with respect to the actual SA scheme $(\theta_n)_{n \in \mathbb{N}}$. Finally Theorem 3.7 can be seen as a Riemannian counterpart of [Pfl86, Theorem 1]. In the following section, we illustrate our results on SGD.

3.4 Application to SGD

We assume throughout this section that **A6**-(i)-(ii) holds. We apply the results of Section 3.2 and Section 3.3, to the unconstrained stochastic gradient scheme, *i.e.* $(\theta_n)_{n \in \mathbb{N}}$ defined by (3.2) with $h = -\text{grad } f$ and $S = \Theta$. Proofs are postponed to the supplement, Section 3.10.

Geodesically Strongly Convex and Smooth Function First, the objective function $f : \Theta \rightarrow \mathbb{R}$ is subject to the following assumptions.

F1. $f : \Theta \rightarrow \mathbb{R}$ is twice continuously differentiable and $\text{grad } f$ is geodesically L_f -Lipschitz, see (3.3).

F2. f is continuously differentiable on Θ and λ_f -strongly geodesically convex, for some $\lambda_f > 0$, *i.e.* for any $\theta_1, \theta_2 \in \Theta$, $f(\theta_2) \geq f(\theta_1) + \langle \text{Exp}_{\theta_1}^{-1}(\theta_2), \text{grad } f(\theta_1) \rangle_{\theta_1} + \lambda_f \rho_\Theta^2(\theta_1, \theta_2)$.

Under **F2**, f admits a unique minimizer denoted by θ^* . In addition, we have the following inequalities.

Lemma 3.8. *Assume **A6**-(i)-(ii) and **F2**. Then for any $\theta \in \Theta$, we have*

$$\begin{aligned} \|\text{grad } f(\theta)\|_\theta^2 &\geq \lambda_f(f(\theta) - f(\theta^*)) \quad \text{and,} \\ f(\theta) - f(\theta^*) &\geq \lambda_f \rho_\Theta^2(\theta, \theta^*). \end{aligned} \tag{3.13}$$

3.4. Application to SGD

Under **F1** and **F2**, Lemma 3.8 implies that $V(\theta) = f(\theta) - f(\theta^*)$ and $h = -\text{grad } f$ satisfy **A6** with $L \leftarrow L_f$, **A7** with $C_1 \leftarrow 0, C_2 \leftarrow 1$ and **A8**(\emptyset) with $\lambda \leftarrow \lambda_f$. A direct application of Theorem 3.1-(c) leads to the following result.

Corollary 3.9. *Assume **A6**-(i)-(ii), **MD3**, **F1**, **F2**. Consider $(\theta_n)_{n \in \mathbb{N}}$ defined by (3.2) with $h = -\text{grad } f$. Let $\bar{\eta} = [2L_f(1 + \sigma_1^2)]^{-1}$ and $\eta \in (0, \bar{\eta}]$. For any $\theta_0 \in \Theta$, and $n \in \mathbb{N}$,*

$$\mathbb{E}[f(\theta_n) - f(\theta^*)] \leq (1 - \eta\lambda_f/2)^n (f(\theta_0) - f(\theta^*)) + 2\eta L_f \sigma_0^2 / \lambda_f.$$

Then, setting $\eta = \bar{\eta} \wedge [\varepsilon\lambda_f / \{4\sigma_0^2 L_f\}]$, for $\varepsilon \in (0, 1)$, and setting

$$n = \lceil [\log(1/\varepsilon) - \log(f(\theta_0) - f(\theta^*))] / \log(1 - \eta\lambda_f/2) \rceil,$$

we get $\mathbb{E}[f(\theta_n) - f(\theta^*)] \leq \varepsilon$.

Corollary 3.9 shows that (3.2) has a computational complexity of order $\mathcal{O}(\log(1/\varepsilon)\varepsilon^{-1})$ to minimize f , without any boundedness assumptions on Θ , contrary to [ZS16]. In addition, Lemma 3.8 also implies that **A10** and **A11** hold if f is three times continuously differentiable and therefore Theorem 3.7 can be applied.

Geodesically Quasi-Convex Function with Bounded Gradient Consider the following assumption.

F3. *f is twice continuously differentiable. Further, there exists $\tilde{\lambda}_f > 0$ such that for any $\theta \in \Theta$, $-\langle \text{Exp}_\theta^{-1}(\theta^*), \text{grad } f(\theta) \rangle_\theta \geq \tilde{\lambda}_f V_1(\theta)$, where V_1 is defined by (3.9) with $\delta = 1$. In addition, there exists $C_f > 0$ such that for any $\theta \in \Theta$, $\|\text{grad } f(\theta)\|_\theta^2 \leq C_f(\rho_\Theta^2(\theta^*, \theta) \wedge 1)$.*

In the Euclidean case, this assumption corresponds to weak quasi-convexity as considered in [HMR19, Definition 2.1], $\langle \text{grad } f(\theta), \theta - \theta^* \rangle \geq \tau[f(\theta) - f(\theta^*)]$ where the Lyapunov function is $\theta \mapsto f(\theta) - f(\theta^*)$. Note that the geodesical quasi-convexity condition **F3** is a weaker version of the usual geodesical convexity in the SGD setting [see ZS16]. By introducing **F3**, we can relax the condition **F1** using the following result.

Lemma 3.10. *Assume **A7** and **F2**. Suppose in addition that f is twice continuously differentiable and there exists $M_f > 0$ such that for any $\theta \in \Theta$, $\|\text{grad } f(\theta)\|_\theta^2 \leq M_f \rho_\Theta^2(\theta^*, \theta)$. Let $\tilde{f} = \{f - f(\theta^*) + 1\}^{1/2}$. Then \tilde{f} satisfies **F3** with $C_f \leftarrow (M_f/4)[1 \wedge \lambda_f]$ and $\tilde{\lambda}_f \leftarrow \lambda_f / (2M_f^{1/2})$.*

Note that the condition introduced in Lemma 3.10 is a relaxation of the condition that $\text{grad } f$ is geodesically Lipschitz. Indeed, by [Jos05, Theorem 5.6.1], for $\theta^* \in \Theta$, $\theta \mapsto \rho_\Theta^2(\theta^*, \theta)$ satisfies the conditions of Lemma 3.10 but its gradient is not geodesically Lipschitz. A non-asymptotic bound is now given in terms of the distance-like function, defined for any $\theta_1, \theta_2 \in \Theta$ by

$$D_\Theta^2(\theta_1, \theta_2) = \rho_\Theta^2(\theta_1, \theta_2) / (1 + \rho_\Theta^2(\theta_1, \theta_2)). \quad (3.14)$$

Proposition 3.11. *Assume that **A 7**, **MD3**, **F3** hold. Let $\bar{\eta} = [(8C_f/\tilde{\lambda}_f)(1+\kappa)(1+\sigma_1^2)]^{-1}$ and $\eta \in (0, \bar{\eta}]$. Consider $(\theta_n)_{n \in \mathbb{N}}$ defined by (3.2) with $h = -\text{grad } f$ and $S = \Theta$. Then, for any $\theta_0 \in \Theta$ and $n \in \mathbb{N}^*$,*

$$n^{-1} \sum_{k=0}^{n-1} \mathbb{E} [D_{\Theta}^2(\theta^*, \theta_n)] \leq 4V_1(\theta_0)/(n\eta\tilde{\lambda}_f) + 4\eta(1+\kappa)\sigma_0^2/\tilde{\lambda}_f,$$

where κ is given in **A 7**, and V_1 is defined by (3.9) with $\delta = 1$.

To the authors' knowledge, such a bound is novel even in a deterministic setting.

Application to the Riemannian Barycenter Problem To conclude our study, we consider the problem of computing the Riemannian barycenter θ^* of a probability distribution π on a Hadamard manifold Θ . First, we look at the discrete case:

$$\pi = M_{\pi}^{-1} \sum_{i=1}^{M_{\pi}} \delta_{\bar{\theta}_i}, \quad (3.15)$$

where $M_{\pi} \in \mathbb{N}^*$ and $\{\bar{\theta}_i\}_{i=1}^{M_{\pi}} \in \Theta^{M_{\pi}}$. The Riemannian barycenter θ^* or Karcher mean of π [Arn+12] is the unique global minimum of the function $f_{\pi} : \theta \mapsto \sum_{i=1}^{M_{\pi}} \rho_{\Theta}^2(\theta, \bar{\theta}_i)/(2M_{\pi})$. By [Jos05, Theorem 5.6.1], $\text{grad } f_{\pi}(\theta) = -M_{\pi}^{-1} \sum_{i=1}^{M_{\pi}} \text{Exp}_{\theta}^{-1}(\bar{\theta}_i)$ for any $\theta \in \Theta$ and f_{π} satisfies **F 2** with $\lambda_f = 1/2$ using [Dur+20, Lemma 10]. Therefore, by Lemma 3.10, Proposition 3.11 can be applied. In addition, we get the following result, as an application of Proposition 3.4 and Theorem 3.1-(c).

Proposition 3.12. *Assume **A 7**. Let θ_{π}^* be the Riemannian barycenter of the probability measure π in (3.15) on the Hadamard manifold Θ , and let $(\theta_n)_{n \in \mathbb{N}}$ be given by $\theta_{n+1} = \text{Exp}_{\theta_n}(\eta \text{Exp}_{\theta_n}^{-1}(X_{n+1}))$, where $(X_n)_{n \in \mathbb{N}^*}$ is a sequence of i.i.d. random variables with distribution π . Then, for any $\eta \in (0, 1/(CL_{\pi}^3)]$, $\theta_0 \in \Theta$ and $n \in \mathbb{N}$,*

$$\mathbb{E}[\rho_{\Theta}^2(\theta_n, \theta_{\pi}^*)] \leq (1 - \eta/4)^n \rho_{\Theta}^2(\theta_0, \theta_{\pi}^*) + C\eta L_{\pi} D^2,$$

where $L_{\pi} = (1 + D)(1 + \kappa \coth(\kappa D))$, C is a universal constant, and $D = \max_{i=1, \dots, M_{\pi}} \rho_{\Theta}(\theta_0, \bar{\theta}_i)$.

Secondly, we tackle the general case where π is not required to be discrete or compactly supported. In this case, the mapping that we are looking to minimize is

$$f_{\pi} : \theta \mapsto (1/2) \int_{\Theta} \rho_{\Theta}^2(\theta, \nu) \pi(d\nu). \quad (3.16)$$

The function f_{π} is well-defined and finite under the following assumption.

MD7. *There exists $\theta \in \Theta$ such that $\int_{\Theta} \rho_{\Theta}^2(\theta, \nu) \pi(d\nu) < +\infty$.*

Note that by the triangle inequality, **MD7** is equivalent to for any $\theta \in \Theta$ such that $\int_{\Theta} \rho_{\Theta}^2(\theta, \nu) \pi(d\nu) < +\infty$ and therefore f_{π} is finite. Using the Lebesgue's dominated convergence theorem and [Jos05, Theorem 5.6.1], we can compute its Riemannian gradient given for any $\theta \in \Theta$ by, $\text{grad } f_{\pi}(\theta) = -\int_{\Theta} \text{Exp}_{\theta}^{-1}(\nu) \pi(d\nu)$. Then, f_{π} satisfies **F2** with $\lambda_f = 1/2$ and admits a unique minimizer θ_{π}^* . However, $\text{grad } f$ does not satisfy **F1** in general. More

3.5. Numerical Experiments

precisely, it fails to be geodesically Lipschitz, see [Jos05, Theorem 5.6.1]. In the Euclidean setting, several modifications of SGD have been suggested to rescale the gradient such as RMSProp, AdaGrad and Adam [DHS11; Geo14; KB17]. Inspired by these methods, we consider the stochastic approximation scheme (3.2) with $S = \Theta$ and

$$H_\theta(X_{n+1}) = (1/2)\text{Exp}_\theta^{-1}\left(X_{n+1}^{(1)}\right) \{\rho_\Theta^2(\theta, X_{n+1}^{(2)})/2 + 1\}^{-1/2}, \quad (3.17)$$

where $X_{n+1} = (X_{n+1}^{(1)}, X_{n+1}^{(2)})$ and $(X_k^{(1)}, X_k^{(2)})_{k \in \mathbb{N}^*}$ is an i.i.d. sequence of pairs of independent random variables with distribution π . The following result establishes non-asymptotic convergence bounds for the resulting recursion.

Theorem 3.13. *Assume A 7 and MD 7. Let θ_π^* be the Riemannian barycenter of the probability measure π . Let $(\theta_n)_{n \in \mathbb{N}}$ be given by (3.2) with $S = \Theta$ and H defined by (3.17). Then, for any $n \in \mathbb{N}$,*

$$n^{-1} \sum_{k=0}^{n-1} \mathbb{E} [D_\Theta^2(\theta_k, \theta_\pi^*)] \leq 4V_1(\theta_0)C_\pi^{1/2} / (\eta n) + 4\eta B_\pi,$$

where V_1 is defined by (3.9) with $\delta \leftarrow 1$, $\theta^* \leftarrow \theta_\pi^*$, $C_\pi = 1 + 2f_\pi(\theta_\pi^*)$, $B_\pi = (1 + \kappa)(f_\pi(\theta_\pi^*) + 1)(f_\pi(\theta_\pi^*) + 2)C_\pi^{-1/2}$ and D_Θ^2 is defined in (3.14).

3.5 Numerical Experiments

We consider in our experiments the Karcher mean estimation problem on $\Theta = \text{Sym}_{50}^+(\mathbb{R}) \subset \mathbb{R}^{50 \times 50}$, the symmetric definite positive matrix manifold (SPD) equipped with its affine-invariant metric, see [PFA06]. Note that the dimension of Θ is 1275.

We first consider the case where $\pi = (15)^{-1} \sum_{i=1}^{15} \delta_{x_i}$ is a discrete distribution, where $\{x_i\}_{i=1}^{15}$ are random samples from the Wishart distribution $\mathbf{W}(50, \text{Id})$ i.e. with 50 degrees of freedom and scale matrix identity. The Karcher mean θ_π^* associated with π is estimated using the Matrix Means Toolbox [BI13].

Figure 3.1 represents the behavior of the squared distance to the barycenter θ_π^* for a single path and three step-sizes $\eta \in \{10^{-3}, 4 \times 10^{-3}, 10^{-2}\}$. As expected from Proposition 3.12, two regimes can be observed. At first, the squared-distance to the barycenter exponentially decreases and then the iterates oscillate in a $\mathcal{O}(\eta^{1/2})$ -neighborhood of θ_π^* . In addition, the rate of convergence in the exponential decay depends on the step-size.

In Figure 3.2, we aim at illustrating (3.7), Theorem 3.6 and Theorem 3.7. To this end, 1000 replications of the previous experiment are performed to obtain $\{(\theta_n^{(i)}) : i \in \{1, \dots, 1000\}\}$ for $n = \lceil 10/\eta \rceil$ and $\eta \in \{1, 2.8, 4.6, 6.4, 8.2, 10\} \times 10^{-2}$. These samples are used to estimate the mean and the variance of $\rho_\Theta^2(\theta, \theta_\pi^*)$, for θ following the stationary distribution μ^η . As expected, the mean and the variance are both linear with respect to the step-size η , further confirming that the iterates remain in a neighborhood of diameter $\mathcal{O}(\eta^{1/2})$ to the ground truth.

Secondly, we examine the barycenter problem for $\pi = \mathbf{W}(50, \text{Id})$, following the scheme introduced in (3.17). The estimation of θ_π^* , relative to the new distribution π , is now done with a 100-batch-size version of our methodology, with 10^6 iterations and $\eta = 10^{-4}$.




Figure 3.1: Paths of the algorithm in Proposition 3.12




Figure 3.2




Figure 3.3: Paths of the algorithm in Theorem 3.13

As a counterpart to Figure 3.1, in Figure 3.3 we are interested in the mean values of $(D_{\Theta}^2(\theta_n, \theta_{\pi}^*))_{n \in \mathbb{N}}$ along a single path for three step-sizes $\eta \in \{10^{-3}, 4 \times 10^{-3}, 10^{-2}\}$, with respective burn-ins $\{13, 3.3, 1.645\} \times 10^3$. As predicted by Theorem 3.13, an initial decrease in $\mathcal{O}(n^{-1})$ is followed by a plateau in $\mathcal{O}(\eta)$. We can observe that compared to Figure 3.1, averaging smoothes oscillations.

Finally, we also perform the experiment corresponding to Figure 3.2 for the discrete setting to illustrate numerically that the conclusions of (3.7), Theorem 3.6 and Theorem 3.7 still hold. However, due to space constraints and since the conclusions are the same than for Figure 3.2, the corresponding figure is postponed to the supplement Figure 3.4.

Acknowledgments

AD and EM acknowledge support of the Lagrange Mathematical and Computing Research Center.

3.6 Assumptions

On the manifold

A 8. Assume one of the following conditions.

(i) Θ is a Hadamard manifold, i.e. a complete, simply connected Riemannian manifold with non-positive sectional curvature. In addition, S is a closed geodesically convex subset of Θ with non-empty interior.

(ii) Θ is a complete, connected Riemannian manifold and $S = \Theta$.

A 9. Θ is a Hadamard manifold. In addition, there exists $\kappa > 0$ such that the sectional curvature of Θ is bounded below by $-\kappa^2$.

On the distribution of the data

MD 8. The sequence $(X_n)_{n \in \mathbb{N}^*}$ is independent and identically distributed (i.i.d.). In addition, for any $\theta \in \Theta$, $\mathbb{E}[e_\theta(X_1)] = 0$ and there exist $\sigma_0^2, \sigma_1^2 > 0$ such that for any $\theta \in S$, $\mathbb{E}[\|e_\theta(X_1)\|_\theta^2] \leq \sigma_0^2 + \sigma_1^2 \|h(\theta)\|_\theta^2$.

MD 9. (i) \mathbb{P} -almost surely, the vector field $\theta \mapsto e_\theta(X_1)$ is continuous on Θ .

(ii) For any $\theta \in \Theta$, Leb_θ and the distribution of $e_\theta(X_1)$ are mutually absolutely continuous.

MD 10. Σ is a continuous tensor field of type $(2, 0)$ on Θ .

MD 11. There exist $\varepsilon_e > 0$, $\tilde{\sigma}_0^2, \tilde{\sigma}_1^2 \geq 0$ such that for any $\theta \in \Theta$, $\mathbb{E}[\|e_\theta(X_1)\|_\theta^{2+\varepsilon_e}] \leq \tilde{\sigma}_0^2 + \tilde{\sigma}_1^2 V(\theta)$.

MD 12. There exists $\theta \in \Theta$ such that

$$\int_{\Theta} \rho_{\Theta}^2(\theta, \nu) \pi(d\nu) < +\infty.$$

On the Lyapunov function V and the mean field function h

A 12. (i) For any $\theta \in \Theta$, $V \circ \text{proj}_S(\theta) \leq V(\theta)$.

(ii) V is continuously differentiable on Θ and its Riemannian gradient $\text{grad} V$ is geodesically L -Lipschitz, i.e., there exists $L \geq 0$ such that for any $\theta_0, \theta_1 \in \Theta$, and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ such that $\gamma(0) = \theta_0$ and $\gamma(1) = \theta_1$,

$$\|\text{grad} V(\theta_1) - T_{01}^{\gamma} \text{grad} V(\theta_0)\|_{\theta_1} \leq L \ell(\gamma),$$

where $\ell(\gamma) = \|\dot{\gamma}(0)\|_{\theta_0}$ is the length of the geodesic.

3.7. Supplementary notation

(iii) V is proper on S , i.e., for any $M \geq 0$, there exists a compact set $K \subset S$ such that for any $\theta \in S \setminus K$, $V(\theta) > M$.

A13. There exist $C_1 \geq 0$ and $C_2 > 0$ such that for any $\theta \in S$,

$$\|h(\theta)\|_\theta^2 + C_2 \langle \text{grad } V(\theta), h(\theta) \rangle_\theta \leq C_1.$$

A14 (K^*). There exists $\lambda > 0$ such that for any $\theta \in S$,

$$\langle \text{grad } V(\theta), h(\theta) \rangle_\theta \leq -\lambda V(\theta) \mathbb{1}_{S \setminus K^*}(\theta)$$

.

A15. There exists $\theta^* \in S$ such that for any $r > 0$, **A8**($\bar{B}(\theta^*, r)$) holds and that there exists $c_r > 0$ satisfying for any $\theta \in S \setminus \bar{B}(\theta^*, r)$, $c_r \leq V(\theta)$.

A16. There exist a linear mapping $\mathbf{A} : T_{\theta^*}\Theta \rightarrow T_{\theta^*}\Theta$ and a map $\mathcal{H} : \Theta \rightarrow T_{\theta^*}\Theta$, such that for any $\theta \in \Theta$,

$$h(\theta) = T_{01}^\gamma \left(\mathbf{A} \text{Exp}_{\theta^*}^{-1}(\theta) + \mathcal{H}(\theta) \right),$$

where θ^* is defined in **A9**, T_{01}^γ denotes parallel transport along the geodesic $\gamma : [0, 1] \rightarrow \Theta$ with $\gamma(0) = \theta^*$ and $\gamma(1) = \theta$, and $\lim_{\theta \rightarrow \theta^*} \{\|\mathcal{H}(\theta)\|_{\theta^*} / \rho_\Theta(\theta^*, \theta)\} = 0$. In addition, the eigenvalues of the matrix \mathbf{A} all have strictly negative real parts. Finally, there exists $C_3 > 0$ such that for any $\theta \in \Theta$, $\|h(\theta)\|_\theta \leq C_3 \rho_\Theta(\theta^*, \theta)$.

A17. There exists θ^* such that **A8**($\{\theta^*\}$) holds and there exists $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that for any $\theta \in \Theta$, $V(\theta) \geq \phi(\rho_\Theta(\theta^*, \theta))$ and for any $r > 0$, $\inf_{[r, +\infty)} \phi > 0$. In addition, there exists $\bar{a} > 0$, such that $\lim_{r \rightarrow +\infty} \sup_{a \leq \bar{a}} a / \phi(a^{1/2}r) = 0$.

On the objective function f in the gradient case

F4. $f : \Theta \rightarrow \mathbb{R}$ is twice continuously differentiable and $\text{grad } f$ is geodesically L_f -Lipschitz, see (3.3).

F5. f is λ_f -strongly geodesically convex, for some $\lambda_f > 0$, i.e. for any $\theta_1, \theta_2 \in \Theta$, $f(\theta_2) \geq f(\theta_1) + \langle \text{Exp}_{\theta_1}^{-1}(\theta_2), \text{grad } f(\theta_1) \rangle_{\theta_1} + \lambda_f \rho_\Theta^2(\theta_1, \theta_2) / 2$.

F6. f is twice continuously differentiable. There exists $\tilde{\lambda}_f > 0$ such that for any $\theta \in \Theta$, $-\langle \text{Exp}_\theta^{-1}(\theta^*), \text{grad } f(\theta) \rangle_\theta \geq \tilde{\lambda}_f V_1(\theta)$, where V_1 is defined by (3.9) with $\delta = 1$. In addition, there exists $C_f > 0$ such that for any $\theta \in \Theta$, $\|\text{grad } f(\theta)\|_\theta^2 \leq C_f (\rho_\Theta^2(\theta^*, \theta) \wedge 1)$.

3.7 Supplementary notation

Denote the unit tangent space $U_\theta\Theta = \{u \in T_\theta\Theta : \|u\|_\theta = 1\}$. The cut-locus of θ , $\text{Cut}(\theta) \subset \Theta$ [Lee19, p. 308] and the injectivity domain $\text{ID}(\theta) \subset T_\theta\Theta$ [Lee19, p. 310] are two notions that inform us about the length-minimizing properties of geodesics, and therefore provide

the domain of definition of the Riemannian exponential. On a complete and connected manifold, [Lee19, Theorem 10.34] holds, meaning the restriction $(\text{Exp}_\theta)|_{\text{ID}(\theta)} : \text{ID}(\theta) \rightarrow \Theta$ is a diffeomorphism onto its image $\Theta \setminus \text{Cut}(\theta)$. We simply denote $\text{Exp}_\theta^{-1} : \Theta \setminus \text{Cut}(\theta) \rightarrow \text{ID}(\theta)$ its inverse. Under the assumption that Θ is complete, simply connected and of non-positive sectional curvature, *i.e.* a Hadamard manifold, [Lee19, Proposition 12.9] proves that $\text{Cut}(\theta) = \emptyset$ and $\text{ID}(\theta) = T_\theta\Theta$ for any $\theta \in \Theta$.

For a measure μ on a measurable space (Y, \mathcal{Y}) , denote by $\mu(g)$ the integral of a measurable function $g : Y \rightarrow \mathbb{R}$ with respect to μ , when it exists.

3.8 Proofs of Section 3.2

Under **A6** and **MD3**, for any $\eta > 0$, we denote by Q_η the Markov kernel associated with $(\theta_n)_{n \in \mathbb{N}}$ defined by (3.2) given for any $A \in \mathcal{B}(S)$ and $\theta \in S$ by

$$Q_\eta(\theta, A) = \mathbb{E}[\mathbb{1}_A(\text{Exp}_\theta\{\eta H_\theta(X_1)\})] . \quad (3.18)$$

Useful notions, definitions and results relative to Markov chain theory are given in Section 3.11.1.

Lemma 3.14. *Assume **A6**, **MD3**, **A6-(i)-(ii)**. Then for any $\eta > 0$ and $\theta_0 \in S$,*

$$Q_\eta V(\theta_0) \leq V(\theta_0) + \eta \langle \text{grad } V(\theta_0), h(\theta_0) \rangle_{\theta_0} + L\eta^2 \left[\|h(\theta_0)\|_{\theta_0}^2 + \sigma_0^2 + \sigma_1^2 \|h(\theta_0)\|_{\theta_0}^2 \right] . \quad (3.19)$$

Proof. Let $\theta_0 \in S$, and $\eta > 0$. Consider

$$\theta_{1/2} = \text{Exp}_{\theta_0}[\eta H_{\theta_0}(X_1)] , \theta_1 = \text{proj}_S(\theta_{1/2}) . \quad (3.20)$$

First, by definition of Q_η and **A6-(i)**, we have

$$Q_\eta V(\theta_0) = \mathbb{E}[V(\theta_1)] \leq \mathbb{E}[V(\theta_{1/2})] . \quad (3.21)$$

Second, using **A6**, **A6-(ii)**, [Dur+20, Lemma 1] and (3.20), we obtain

$$V(\theta_{1/2}) \leq V(\theta_0) + \eta \langle \text{grad } V(\theta_0), H_{\theta_0}(X_1) \rangle_{\theta_0} + (L/2)\eta^2 \|H_{\theta_0}(X_1)\|_{\theta_0}^2 .$$

Plugging this result in (3.21) and using **MD3** completes the proof of (3.19). \square

3.8.1 Proof of Theorem 3.1

(a) Using Lemma 3.14 and **A7** we have for any $\theta_0 \in S$ and $\eta > 0$,

$$Q_\eta V(\theta_0) \leq V(\theta_0) + \eta \{1 - C_2 L \eta (1 + \sigma_1^2)\} \langle \text{grad } V(\theta_0), h(\theta_0) \rangle_{\theta_0} + L\eta^2 [\sigma_0^2 + C_1 (1 + \sigma_1^2)] .$$

Letting $\bar{\eta} = [2C_2 L (1 + \sigma_1^2)]^{-1}$, then for any $\eta \in (0, \bar{\eta}]$, we have $1 - C_2 L \eta (1 + \sigma_1^2) \geq 1/2$. Therefore, using also that $\langle \text{grad } V(\theta_0), h(\theta_0) \rangle_{\theta_0} \leq 0$, we obtain,

$$Q_\eta V(\theta_0) \leq V(\theta_0) + (\eta/2) \langle \text{grad } V(\theta_0), h(\theta_0) \rangle_{\theta_0} + L\eta^2 [\sigma_0^2 + C_1 (1 + \sigma_1^2)] . \quad (3.22)$$

Therefore, by the Markov property, for any $k \in \mathbb{N}^*$, $\eta \in (0, \bar{\eta}]$ and $\theta_0 \in \mathcal{S}$ we get,

$$-(\eta/2) \int_{\Theta} \langle \text{grad } V(\theta), h(\theta) \rangle_{\theta} Q_{\eta}^{k-1}(\theta_0, d\theta) \leq Q_{\eta}^{k-1}V(\theta_0) - Q_{\eta}^k V(\theta_0) + L\eta^2[\sigma_0^2 + C_1(1 + \sigma_1^2)].$$

Summing these inequalities for $k \in \{1, \dots, n\}$ concludes the proof of (a) upon using that V is a non-negative function.

(b) We prove (3.5) by using **A8**(\mathcal{K}^*) in (3.4) and dividing both sides by $\lambda > 0$.

(c) We start by using **A8**(\mathcal{K}^*) in (3.22). For any $\eta \in (0, \bar{\eta}]$ and $\theta_0 \in \mathcal{S}$, we have

$$Q_{\eta}V(\theta_0) \leq V(\theta_0) \left[1 - (\lambda\eta/2) \mathbb{1}_{\mathcal{S} \setminus \mathcal{K}^*}(\theta_0) \right] + \eta^2 b/2, \quad (3.23)$$

where $b = 2L[\sigma_0^2 + C_1(1 + \sigma_1^2)]$. By adding and subtracting $V(\theta_0)(\lambda\eta/2)\mathbb{1}_{\mathcal{K}^*}(\theta_0)$ in the right-hand side of (3.23), we have,

$$Q_{\eta}V(\theta_0) \leq V(\theta_0)[1 - \eta a] + \eta(b\eta/2 + a \|V\|_{\mathcal{K}^*}), \quad (3.24)$$

where $a = \lambda/2$. Therefore, by a straightforward induction on $n \in \mathbb{N}$, using the Markov property, we get, for any $n \in \mathbb{N}$, $\eta \in (0, \bar{\eta}]$ and $\theta_0 \in \mathcal{S}$,

$$\begin{aligned} \mathbb{E}[V(\theta_n)] &\leq \{1 - \eta a\}^n V(\theta_0) + \eta(b\eta/2 + a \|V\|_{\mathcal{K}^*}) \sum_{k=0}^{n-1} [1 - \eta a]^k \\ &\leq \{1 - \eta a\}^n V(\theta_0) + \{\|V\|_{\mathcal{K}^*} + (b\eta/2a)\}, \end{aligned}$$

which concludes the proof of (c) and Theorem 3.1.

3.8.2 An alternative to Theorem 3.1-(b)

Consider the following condition for some compact set $\mathcal{K}^* \subset \mathcal{S}$.

HS1 (\mathcal{K}^*). *There exists $\lambda > 0$ such that for any $\theta \in \mathcal{S}$, $\langle \text{grad } V(\theta), h(\theta) \rangle_{\theta} \leq -\lambda \|h(\theta)\|_{\theta}^2 \mathbb{1}_{\mathcal{S} \setminus \mathcal{K}^*}(\theta)$.*

Theorem 3.15. *Assume **A6**, **MD3**, **A6-(i)-(ii)** and **HS1**(\mathcal{K}^*) hold for some compact set $\mathcal{K}^* \subset \mathcal{S}$, and define $\|h\|_{\mathcal{K}^*} = \sup\{\|h(\theta)\|_{\theta} : \theta \in \mathcal{K}^*\}$ if $\mathcal{K}^* \neq \emptyset$ and $\|h\|_{\mathcal{K}^*} = 0$ otherwise. Then for any $\eta \in (0, \check{\eta}]$ and $\theta_0 \in \mathcal{S}$, and $n \in \mathbb{N}^*$,*

$$n^{-1} \sum_{k=0}^{n-1} \mathbb{E}[\mathbb{1}_{\mathcal{S} \setminus \mathcal{K}^*}(\theta_k) \|h(\theta_k)\|_{\theta_k}^2] \leq V(\theta_0)/(an\eta) + \eta \tilde{b}/a,$$

where $(\theta_n)_{n \in \mathbb{N}}$ is defined by (3.2) starting from θ_0 , $\check{\eta} = \lambda/[2(1 + \sigma_1^2)L]$, $a = \lambda/2$ and $\tilde{b} = L((1 + \sigma_1^2) \|h\|_{\mathcal{K}^*} + \sigma_0^2)$.

Proof. By Lemma 3.14 and **HS1**(\mathcal{K}^*), for any $\eta \in (0, \bar{\eta}]$ and $\theta_0 \in \mathcal{S}$, we have

$$Q_{\eta}V(\theta_0) \leq V(\theta_0) - \eta\lambda \|h(\theta_0)\|_{\theta_0}^2 \mathbb{1}_{\mathcal{S} \setminus \mathcal{K}^*}(\theta_0) + L\eta^2 \left[\|h(\theta_0)\|_{\theta_0}^2 + \sigma_0^2 + \sigma_1^2 \|h(\theta_0)\|_{\theta_0}^2 \right].$$

Therefore, by the Markov property, for any $k \in \mathbb{N}^*$, $\eta \in (0, \bar{\eta}]$ and $\theta_0 \in \mathcal{S}$, we get

$$\begin{aligned} (\eta\lambda/2) \int_{\Theta} \{\mathbb{1}_{\mathcal{S} \setminus \mathcal{K}^*}(\theta) \|h(\theta)\|_{\theta}^2\} Q_{\eta}^{k-1}(\theta_0, d\theta) \\ \leq Q_{\eta}^{k-1}V(\theta_0) - Q_{\eta}^kV(\theta_0) + L\eta^2((1 + \sigma_1^2) \|h\|_{\mathcal{K}^*} + \sigma_0^2). \end{aligned}$$

Summing these inequalities for $k \in \{1, \dots, n\}$ concludes the proof upon using that V is a non-negative function. \square

3.8.3 Proof of Theorem 3.2

Lemma 3.16. *Assume **A6**, **MD3** and **MD4-(i)**. Then the Markov kernel Q_{η} on $\mathcal{S} \times \mathcal{B}(\mathcal{S})$ is Feller, i.e. for any measurable bounded function $f : \mathcal{S} \rightarrow \mathbb{R}$, $Q_{\eta}f$ is continuous from \mathcal{S} to \mathbb{R} .*

Proof. The proof is an easy consequence of the Lebesgue dominated convergence theorem, since h is continuous and **MD4-(i)** holds. \square

For the next lemma, we introduce $\mu_{\mathcal{S}}$, the restriction to \mathcal{S} of the Riemannian measure μ_{Θ} associated with the volume form on Θ .

Lemma 3.17. *Assume **A6**, **MD3** and **MD4-(ii)**. Then Q_{η} is $\mu_{\mathcal{S}}$ -irreducible and aperiodic.*

Proof. We consider first the case **A6-(i)**, where Θ is a Hadamard manifold. Let $A \in \mathcal{B}(\mathcal{S})$ be a Borel set of \mathcal{S} , such that $\mu_{\mathcal{S}}(A) > 0$. We only need to show that for any $\theta_0 \in \Theta$, $Q_{\eta}(\theta_0, A) > 0$. Indeed, this gives $\mu_{\mathcal{S}}$ -irreducibility by definition and implies that the chain is aperiodic by [MT09, Theorem 5.4.4] since for any $A \in \mathcal{B}(\mathcal{S})$, $\mu_{\mathcal{S}}(A) > 0$, $\theta \in A$, we have $Q_{\eta}(\theta, A) > 0$.

Let $\theta_0 \in \mathcal{S}$. By definition of the scheme (3.2) and $\text{proj}_{\mathcal{S}}$, $Q_{\eta}(\theta_0, A) = \mathbb{P}(\text{proj}_{\mathcal{S}} \circ \text{Exp}_{\theta_0}(\eta\{h(\theta_0) + e_{\theta_0}(X_1)\}) \in A) \geq \mathbb{P}(\text{Exp}_{\theta_0}(\eta\{h(\theta_0) + e_{\theta_0}(X_1)\}) \in A)$. However, using **MD4-(ii)**, the law of $e_{\theta_0}(X_1)$ has a positive density $\phi : T_{\theta_0}\Theta \rightarrow (0, +\infty)$ with respect to Lebesgue's measure Leb_{θ_0} . Denote $(\mathfrak{g}_{ij}(\theta))_{1 \leq i, j \leq d}$ the matrix representing the Riemannian metric at $\theta \in \Theta$ in normal global coordinates at θ_0 . Expressing $\mu_{\mathcal{S}}$ in these coordinates and using [Lee19, p.404 and Proposition 2.41],

$$\begin{aligned} \mathbb{P}(\eta\{h(\theta_0) + e_{\theta_0}(X_1)\} \in \text{Exp}_{\theta_0}^{-1}(A)) &= \int_{\text{Exp}_{\theta_0}^{-1}(A)} \phi(\eta^{-1}v - h(\theta_0)) d\text{Leb}_{\theta_0}(v) \\ &= \int_A \phi(\eta^{-1}\text{Exp}_{\theta_0}^{-1}(\theta) - h(\theta_0)) \{\det(\mathfrak{g}_{ij}(\theta))\}^{-1/2} d\mu_{\mathcal{S}}(\theta) > 0, \end{aligned}$$

since all quantities in the integral are positive and $\mu_{\mathcal{S}}(A) > 0$.

Now assume **A6-(ii)** and keep the notations of the first case. Then $\text{Exp}_{\theta_0} : T_{\theta_0}\Theta \rightarrow \Theta$ is no longer a diffeomorphism. However, $(\text{Exp}_{\theta_0})|_{\text{ID}(\theta_0)} : \text{ID}(\theta_0) \rightarrow \Theta \setminus \text{Cut}(\theta_0)$ is a diffeomorphism, see [Lee19, Theorem 10.34]. Moreover, as $\text{Cut}(\theta_0)$ is a set of measure zero, see again [Lee19, Theorem 10.34], considering $\tilde{A} = A \setminus \text{Cut}(\theta_0)$ allows the previous proof to give the desired result. \square

3.8. Proofs of Section 3.2

Proof of Theorem 3.2. First, we prove that the chain is Harris-recurrent. For that, we start by proving, for any $\theta_0 \in \mathcal{S}$,

$$\mathbb{P} \left[\bigcup_{k \in \mathbb{N}^*} \bigcap_{N \in \mathbb{N}} \bigcup_{n \geq N} \{\theta_n \in \bar{\mathbb{B}}(\theta^*, k)\} \right] = 1, \quad (3.25)$$

where $(\theta_n)_{n \in \mathbb{N}}$ is defined by (3.2) and with initial condition θ_0 .

Theorem 3.1-(3.6) implies that for any $\theta_0 \in \Theta$, $\sup_{n \in \mathbb{N}} Q_\eta^n V(\theta_0) < +\infty$; since $\|V\|_{\mathcal{K}^*} = \sup_{\mathcal{K}^*} V < +\infty$ because V is assumed to be continuous. Therefore $\liminf_{n \rightarrow +\infty} V(\theta_n)$ is integrable by Fatou's lemma. Thus, for any $k \in \mathbb{N}^*$, using Markov's inequality,

$$\mathbb{P} \left(\liminf_{n \rightarrow +\infty} V(\theta_n) > k \right) \leq \mathbb{E} \left[\liminf_{n \rightarrow +\infty} V(\theta_n) \right] / k.$$

However, $\{\liminf_{n \rightarrow +\infty} V(\theta_n) \leq k\} = \bigcap_{N \in \mathbb{N}} \bigcup_{n \geq N} \{\theta_n \in V^{-1}([0, k])\}$. Thus, for any $k \in \mathbb{N}^*$,

$$\mathbb{P} \left(\bigcap_{N \in \mathbb{N}} \bigcup_{n \geq N} \{\theta_n \in V^{-1}([0, k])\} \right) \geq 1 - \mathbb{E} \left[\liminf_{n \rightarrow +\infty} V(\theta_n) \right] / k.$$

Now, taking the union of these events for any $k \in \mathbb{N}^*$ gives

$$\mathbb{P} \left(\bigcup_{k \in \mathbb{N}^*} \bigcap_{N \in \mathbb{N}} \bigcup_{n \geq N} \{\theta_n \in V^{-1}([0, k])\} \right) = 1. \quad (3.26)$$

Nonetheless, using A6-(iii), for any $k \in \mathbb{N}^*$, $V^{-1}([0, k])$ is a subset of a compact set, therefore it is bounded. Thus, for any $k \in \mathbb{N}^*$, there exists $k' \in \mathbb{N}^*$ such that $V^{-1}([0, k]) \subset \bar{\mathbb{B}}(\theta^*, k')$. This gives the following,

$$\bigcup_{k \in \mathbb{N}^*} \bigcap_{N \in \mathbb{N}} \bigcup_{n \geq N} \{\theta_n \in V^{-1}([0, k])\} \subset \bigcup_{k \in \mathbb{N}^*} \bigcap_{N \in \mathbb{N}} \bigcup_{n \geq N} \{\theta_n \in \bar{\mathbb{B}}(\theta^*, k)\}.$$

Combining this with (3.26) gives (3.25).

Equation (3.25) gives that the chain is non-evanescent [MT09, Section 9.2.1]. Since Q_η is Feller (see Lemma 3.16), this result and [MT09, Theorem 9.2.2] imply that Q_η is Harris recurrent.

We now show that Q_η is \tilde{V} -uniformly geometrically ergodic (see Section 3.11.1) setting $\tilde{V} = 1 + V$. First, by Theorem 3.1 and (3.24) obtained in the proof above, we have that for any $\theta_0 \in \mathcal{S}, \eta \in (0, \bar{\eta}]$,

$$Q_\eta \tilde{V}(\theta_0) \leq (1 - \eta a) \tilde{V}(\theta_0) + \eta(\eta b/2 + a(1 + \|V\|_{\mathcal{K}^*})),$$

where $a, b, \bar{\eta}$ and $\|V\|_{\mathcal{K}^*}$ are defined in Theorem 3.1. Then, by A6-(iii) there exists $\tilde{r} > 0$, such that for any $\theta_0 \in \mathcal{S}$,

$$Q_\eta \tilde{V}(\theta_0) \leq (1 - a\eta/2) \tilde{V}(\theta_0) + \eta(\eta b/2 + a(1 + \|V\|_{\mathcal{K}^*})) \mathbb{1}_{\bar{\mathbb{B}}(\theta^*, \tilde{r})}(\theta_0).$$

Then, since Q_η is Feller by Lemma 3.16 and $\mu_{\mathcal{S}}$ -irreducible by Lemma 3.17, using [MT09, Proposition 6.2.8 (ii)], $\bar{\mathbb{B}}(\theta^*, r)$ is petite since it is compact by the Hopf-Rinow theorem [Jos05, Theorem 1.7.1] and \mathcal{S} has non-empty interior by A6. Therefore, an application of [MT09, Theorem 16.0.1] proves that the chain is \tilde{V} -uniformly geometrically ergodic. \square

3.8.4 Proof of Theorem 3.3

Lemma 3.18. *Assume **A6**, **MD3** **MD4**, **A6**, **A7** and **A8**(\mathbf{K}^*) hold for some compact set $\mathbf{K}^* \subset \mathbf{S}$. Then for any $\eta \in (0, \bar{\eta}]$,*

$$\mu^\eta[V \mathbb{1}_{\mathbf{S} \setminus \mathbf{K}^*}] \leq 2\eta L \{\sigma_0^2 + C_1(1 + \sigma_1^2)\} / \lambda ,$$

where $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1}$.

Proof. For any $\eta \in (0, \bar{\eta}]$ and $M \geq 0$, setting $V_M = M \wedge V$, (3.23) implies using Jensen inequality, for any $\theta_0 \in \Theta$,

$$Q_\eta V_M(\theta_0) \leq (1 - \eta a \mathbb{1}_{\mathbf{S} \setminus \mathbf{K}^*}(\theta_0)) V_M(\theta_0) + \eta^2 b / 2 ,$$

where $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1}$, $b = 2L\{\sigma_0^2 + C_1(1 + \sigma_1^2)\}$ and $a = \lambda/2$. Using that μ^η is invariant for Q_η by Theorem 3.2 and V_M is bounded, we get $\mu^\eta[V_M \mathbb{1}_{\mathbf{S} \setminus \mathbf{K}^*}] \leq \eta b / (2a)$. By the monotone convergence theorem, taking $M \rightarrow +\infty$, we have $\mu^\eta[V \mathbb{1}_{\mathbf{S} \setminus \mathbf{K}^*}] \leq \eta b / (2a)$, which concludes the proof. \square

Proof of Theorem 3.3. (a) Using Lemma 3.18 and $V(\theta) \geq c > 0$ for any $\theta \in \mathbf{S} \setminus \mathbf{K}^*$, we obtain

$$\mu^\eta \{\mathbf{S} \setminus \mathbf{K}^*\} \leq \eta b / (2ac) ,$$

which concludes the proof of (a) taking the limit $\eta \rightarrow 0$.

(b) Let $(\eta_n)_{n \in \mathbb{N}}$ be a sequence converging to zero such that for any $n \in \mathbb{N}$, $\eta_n \in (0, \bar{\eta}]$. We start by proving that $(\mu^{\eta_n})_{n \in \mathbb{N}}$ is tight. Let $\varepsilon > 0$. On one hand, let $r > 0$ and $\mathbf{K}_0 = \bar{\mathbf{B}}(\theta^*, r)$. Then, using Theorem 3.3-(a), there exists $N \in \mathbb{N}$ such that for any $n \geq N$, $\mu^{\eta_n}(\mathbf{K}_0) \geq 1 - \varepsilon$. On the other hand, $(\mu^{\eta_n})_{n \in \{0, \dots, N-1\}}$ is tight, *i.e.* there exists a compact set $\tilde{\mathbf{K}} \subset \Theta$ such that for any $n \in \{1, \dots, N-1\}$, $\mu^{\eta_n}(\tilde{\mathbf{K}}) \geq 1 - \varepsilon$. Finally, taking $\mathbf{K} = \mathbf{K}_0 \cup \tilde{\mathbf{K}}$ gives the tightness of $(\mu^{\eta_n})_{n \in \mathbb{N}}$. Now, let μ be a limit point of $(\mu^{\eta_n})_{n \in \mathbb{N}}$. Using Theorem 3.3-(a), and Lebesgue's dominated convergence theorem letting $r \rightarrow 0$, gives $\mu(\{\theta^*\}) = 1$, *i.e.* $\mu = \delta_{\theta^*}$. In conclusion, for any $(\eta_n)_{n \in \mathbb{N}}$ converging to zero, $(\mu^{\eta_n})_{n \in \mathbb{N}}$ converges weakly to the Dirac at θ^* . \square

3.8.5 Proof of Proposition 3.4

First, we check **A6**-(i). Using [Stu03, Proposition 2.6], $\text{proj}_{\mathbf{S}}$ is a contraction with respect to ρ_Θ , which implies that for any $\theta \in \Theta$,

$$\rho_\Theta^2(\theta^*, \text{proj}_{\mathbf{S}}(\theta)) = \rho_\Theta^2(\text{proj}_{\mathbf{S}}(\theta^*), \text{proj}_{\mathbf{S}}(\theta)) \leq \rho_\Theta^2(\theta^*, \theta) .$$

This implies, since $\mathbf{S} \subset \mathbf{H}$, that

$$V_2(\text{proj}_{\mathbf{S}}(\theta)) = \rho_\Theta^2(\theta^*, \text{proj}_{\mathbf{S}}(\theta)) \leq \chi_{\mathbf{H}}(\theta) \rho_\Theta^2(\theta^*, \theta) + (1 - \chi_{\mathbf{H}}(\theta)) \text{diam}^2(\bar{\mathbf{H}}) = V_2(\theta) ,$$

which gives **A6**-(i).

To prove **A6-(ii)**, we calculate the operator norm of the Hessian of V_2 and conclude by [Dur+20, Lemma 10]. Using **A7** and [Jos05, Theorem 5.6.1], $\theta \mapsto \rho_{\Theta}^2(\theta^*, \theta)$ is smooth and its gradient on Θ is given by $\theta \mapsto -2\text{Exp}_{\theta}^{-1}(\theta^*)$. Therefore, for any $\theta \in \Theta$,

$$\text{grad } V_2(\theta) = [\rho_{\Theta}^2(\theta^*, \theta) - D_{\mathbb{H}}^2] \text{grad } \chi_{\mathbb{H}}(\theta) - 2\chi_{\mathbb{H}}(\theta) \text{Exp}_{\theta}^{-1}(\theta^*) .$$

Using now **A7**, [Jos05, Theorem 5.6.1] and Cauchy-Schwarz's inequality brings, for any $\theta \in \Theta, v \in T_{\theta}\Theta$,

$$\begin{aligned} \|(\text{Hess } V_2)_{\theta}(v, v)\|_{\theta} &\leq 2\kappa\rho_{\Theta}(\theta^*, \theta) \coth(\kappa\rho_{\Theta}(\theta^*, \theta))\chi_{\mathbb{H}}(\theta) \|v\|_{\theta}^2 + 4\rho_{\Theta}(\theta^*, \theta) \|\text{grad } \chi_{\mathbb{H}}(\theta)\|_{\theta} \|v\|_{\theta}^2 \\ &\quad + \|(\text{Hess } \chi_{\mathbb{H}})_{\theta}(v, v)\|_{\theta} \left| \rho_{\Theta}^2(\theta^*, \theta) - D_{\mathbb{H}}^2 \right| . \end{aligned}$$

However, one can choose $\chi_{\mathbb{H}}$ such that for any $\theta \in \Theta$ satisfying $\inf_{\theta' \in \mathbb{H}} \rho_{\Theta}(\theta', \theta) \geq 1$, it holds that $\chi_{\mathbb{H}}(\theta) = 0$. Therefore, for any $\theta \in \Theta$, $\rho_{\Theta}(\theta^*, \theta)\chi_{\mathbb{H}}(\theta) \leq D_{\mathbb{H}} + 1$. Since $\chi_{\mathbb{H}}$ is smooth with compact support, there exists a constant $M > 0$ such that for any $\theta \in \Theta$ and $v \in T_{\theta}\Theta$,

$$\|\text{grad } \chi_{\mathbb{H}}(\theta)\|_{\theta} \leq M \quad \text{and} \quad \|(\text{Hess } \chi_{\mathbb{H}})_{\theta}(v, v)\|_{\theta} \leq M \|v\|_{\theta}^2 .$$

Therefore, combining these expressions brings for any $\theta \in \Theta$ and $v \in T_{\theta}\Theta$,

$$\|(\text{Hess } V_2)_{\theta}(v, v)\|_{\theta} \leq 6(M + 1)(D_{\mathbb{H}} + 1)[1 + \kappa \coth(\kappa D_{\mathbb{H}})] \|v\|_{\theta}^2 ,$$

thus proving by [Dur+20, Lemma 10] and setting $C_{\chi} = 6(M + 1)$, that **A6-(ii)** holds with $L \leftarrow C_{\chi}(1 + D_{\mathbb{H}})[1 + \kappa \coth(\kappa D_{\mathbb{H}})]$.

We now turn on checking **A8**($\bar{\mathbb{B}}(\theta^*, r)$). Since $\text{grad } \chi_{\mathbb{H}}(\theta) = 0$ for any $\theta \in \mathbb{S}$, we get that V_2 is smooth and for any $\theta \in \mathbb{S}$, $\text{grad } V_2(\theta) = -2\text{Exp}_{\theta}^{-1}(\theta^*)$. Therefore **A8**($\bar{\mathbb{B}}(\theta^*, r)$) holds by (3.8).

3.8.6 Proof of Proposition 3.5

First, we check **A6-(i)**. Using [Stu03, Proposition 2.6], $\text{proj}_{\mathbb{S}}$ is a contraction with respect to ρ_{Θ} , which implies that $\theta \in \Theta$,

$$\rho_{\Theta}(\theta^*, \text{proj}_{\mathbb{S}}(\theta)) = \rho_{\Theta}(\text{proj}_{\mathbb{S}}(\theta^*), \text{proj}_{\mathbb{S}}(\theta)) \leq \rho_{\Theta}(\theta^*, \theta) .$$

Then the proof of **A6-(i)** is completed using that $x \mapsto \delta^2\{(x/\delta)^2 + 1\}^{1/2} - \delta^2$ is increasing.

Next, using **A7**, [Dur+20, Lemma 16], we have for any $\theta \in \Theta, v \in T_{\theta}\Theta \setminus \{0\}$,

$$0 < \text{Hess } V_1(\theta)(v, v) \leq (1 + \kappa\delta) \|v\|_{\theta}^2 .$$

Therefore, using [Dur+20, Lemma 10], **A6-(ii)** holds for $L = 1 + \kappa\delta$. It is easy to see that as $\rho_{\Theta}(\theta^*, \theta) \rightarrow \infty$, $V_1(\theta) \rightarrow +\infty$, meaning **A6-(iii)** holds by the Hopf-Rinow theorem [Jos05, Theorem 1.7.1].

Regarding **A8**($\bar{\mathbb{B}}(\theta^*, r)$), using [Dur+20, Lemma 16], we have for any $\theta \in \Theta$,

$$\text{grad } V_1(\theta) = -\text{Exp}_{\theta}^{-1}(\theta^*) / \left\{ (\rho_{\Theta}(\theta^*, \theta)/\delta)^2 + 1 \right\}^{1/2} , \quad (3.27)$$

Therefore for any $\theta \in \Theta$, we get

$$\langle \text{grad } V_1(\theta), h(\theta) \rangle_\theta = - \left\langle \text{Exp}_\theta^{-1}(\theta^*), h(\theta) \right\rangle_\theta / \left\{ (\rho_\Theta(\theta^*, \theta)/\delta)^2 + 1 \right\}^{1/2}.$$

Then, under the condition (3.8), we obtain

$$\begin{aligned} \langle \text{grad } V_1(\theta), h(\theta) \rangle_\theta &\leq -\lambda_\rho \rho_\Theta^2(\theta^*, \theta) \mathbb{1}_{S \setminus \bar{B}(\theta^*, r)}(\theta) / \left\{ (\rho_\Theta(\theta^*, \theta)/\delta)^2 + 1 \right\}^{1/2} \\ &\leq -\lambda_\rho V_1(\theta) \mathbb{1}_{S \setminus \bar{B}(\theta^*, r)}(\theta), \end{aligned}$$

where we used that

$$V_1(\theta) \leq \rho_\Theta^2(\theta^*, \theta) / \left\{ (\rho_\Theta(\theta^*, \theta)/\delta)^2 + 1 \right\}^{1/2},$$

since for any $a > 0$ and $x \geq 0$, $(ax^2 + 1)^{1/2} - 1 = a \int_0^x t \{at^2 + 1\}^{-1/2} dt \leq ax^2 / \{ax^2 + 1\}^{1/2}$.

3.9 Proofs of Section 3.3

For any $K \in \mathbb{R}_+$, consider a smooth function with compact support $\chi_K : \mathbb{R}_+ \rightarrow [0, 1]$ such that $\chi_K(t) = 1$ for any $t \leq K$ and $\chi_K(t) = 0$ for any $t \geq K + 1$.

Lemma 3.19. *Assume A 6-(ii) and MD 3.*

(a) *Then, for any smooth function with compact support $g : \Theta \rightarrow \mathbb{R}$, any $\eta > 0$ and $\theta_0 \in \Theta$,*

$$Q_\eta g(\theta_0) = g(\theta_0) + \eta \langle \text{grad } g(\theta_0), h(\theta_0) \rangle_{\theta_0} + (\eta^2/2) [\text{Hess } g : \Sigma + h \otimes h](\theta_0) + (\eta^2/6) \mathcal{R}_{g,\eta}(\theta_0), \quad (3.28)$$

where for any $K > 0$,

$$|\mathcal{R}_{g,\eta}(\theta_0)| \leq 8\eta \mathbb{E} \left[\|\nabla \text{Hess } g\|_{\gamma,\infty} \mathbb{1}_{A_{\theta_0}^c} \|H_K\|_{\theta_0}^3 \right] + 16 \|\text{Hess } g\|_\infty \mathbb{E} \left[\|Y_K\|_{\theta_0}^2 \right], \quad (3.29)$$

$$H_K = h(\theta_0) + e_{\theta_0}(X_1) \chi_K(\|e_{\theta_0}(X_1)\|_{\theta_0}), \quad Y_K = e_{\theta_0}(X_1) \{1 - \chi_K(\|e_{\theta_0}(X_1)\|_{\theta_0})\}, \quad (3.30)$$

$$\|\text{Hess } g\|_\infty = \sup\{|\text{Hess } g_\theta(u, u)| : \theta \in \Theta, u \in U_\theta\Theta\},$$

$$\|\nabla \text{Hess } g\|_{\gamma,\infty} = \sup\{|\nabla \text{Hess } g_{\gamma(t)}(u, u, u)| : t \in [0, 1], u \in U_{\gamma(t)}\Theta\},$$

$A_{\theta_0} = \{\|H_K\|_{\theta_0} \leq \|Y_K\|_{\theta_0}\}$ and $\gamma : [0, 1] \rightarrow \Theta$ is defined for any $t \in [0, 1]$ by $\gamma(t) = \text{Exp}_{\theta_0}(t\eta H_{\theta_0}(X_1))$.

(b) *Assume in addition that there exist $C_3 > 0$ and $\theta^* \in \Theta$ such that for any $\theta \in \Theta$, $\|h(\theta)\|_\theta \leq C_3 \rho_\Theta(\theta^*, \theta)$. Then, for any smooth function with compact support $g : \Theta \rightarrow \mathbb{R}$, any $\eta \in (0, (4C_3)^{-1}]$ and $\theta_0 \in \Theta$, (3.28) holds, with for any $K > 0$,*

$$|\mathcal{R}_{g,\eta}(\theta_0)| \leq 8\eta \mathbb{1}_{\mathcal{K}_K}(\theta_0) \mathbb{E} \left[\|\nabla \text{Hess } g\|_{\gamma,\infty} \mathbb{1}_{A_{\theta_0}^c} \|H_K\|_{\theta_0}^3 \right] + 16 \|\text{Hess } g\|_\infty \mathbb{E} \left[\|Y_K\|_{\theta_0}^2 \right], \quad (3.31)$$

where we take the notation of (a) and \mathcal{K}_K is a compact subset of Θ .

Proof. (a) Let $g : \Theta \rightarrow \mathbb{R}$ be a smooth function with compact support and $\theta_0 \in \Theta$. Using (3.2), A6-(ii) and the definition of Q_η (3.18), we have

$$Q_\eta g(\theta_0) = \mathbb{E} [g \{ \text{Exp}_{\theta_0} [\eta H_{\theta_0}(X_1)] \}] . \quad (3.32)$$

Consider the geodesic $\gamma : [0, 1] \rightarrow \Theta$ defined for any $t \in [0, 1]$ by $\gamma(t) = \text{Exp}_{\theta_0}(t\eta H_{\theta_0}(X_1))$. For any $t \in [0, 1]$, let $g(t) = (g \circ \gamma)(t)$. We compute now its derivatives to derive a Taylor expansion. Using [Lee19, Proposition 4.15-(ii) and Theorem 4.24-(iii)], we have for any $t \in [0, 1]$,

$$g'(t) = D_t(g \circ \gamma)(t) = \langle \text{grad } g(\gamma(t)), \dot{\gamma}(t) \rangle_{\gamma(t)} .$$

By definition of the Hessian [Lee19, Example 4.22] and using $D_t \dot{\gamma}(t) = 0$, Proposition 3.27-(3.89)-(iv), we get for any $t \in [0, 1]$,

$$g''(t) = [D_t^2 g](t) = \text{Hess } g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t)) ,$$

In addition, using $D_t \dot{\gamma}(t) = 0$ and Proposition 3.27-(3.89)-(iv), we obtain for any $t \in [0, 1]$,

$$g^{(3)}(t) = [D_t^3 g](t) = \nabla \text{Hess } g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t), \dot{\gamma}(t)) ,$$

where $\nabla \text{Hess } g$ is the total covariant derivative of $\text{Hess } g$ [Lee19, Proposition 4.17]. Finally, for any $K > 0$, consider the two random tangent vectors at θ_0 defined in (3.30). Now, writing the first-order Taylor expansion of $g : [0, 1] \rightarrow \mathbb{R}$, at $t = 1$ on the event $A_{\theta_0} = \{ \|H_K\|_{\theta_0} \leq \|Y_K\|_{\theta_0} \}$, the second-order one on the complement, and summing both expansions, we get

$$\begin{aligned} g(\text{Exp}_{\theta_0}(\eta H_{\theta_0}(X_1))) &= g(\theta_0) + \eta \langle \text{grad } g(\theta_0), H_{\theta_0}(X_1) \rangle_{\theta_0} \\ &\quad + (\eta^2/2) \text{Hess } g_{\theta_0}(H_{\theta_0}(X_1), H_{\theta_0}(X_1)) + \mathcal{R}_{g,\eta}(\theta_0, X_1)/6 , \end{aligned} \quad (3.33)$$

where the remainder term is given by

$$\begin{aligned} \mathcal{R}_{g,\eta}(\theta_0, X_1) &= \mathbb{1}_{A_{\theta_0}^c} \int_0^1 \nabla \text{Hess } g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t), \dot{\gamma}(t)) dt \\ &\quad + \mathbb{1}_{A_{\theta_0}} \left[\int_0^1 \text{Hess } g_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t)) dt - 3\eta^2 \text{Hess } g_{\theta_0}(H_{\theta_0}(X_1), H_{\theta_0}(X_1)) \right] . \end{aligned}$$

We bound the remainder as follows. Since g has compact support, $\text{Hess } g$ and $\nabla \text{Hess } g$ have an operator norm uniformly bounded over Θ , which we express in the following way. For any $\theta \in \Theta$, consider the unit tangent space at θ , $U_\theta \Theta = \{v \in T_\theta \Theta : \|v\|_\theta = 1\}$, let $\|\text{Hess } g\|_\infty = \sup\{|\text{Hess } g_\theta(v, v)| : \theta \in \Theta, v \in U_\theta \Theta\}$ and $\|\nabla \text{Hess } g\|_{\gamma, \infty} = \sup\{|\nabla \text{Hess } g_{\gamma(t)}(v, v, v)| : t \in [0, 1], v \in U_{\gamma(t)} \Theta\}$. Then, using [Lee19, Corollary 5.6-(b)], and $\dot{\gamma}(0) = \eta H_{\theta_0}(X_1)$,

$$\begin{aligned} |\mathcal{R}_{g,\eta}(\theta_0, X_1)| &\leq \mathbb{1}_{A_{\theta_0}^c} \|\nabla \text{Hess } g\|_{\gamma, \infty} \int_0^1 \|\dot{\gamma}(t)\|_{\gamma(t)}^3 dt \\ &\quad + \mathbb{1}_{A_{\theta_0}} \|\text{Hess } g\|_\infty \left[\int_0^1 \|\dot{\gamma}(t)\|_{\gamma(t)}^2 dt + 3\eta^2 \|H_{\theta_0}(X_1)\|_{\theta_0}^2 \right] \\ &= \mathbb{1}_{A_{\theta_0}^c} \|\nabla \text{Hess } g\|_{\gamma, \infty} \eta^3 \|H_{\theta_0}(X_1)\|_{\theta_0}^3 + 4\mathbb{1}_{A_{\theta_0}} \|\text{Hess } g\|_\infty \eta^2 \|H_{\theta_0}(X_1)\|_{\theta_0}^2 . \end{aligned}$$

Moreover, using that $H_K + Y_K = H_{\theta_0}(X_1)$ and the definition of A_{θ_0} ,

$$|\mathcal{R}_{g,\eta}(\theta_0, X_1)| \leq 8 \mathbb{1}_{A_{\theta_0}^c} \|\nabla \text{Hess } g\|_{\gamma, \infty} \eta^3 \|H_K\|_{\theta_0}^3 + 16 \|\text{Hess } g\|_{\infty} \eta^2 \|Y_K\|_{\theta_0}^2 . \quad (3.34)$$

Now, using **MD3**,

$$\mathbb{E} [\langle \text{grad } g(\theta_0), H_{\theta_0}(X_1) \rangle_{\theta_0}] = \langle \text{grad } g(\theta_0), h(\theta_0) \rangle_{\theta_0} . \quad (3.35)$$

In addition, since

$$\text{Hess } g_{\theta_0}(H_{\theta_0}(X_1), H_{\theta_0}(X_1)) = [\text{Hess } g : H_{\theta_0}(X_1) \otimes H_{\theta_0}(X_1)] ,$$

it follows by a further application of **MD3**, that

$$\mathbb{E} [\text{Hess } g_{\theta_0}(H_{\theta_0}(X_1), H_{\theta_0}(X_1))] = [\text{Hess } g : h \otimes h + \Sigma] (\theta_0) , \quad (3.36)$$

where $\Sigma(\theta_0)$ is defined in (3.10). Using that $\|H_K\|_{\theta_0} \leq K + \|h(\theta_0)\|_{\theta_0}$, and **MD3** in (3.34), we obtain that for any $\theta_0 \in \Theta$, $\mathbb{E}[|\mathcal{R}_{g,\eta}(\theta_0, X_1)|] < +\infty$. Then, by (3.33), (3.35) and (3.36), it follows from (3.32),

$$Q_{\eta}g(\theta_0) = g(\theta_0) + \eta \langle \text{grad } g(\theta_0), h(\theta_0) \rangle_{\theta_0} + (\eta^2/2) [\text{Hess } g : h \otimes h + \Sigma] (\theta_0) + \eta^2 \mathcal{R}_{g,\eta}(\theta_0)/6 ,$$

where we define $\mathcal{R}_{g,\eta}(\theta_0) = \eta^{-2} \mathbb{E}[\mathcal{R}_{g,\eta}(\theta_0, X_1)]$. The desired bound on the remainder in (3.29), is a simple consequence of (3.34).

(b) In addition to the results of (a) and specifically (3.29), we need to prove that, since g has compact support, there exists a compact set $K_K \subset \Theta$ such that $\|\nabla \text{Hess } g\|_{\gamma, \infty} \mathbb{1}_{A_{\theta_0}^c} = 0$ for any $\theta_0 \notin K_K$.

Using that $\|h(\theta)\|_{\theta} \leq C_3 \rho_{\Theta}(\theta^*, \theta)$, we obtain that on $A_{\theta_0}^c$, $\|H_{\theta_0}(X_1)\|_{\theta} \leq 2(C_3 \rho_{\Theta}(\theta^*, \theta) + K)$. In addition, by [Lee19, Corollary 6.12], $\rho_{\Theta}(\theta, \gamma(t)) = t\eta \|H_{\theta_0}(X_1)\|_{\theta}$ for any $t \in [0, 1]$, therefore for any $t \in [0, 1]$ and $\eta \in (0, (4C_3)^{-1})$

$$\rho_{\Theta}(\theta^*, \gamma(t)) \geq \rho_{\Theta}(\theta^*, \theta) - \rho_{\Theta}(\theta, \gamma(t)) \geq (1 - 2\eta t C_3) \rho_{\Theta}(\theta^*, \theta) - 2\eta K \geq \rho_{\Theta}(\theta^*, \theta)/2 - K/(2C_3) .$$

Consider now $R \geq 0$ such that for any $\theta \notin \bar{B}(\theta^*, R)$, $g(\theta) = 0$. Then, setting $K_K = \bar{B}(\theta^*, 2(R + K/(2C_3)))$, we obtain that for any $\theta_0 \notin K_K$ and $t \in [0, 1]$, $\gamma(t) \notin \bar{B}(\theta^*, R)$ and therefore, $\nabla \text{Hess } g_{\gamma(t)} = 0$, which yields $\|\nabla \text{Hess } g\|_{\gamma, \infty} \mathbb{1}_{A_{\theta_0}^c} = 0$ for any $\theta_0 \notin K_K$. Finally K_K is a compact subset of Θ by [Jos05, Theorem 1.7.1]. \square

3.9.1 Proof of Theorem 3.6

Let $g : \Theta \rightarrow \mathbb{R}$ be a smooth function. Since we assume that Θ is compact, g is smooth with compact support. Therefore, using Lemma 3.19-(a) for any $\theta \in \Theta$ and $\eta > 0$, we have,

$$Q_{\eta}g(\theta) = g(\theta) + \eta \langle \text{grad } g(\theta), h(\theta) \rangle_{\theta} + (\eta^2/2) [\text{Hess } g : \Sigma + h \otimes h] (\theta) + (\eta^2/6) \mathcal{R}_{g,\eta}(\theta) , \quad (3.37)$$

where using (3.29), Hölder inequality and MD3 gives,

$$|\mathcal{R}_{g,\eta}(\theta)| \leq 32\eta(\|h(\theta)\|_\theta^3 + K^3) \sup\{|\nabla \text{Hess } g_\theta(u, u, u)| : \theta \in \Theta, u \in U_\theta \Theta\} \\ + 16 \|\text{Hess } g\|_\infty \left(\sigma_0^2 + \sigma_1^2 \|h(\theta)\|_\theta^2 \right) .$$

Next, let $\eta \in (0, \bar{\eta}]$, where $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1}$. Note that since Θ is compact, g is smooth, h and Σ are continuous, all the functions appearing in (3.37) are bounded. Therefore, integrating (3.37) with respect to μ^η given by Theorem 3.2 and using that μ^η is invariant with respect to Q_η , we obtain,

$$- \int_\Theta \langle \text{grad } g(\theta), h(\theta) \rangle_\theta \mu^\eta(d\theta) = (\eta/2) \int_\Theta [\text{Hess } g : \Sigma + h \otimes h](\theta) \mu^\eta(d\theta) + (\eta/6) \int_\Theta \mathcal{R}_{g,\eta}(\theta) \mu^\eta(d\theta) .$$

Using that $\theta \mapsto [\text{Hess } g : \Sigma + h \otimes h](\theta)$ is bounded and continuous over Θ , Theorem 3.3-(b) and that $h(\theta^*) = 0$, by weak convergence of μ^η to δ_{θ^*} when $\eta \rightarrow 0$, we have,

$$\lim_{\eta \rightarrow 0} \int_\Theta [\text{Hess } g : \Sigma + h \otimes h](\theta) \mu^\eta(d\theta) = [\text{Hess } g : \Sigma + h \otimes h](\theta^*) = [\text{Hess } g : \Sigma](\theta^*) .$$

Equivalently, there exists $\mathcal{R}_{\text{Hess } g} : (0, \bar{\eta}] \rightarrow \mathbb{R}$ such that for any $\eta \in (0, \bar{\eta}]$, we have

$$\int_\Theta [\text{Hess } g : \Sigma + h \otimes h](\theta) \mu^\eta(d\theta) = [\text{Hess } g : \Sigma](\theta^*) + \mathcal{R}_{\text{Hess } g}(\eta) ,$$

where $\lim_{\eta \rightarrow 0} |\mathcal{R}_{\text{Hess } g}(\eta)| = 0$.

To conclude, we prove that $\limsup_{\eta \rightarrow 0} |\int_\Theta \mathcal{R}_{g,\eta}(\theta) \mu^\eta(d\theta)| = 0$. Let $K \geq 0$. By (3.29), since $\theta_0 \mapsto \mathbb{E}[\mathbb{1}_{A_{\theta_0}^c} \|H_K\|_{\theta_0}^3]$ is uniformly bounded over Θ by definition (3.30) and since h is continuous, we have that

$$\limsup_{\eta \rightarrow 0} \left| \int_\Theta \mathcal{R}_{g,\eta}(\theta) \mu^\eta(d\theta) \right| \leq 16 \|\text{Hess } g\|_\infty \limsup_{\eta \rightarrow 0} \int_\Theta \mathbb{E} \left[\|e_\theta(X_1)\|_\theta^2 \{1 - \chi_K(\theta)\} \right] \mu^\eta(d\theta) \\ \leq 16 \|\text{Hess } g\|_\infty \mathbb{E} \left[\|e_{\theta^*}(X_1)\|_{\theta^*}^2 \{1 - \chi_K(\theta^*)\} \right] ,$$

using Theorem 3.3-(b), that $\theta \mapsto \mathbb{E}[\|e_\theta(X_1)\|_\theta^2]$ and χ_K are continuous and bounded by MD5 since $\mathbb{E}[\|e_\theta(X_1)\|_\theta^2] = \text{Tr}(\Sigma(\theta))$ for any $\theta \in \Theta$ and Θ is compact. Taking $K \rightarrow +\infty$ completes the proof.

3.9.2 Proof of Theorem 3.7

We introduce an auxiliary chain $(U_n)_{n \in \mathbb{N}}$ as an intermediate step between $(\theta_n)_{n \in \mathbb{N}}$ and $(\bar{U}_n)_{n \in \mathbb{N}}$ for which we recall the definition below. Define for any $\eta > 0, n \in \mathbb{N}$,

$$U_n = \text{Exp}_{\theta^*}^{-1}(\theta_n) \quad \text{and} \quad \bar{U}_n = \eta^{-1/2} \text{Exp}_{\theta^*}^{-1}(\theta_n) = \eta^{-1/2} U_n , \quad (3.38)$$

where $(\theta_n)_{n \in \mathbb{N}}$ is defined by (3.2) with $S = \Theta$ i.e. $\text{proj}_S = \text{Id}$. Note that $(U_n)_{n \in \mathbb{N}}$ and $(\bar{U}_n)_{n \in \mathbb{N}}$ are Markov chains with state space $\text{T}_{\theta^*} \Theta$, as Exp_{θ^*} is a bijection. Conversely,

since $\text{Exp}_{\theta^*}^{-1}$ and $\eta^{-1/2}\text{Exp}_{\theta^*}^{-1}$ are bijections from Θ to $\mathsf{T}_{\theta^*}\Theta$ under **A 6-(i)**, $(\theta_n)_{n \in \mathbb{N}}$ is a deterministic function of $(U_n)_{n \in \mathbb{N}}$ or $(\bar{U}_n)_{n \in \mathbb{N}}$. Therefore, the convergence of these three processes is expected to be the same. This is the content of the following result. Denote by R_η and \bar{R}_η the Markov kernels on $\mathsf{T}_{\theta^*}\Theta \times \mathcal{B}(\mathsf{T}_{\theta^*}\Theta)$, associated with $(U_n)_{n \in \mathbb{N}}$ and $(\bar{U}_n)_{n \in \mathbb{N}}$ respectively.

Lemma 3.20. *Assume **A 6-(i)-(ii)**, **MD3**, **MD4**, **A6**, **A7** and **A8(K^{*})** for some compact set $\mathsf{K}^* \subset \mathsf{S}$. Let $\eta \in (0, \bar{\eta}]$ where $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1}$. For any measurable and bounded function $g : \mathsf{T}_{\theta^*}\Theta \rightarrow \mathbb{R}$ and any $u_0, \bar{u}_0 \in \mathsf{T}_{\theta^*}\Theta$, R_η and \bar{R}_η satisfy*

$$R_\eta g(u_0) = Q_\eta g(\text{Exp}_{\theta^*}(u_0)) \quad \text{and} \quad \bar{R}_\eta g(\bar{u}_0) = R_\eta g_\eta(\eta^{1/2}\bar{u}_0), \quad (3.39)$$

where $g : \theta \mapsto g[\text{Exp}_{\theta^*}^{-1}(\theta)]$ and $g_\eta : u \mapsto g(\eta^{-1/2}u)$ are defined over Θ and $\mathsf{T}_{\theta^*}\Theta$ respectively, and Q_η is the Markov kernel associated with $(\theta_n)_{n \in \mathbb{N}}$. In addition, R_η and \bar{R}_η both admit a unique stationary distribution ν^η and $\bar{\nu}^\eta$ respectively, defined for any $A \in \mathcal{B}(\mathsf{T}_{\theta^*}\Theta)$ by

$$\nu^\eta(A) = \mu^\eta(\text{Exp}_{\theta^*}(A)) \quad \text{and} \quad \bar{\nu}^\eta(A) = \nu^\eta(\eta^{1/2}A). \quad (3.40)$$

Finally, both R_η and \bar{R}_η are Harris-recurrent and geometrically ergodic, i.e. there exist $C, \bar{C} : \mathsf{T}_{\theta^*}\Theta \rightarrow \mathbb{R}$ and $\rho, \bar{\rho} \in \mathbb{R}_+^*$ such that for any $u, \bar{u} \in \mathsf{T}_{\theta^*}\Theta$,

$$\|\delta_u R_\eta - \nu^\eta\|_{\text{TV}} \leq C(u)\rho^n \quad \text{and} \quad \|\delta_{\bar{u}} \bar{R}_\eta - \bar{\nu}^\eta\|_{\text{TV}} \leq \bar{C}(\bar{u})\bar{\rho}^n.$$

Proof. Let $g : \mathsf{T}_{\theta^*}\Theta \rightarrow \mathbb{R}$ be a measurable and bounded function and $u_0 \in \mathsf{T}_{\theta^*}\Theta$. Consider $(U_n)_{n \in \mathbb{N}}$ defined by (3.38) with $\theta_0 = \text{Exp}_{\theta^*}(u_0)$. Using (3.38), we have by definition

$$\mathbb{E}[g(U_1)] = \mathbb{E}\left[g\left(\text{Exp}_{\theta^*}^{-1}(\theta_1)\right)\right] = Q_\eta\left(g \circ \text{Exp}_{\theta^*}^{-1}\right)\left(\text{Exp}_{\theta^*}(u_0)\right).$$

Moreover, let $\bar{u}_0 \in \mathsf{T}_{\theta^*}\Theta$ and consider $(\bar{U}_n)_{n \in \mathbb{N}}$ defined by (3.38) with $U_0 = \eta^{1/2}\bar{u}_0$. Using (3.38), we have by definition

$$\mathbb{E}[g(\bar{U}_1)] = \mathbb{E}\left[g\left(\eta^{-1/2}U_1\right)\right] = R_\eta g_\eta\left(\eta^{1/2}\bar{u}_0\right),$$

where $g_\eta : u \mapsto g(\eta^{-1/2}u)$ is defined over $\mathsf{T}_{\theta^*}\Theta$, therefore proving (3.39).

We show that ν^η and $\bar{\nu}^\eta$ are invariant for R_η and \bar{R}_η respectively. Indeed, for any $A \in \mathcal{B}(\mathsf{T}_{\theta^*}\Theta)$, we have by (3.38), (3.39) and (3.40)

$$\begin{aligned} \nu^\eta R_\eta(A) &= \int_{\mathsf{T}_{\theta^*}\Theta} d\nu^\eta(u) R_\eta(u, A) = \int_{\Theta} d\mu^\eta(\theta) R_\eta\left(\text{Exp}_{\theta^*}^{-1}(\theta), A\right) \\ &= \int_{\Theta} d\mu^\eta(\theta) Q_\eta(\theta, \text{Exp}_{\theta^*}(A)) = \mu^\eta(\text{Exp}_{\theta^*}(A)) = \nu^\eta(A). \end{aligned}$$

Therefore ν^η is invariant for R_η . Similarly, we show that $\bar{\nu}^\eta$ is invariant for \bar{R}_η . Using again (3.38), (3.39) and (3.40), for any $A \in \mathcal{B}(\mathsf{T}_{\theta^*}\Theta)$ we have,

$$\bar{\nu}^\eta \bar{R}_\eta(A) = \int_{\mathsf{T}_{\theta^*}\Theta} d\bar{\nu}^\eta(u) \bar{R}_\eta\left(\eta^{-1/2}u, A\right) = \int_{\mathsf{T}_{\theta^*}\Theta} d\nu^\eta(u) R_\eta\left(u, \eta^{1/2}A\right) = \bar{\nu}^\eta(A).$$

Finally, since $(\theta_n)_{n \in \mathbb{N}}$, $(U_n)_{n \in \mathbb{N}}$ and $(\bar{U}_n)_{n \in \mathbb{N}}$ are deterministic functions of each other and since Theorem 3.2 proves that $(\theta_n)_{n \in \mathbb{N}}$ is geometrically ergodic and Harris-recurrent, the same holds for $(U_n)_{n \in \mathbb{N}}$ and $(\bar{U}_n)_{n \in \mathbb{N}}$ and their invariant distributions are unique. \square

3.9. Proofs of Section 3.3

For any smooth function with compact support $g : T_{\theta^*}\Theta \rightarrow \mathbb{R}$, $\bar{u}_0 \in T_{\theta^*}\Theta$ and $\eta > 0$ consider the 2-tensor $(C^2(g, \bar{u}_0, \eta)_{ij})_{i,j \in \{1, \dots, d\}}$ defined by, for any $i, j \in \{1, \dots, d\}$,

$$C^2(g, \bar{u}_0, \eta)_{ij} = \partial_{ij}^2 g(\bar{u}_0) - \eta^{1/2} \sum_{k=1}^d \Gamma_{ij}^k \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \partial_k g(\bar{u}_0), \quad (3.41)$$

and, similarly consider the 3-tensor $(C^3(g, \bar{u}_0, \eta)_{ijk})_{i,j,k \in \{1, \dots, d\}}$ defined by, for any $i, j, k \in \{1, \dots, d\}$,

$$\begin{aligned} C^3(g, \bar{u}_0, \eta)_{ijk} &= \partial_{ijk}^3 g(\bar{u}_0) \\ &\quad - \eta^{1/2} \sum_{l=1}^d \left[\Gamma_{ij}^l \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \partial_{kl}^2 g(\bar{u}_0) + \Gamma_{ki}^l \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \partial_{jl}^2 g(\bar{u}_0) \right. \\ &\quad \left. + \Gamma_{kj}^l \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \partial_{il}^2 g(\bar{u}_0) \right] - \eta \sum_{m=1}^d \partial_k \Gamma_{ij}^m \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \partial_m g(\bar{u}_0) \\ &\quad + \eta \sum_{l,m=1}^d \left[\Gamma_{kj}^l \Gamma_{il}^m + \Gamma_{ki}^l \Gamma_{lj}^m \right] \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \partial_m g(\bar{u}_0), \end{aligned} \quad (3.42)$$

where $(\Gamma_{ij}^k)_{i,j,k \in \{1, \dots, d\}}$ are the Christoffel symbols of the Levi-Civita connection ∇ . We derive the following Taylor formulas.

Lemma 3.21. *Assume **A 6-(i)-(ii)**, **MD 3**, **MD 4**, **A 6**, **A 7** and **A 8(K^{*})** for some compact set $K^* \subset S$. Suppose in addition that there exists $C_3 > 0$ such that for any $\theta \in \Theta$, $\|h(\theta)\|_{\theta} \leq C_3 \rho_{\Theta}(\theta^*, \theta)$ and let $\bar{\eta} = [2C_2 L(1 + \sigma_1^2)]^{-1} \wedge (4C_3)^{-1}$. Consider normal coordinates $(u^i)_{i \in \{1, \dots, d\}}$ centered at θ^* and define for any $i, j \in \{1, \dots, d\}$, $h^i : \Theta \rightarrow \mathbb{R}$, $\Sigma_{ij} : \Theta \rightarrow \mathbb{R}$ by $h^i = du^i(h)$ and $\Sigma_{ij} = [du^i \otimes du^j]\{\Sigma\}$. For any smooth function with compact support $g : T_{\theta^*}\Theta \rightarrow \mathbb{R}$, any $\eta \in (0, \bar{\eta}]$ and $\bar{u}_0 \in T_{\theta^*}\Theta$, we have*

$$\begin{aligned} \bar{R}_{\eta} g(\bar{u}_0) &= g(\bar{u}_0) + \eta^{1/2} \sum_{i=1}^d \partial_i g(\bar{u}_0) h^i \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \\ &\quad + \frac{\eta}{2} \sum_{i,j=1}^d \left\{ \partial_{ij}^2 g(\bar{u}_0) - \eta^{1/2} \sum_{k=1}^d \Gamma_{ij}^k \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \partial_k g(\bar{u}_0) \right\} \left[\Sigma_{ij} + h^i h^j \right] \left(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0) \right) \\ &\quad + (\eta/6) \bar{\mathcal{R}}_{g,\eta}(\bar{u}_0), \end{aligned} \quad (3.43)$$

where, setting $\theta_0 = \text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0)$,

$$\left| \bar{\mathcal{R}}_{g,\eta}(\bar{u}_0) \right| \leq 8\eta^{1/2} \mathbb{1}_{K_K}(\theta_0) \mathbb{E} \left[\|C^3(g, \eta)\|_{\gamma} \mathbb{1}_{A_{\theta_0}^c} \|H_K\|_{\theta_0}^3 \right] + 16 \|C^2(g, \eta)\| \mathbb{E} \left[\|Y_K\|_{\theta_0}^2 \right], \quad (3.44)$$

using the definitions of $H_K, Y_K, A_{\theta_0}, K_K$ and γ in Lemma 3.19-(3.30),

$$\begin{aligned} \|C^2(g, \eta)\| &= \sup\{ \|C^2(g, \bar{u}, \eta)[v^{\otimes 2}]\| : \bar{u} \in T_{\theta^*}\Theta, v \in \mathbb{R}^d, \|v\|_2 = 1 \} \\ \|C^3(g, \eta)\|_{\gamma} &= \sup\{ \|C^3(g, \bar{u}, \eta)[v^{\otimes 3}]\| : \bar{u} \in \eta^{-1/2} \text{Exp}_{\theta^*}^{-1}(\gamma([0, 1])), v \in \mathbb{R}^d, \|v\|_2 = 1 \}, \end{aligned} \quad (3.45)$$

where $C^2(g, \bar{u}, \eta)$ and $C^3(g, \bar{u}, \eta)$ are defined in (3.41) and (3.42).

Proof. Using **A 6-(i)** and [Lee19, Proposition 12.9], $(u^i)_{i \in \{1, \dots, d\}}$ are global coordinates on the Hadamard manifold Θ . Let $g : \mathbb{T}_{\theta^*} \Theta \rightarrow \mathbb{R}$ be a smooth function with compact support and $g : \Theta \rightarrow \mathbb{R}$ defined for any $\theta \in \Theta$ by $g(\theta) = g(\text{Exp}_{\theta^*}^{-1}(\theta))$. Note that since $\|\text{Exp}_{\theta^*}^{-1}(\theta)\|_{\theta^*} = \rho_{\Theta}(\theta^*, \theta)$, for any $\theta \in \Theta$ by [Lee19, Corollary 6.12], g is a smooth function with compact support as well. In addition, by definition of the normal coordinates, $g : u \mapsto g(\text{Exp}_{\theta^*}(u))$ is the expression of g in this coordinate system. Using this fact and the definitions of the Riemannian gradient and Hessian [Lee19, Equation 2.14, Example 4.22], we have, for any $\theta_0 \in \Theta$,

$$\begin{aligned} \text{grad } g(\theta_0) &= \sum_{i=1}^d \partial_i g(u_0) \partial u_i, \\ \text{Hess } g(\theta_0) &= \sum_{i,j=1}^d \left\{ \partial_{ij}^2 g(u_0) - \sum_{k=1}^d \Gamma_{ij}^k(\text{Exp}_{\theta^*}(u_0)) \partial_k g(u_0) \right\} du^i \otimes du^j, \end{aligned} \quad (3.46)$$

where $u_0 = \text{Exp}_{\theta^*}^{-1}(\theta_0)$ and $(\Gamma_{ij}^k)_{i,j,k \in \{1, \dots, d\}}$ are the Christoffel symbols. Combining these expressions with Lemma 3.20-(3.39) and Lemma 3.19-(b)-(3.28) gives

$$\begin{aligned} R_{\eta} g(u) &= g(u_0) + \eta \sum_{i=1}^d \partial_i g(u_0) h^i(\text{Exp}_{\theta^*}(u_0)) \\ &+ (\eta^2/2) \sum_{i,j=1}^d \left\{ \partial_{ij}^2 g(u_0) - \sum_{k=1}^d \Gamma_{ij}^k(\text{Exp}_{\theta^*}(u_0)) \partial_k g(u_0) \right\} \left[\Sigma_{ij}(\text{Exp}_{\theta^*}(u_0)) + h^i h^j(\text{Exp}_{\theta^*}(u_0)) \right] \\ &\quad + (\eta^2/6) \tilde{\mathcal{R}}_{g,\eta}(u_0), \end{aligned}$$

where $\tilde{\mathcal{R}}_{g,\eta}(u_0) = \mathcal{R}_{g,\eta}(\theta_0)$ is bounded using (3.31), for $\theta_0 = \text{Exp}_{\theta^*}(u_0)$ and $g : \theta \mapsto g(\text{Exp}_{\theta^*}^{-1}(\theta))$.

Replacing g with $g_{\eta} : u \mapsto g(\eta^{-1/2}u)$ defined over $\mathbb{T}_{\theta^*} \Theta$ and using that for any $i, j \in \{1, \dots, d\}$ and $u_0 \in \mathbb{T}_{\theta^*} \Theta$,

$$\partial_i g_{\eta}(u_0) = \eta^{-1/2} \partial_i g(\eta^{-1/2}u_0) \quad \text{and} \quad \partial_{ij}^2 g_{\eta}(u_0) = \eta^{-1} \partial_{ij}^2 g(\eta^{-1/2}u_0), \quad (3.47)$$

we have for any $u_0 \in \mathbb{T}_{\theta^*} \Theta$,

$$\begin{aligned} R_{\eta} g_{\eta}(u_0) &= g(\eta^{-1/2}u_0) + \eta^{1/2} \sum_{i=1}^d \partial_i g(\eta^{-1/2}u_0) h^i(\text{Exp}_{\theta^*}(u_0)) \\ &+ (\eta/2) \sum_{i,j=1}^d \left\{ \partial_{ij}^2 g\left(\frac{u_0}{\eta^{1/2}}\right) - \eta^{1/2} \sum_{k=1}^d \Gamma_{ij}^k(\text{Exp}_{\theta^*}(u_0)) \partial_k g\left(\frac{u_0}{\eta^{1/2}}\right) \right\} \left[\Sigma_{ij} + h^i h^j \right](\text{Exp}_{\theta^*}(u_0)) \\ &\quad + (\eta^2/6) \tilde{\mathcal{R}}_{g_{\eta},\eta}(u_0). \end{aligned} \quad (3.48)$$

Expressing $\tilde{\mathcal{R}}_{g_{\eta},\eta}(u_0)$ using partial derivatives shows explicitly the dependency on η . Using (3.47) and the equivalent formula for the third order derivative, we have for any $K > 0$,

$$\eta^2 \left| \tilde{\mathcal{R}}_{g_{\eta},\eta}(u_0) \right| \leq 8\eta^3 \mathbb{1}_{K_K}(\theta_0) \mathbb{E} \left[\|\nabla \text{Hess } g_{\eta}\|_{\gamma, \infty} \mathbb{1}_{A_{\theta_0}^c} \|H_K\|_{\theta_0}^3 \right] + 16\eta^2 \|\text{Hess } g_{\eta}\|_{\infty} \mathbb{E} \left[\|Y_K\|_{\theta_0}^2 \right], \quad (3.49)$$

3.9. Proofs of Section 3.3

where $\theta_0 = \text{Exp}_{\theta^*}(u_0)$, $\gamma : [0, 1] \rightarrow \Theta$ is defined by $\gamma(t) = \text{Exp}_{\theta_0}(t\eta H_{\theta_0}(X_1))$, H_K, Y_K and \mathbf{A}_{θ_0} are defined in (3.30). Using (3.46) and Proposition 3.30, we have $\text{Hess } g_\eta(u) = \eta^{-1}\mathbf{C}^2(g, \eta^{-1/2}u_0, \eta)$ and $\nabla \text{Hess } g_\eta(u) = \eta^{-3/2}\mathbf{C}^3(g, \eta^{-1/2}u_0, \eta)$, where \mathbf{C}^2 and \mathbf{C}^3 are defined in (3.41) and (3.42) respectively. This gives

$$\|\nabla \text{Hess } g_\eta\|_{\gamma, \infty} = \eta^{-3/2}\|\mathbf{C}^3(g, \eta)\|_\gamma \quad \text{and} \quad \|\text{Hess } g_\eta\|_\infty = \eta^{-1}\|\mathbf{C}^2(g, \eta)\|, \quad (3.50)$$

where $\|\mathbf{C}^2(g, \eta)\|$ and $\|\mathbf{C}^3(g, \eta)\|_\gamma$ are defined in (3.45). Setting $u_0 = \eta^{1/2}\bar{u}_0$ in (3.48), we get

$$\begin{aligned} R_\eta g_\eta(\eta^{1/2}\bar{u}_0) &= g(\bar{u}_0) + \eta^{1/2} \sum_{i=1}^d \partial_i g(\bar{u}_0) h^i \left(\text{Exp}_{\theta^*}(\eta^{1/2}\bar{u}_0) \right) \\ &+ (\eta/2) \sum_{i,j=1}^d \left\{ \partial_{ij}^2 g(\bar{u}_0) - \eta^{1/2} \sum_{k=1}^d \Gamma_{ij}^k(\text{Exp}_{\theta^*}(\eta^{1/2}\bar{u}_0)) \partial_k g(\bar{u}_0) \right\} \left[\Sigma_{ij} + h^i h^j \right] \left(\text{Exp}_{\theta^*}(\eta^{1/2}\bar{u}_0) \right) \\ &\quad + \eta^2 \tilde{\mathcal{R}}_{g_\eta, \eta}(\eta^{1/2}\bar{u}_0). \end{aligned} \quad (3.51)$$

Therefore, letting $\bar{\mathcal{R}}_{g, \eta}(\bar{u}_0) = \eta \tilde{\mathcal{R}}_{g_\eta, \eta}(\eta^{1/2}\bar{u}_0)$, and combining Lemma 3.20-(3.39), (3.49), (3.50) and (3.51) gives the desired result. \square

Lemma 3.22. *Assume **A 6-(i)-(ii)** and **A10**. Consider normal coordinates $(u^i)_{i \in \{1, \dots, d\}}$ centered at θ^* with respect to the orthonormal basis $(\mathbf{e}_i)_{i \in \{1, \dots, d\}}$ of $\mathbf{T}_{\theta^*}\Theta$. Then h can be expressed in this chart as, for any $\eta > 0$, $\bar{u} \in \mathbf{T}_{\theta^*}\Theta$,*

$$h \left(\text{Exp}_{\theta^*}(\eta^{1/2}\bar{u}) \right) = \sum_{i=1}^d \left\{ \eta^{1/2} \sum_{k=1}^d \mathbf{A}_k^i \bar{u}^k + \mathcal{R}_h^i(\eta^{1/2}\bar{u}) \right\} \partial u_i, \quad (3.52)$$

where \mathbf{A} is defined in **A10**, \bar{u}^k are the components of \bar{u} in $(\mathbf{e}_i)_{i \in \{1, \dots, d\}}$ and for any $i \in \{1, \dots, d\}$, $\lim_{u \rightarrow 0} \{ |\mathcal{R}_h^i(u)| / \|u\|_{\theta^*} \} = 0$.

Proof. Since Θ is a Hadamard manifold, these normal coordinates are defined throughout Θ . Thus, for any $\theta \in \Theta$, it is possible to write,

$$h(\theta) = \sum_{j=1}^d h^j(\theta) \partial u_j(\theta). \quad (3.53)$$

Recall the definition of the metric coefficients in the coordinates $(u^i)_{i \in \{1, \dots, d\}}$ at $\theta \in \Theta$, for any $i, j \in \{1, \dots, d\}$,

$$\mathbf{g}_{ij}(\theta) = \langle \partial u_i(\theta), \partial u_j(\theta) \rangle_\theta. \quad (3.54)$$

Then, taking the scalar product of (3.53) with each ∂u_i , we have for any $i \in \{1, \dots, d\}$,

$$\sum_{j=1}^d \mathbf{g}_{ij}(\theta) h^j(\theta) = \langle h(\theta), \partial u_i(\theta) \rangle_\theta. \quad (3.55)$$

From the Taylor expansion formula for vector fields given by Theorem 3.29 for the geodesic $\gamma : [0, 1] \rightarrow \Theta$ given by $\gamma(0) = \theta^*$ and $\dot{\gamma}(0) = \text{Exp}_{\theta^*}^{-1}(\theta)$, it follows that,

$$\partial u_i(\theta) = T_{01}^Y \left[\mathbf{e}_i + \nabla(\partial u_i)_{\theta^*} \left(\text{Exp}_{\theta^*}^{-1}(\theta) \right) \right] + \mathcal{R}_{\partial u_i}(\theta), \quad (3.56)$$

where the remainder is given by

$$\mathcal{R}_{\partial u_i}(\theta) = \int_0^1 (1-t) T_{t1}^Y \nabla^2(\partial u_i)_{\gamma(t)}(\dot{\gamma}(t), \dot{\gamma}(t)) dt.$$

Let $\|\nabla^2 \partial u_i\|_{\infty, \gamma} = \sup\{|\nabla^2(\partial u_i)_{\gamma(t)}(v, v)| : t \in [0, 1], v \in U_{\gamma(t)}\Theta\}$ which is finite as $\gamma[0, 1]$ is compact. Then using that for any $t \in [0, 1]$, $\|\dot{\gamma}(t)\|_{\gamma(t)} = \rho_{\Theta}(\theta^*, \theta)$ by [Lee19, Corollary 5.6] and that geodesics are length-minimizing curves by A 6-(i); and that the parallel transport map is an isometry [Lee19, p.108], we have

$$|\mathcal{R}_{\partial u_i}(\theta)| \leq (1/2) \|\nabla^2 \partial u_i\|_{\infty, \gamma} \rho_{\Theta}^2(\theta^*, \theta).$$

This proves that $\lim_{\theta \rightarrow \theta^*} |\mathcal{R}_{\partial u_i}(\theta)/\rho_{\Theta}(\theta^*, \theta)| = 0$. By the definition of normal coordinates centered at θ^* , for any $i, j \in \{1, \dots, d\}$, $\nabla_{\partial u_j} \partial u_i = \sum_{k=1}^d \Gamma_{ji}^k \partial u_k$ and $(\Gamma_{ji}^k)_{i,j,k \in \{1, \dots, d\}}$ vanishes at θ^* [Lee19, Proposition 5.24] so (3.56) becomes

$$\partial u_i(\theta) = T_{01}^Y(\mathbf{e}_i) + \mathcal{R}_{\partial u_i}(\theta). \quad (3.57)$$

Taking the scalar product of (3.12) and (3.57), it follows that

$$\langle h(\theta), \partial u_i(\theta) \rangle_{\theta} = \langle \mathbf{A} \text{Exp}_{\theta^*}^{-1}(\theta), \mathbf{e}_i \rangle_{\theta^*} + \tilde{\mathcal{R}}_h^i(\theta), \quad (3.58)$$

since parallel transport preserves scalar products, where $\lim_{\theta \rightarrow \theta^*} \{|\tilde{\mathcal{R}}_h^i(\theta)|/\rho_{\Theta}(\theta^*, \theta)\} = 0$. On the other hand, from (3.54) and (3.57), since the $(\mathbf{e}_i)_{i \in \{1, \dots, d\}}$ are orthonormal,

$$\mathfrak{g}_{ij}(\theta) = \delta_{ij} + \mathcal{R}_{\mathfrak{g}}^{ij}(\theta), \quad (3.59)$$

where $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise and $\lim_{\theta \rightarrow \theta^*} \{|\mathcal{R}_{\mathfrak{g}}^{ij}(\theta)|/\rho_{\Theta}(\theta^*, \theta)\} = 0$. Plugging (3.58) and (3.59) in (3.55), we obtain

$$h^i(\theta) = \sum_{j=1}^d \mathbf{A}_j^i u^j(\theta) + \mathcal{R}_h^i(\theta), \quad (3.60)$$

where $\lim_{\theta \rightarrow \theta^*} |\mathcal{R}_h^i(\theta)| = 0$. Finally, (3.52) is obtained from (3.53)-(3.60), by setting $\theta = \text{Exp}_{\theta^*}(\eta^{1/2} \bar{u})$, for $\bar{u} \in T_{\theta^*}\Theta$, and noting that

$$\begin{aligned} u^j(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u})) &= \langle \text{Exp}_{\theta^*}^{-1}(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u})), \mathbf{e}_j \rangle_{\theta^*} = \eta^{1/2} \bar{u}^j, \\ \rho_{\Theta}(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}), \theta^*) &= \eta^{1/2} \|\bar{u}\|_{\theta^*}, \end{aligned}$$

which follow from [Lee19, Corollary 5.6] and the definition of the coordinates $(u^i)_{i \in \{1, \dots, d\}}$. \square

Lemma 3.23. *Assume **A 6-(i)-(ii)**, **MD3**, **MD4**, **MD5**, **MD6**, **A6**, **A7**, **A10** and **A11** hold. Let $\bar{\eta} = [2C_2L(1 + \sigma_1^2)]^{-1} \wedge (4C_3)^{-1}$. Then the family of distributions $(\bar{\nu}^\eta)_{\eta \in (0, \bar{\eta}]}$, defined by (3.11), is tight.*

Proof. For any $\eta \in (0, \bar{\eta}]$, the conditions of Lemma 3.20 hold, thus the Markov chain $(\bar{U}_n)_{n \in \mathbb{N}}$ is ergodic and its invariant distribution $\bar{\nu}^\eta$ is given by (3.11). For any $r \geq 0$, let $\bar{\mathbb{B}}_r = \{u \in \mathbb{T}_{\theta^*}\Theta : \|u\|_{\theta^*} \leq r\}$ be the tangent closed ball at θ^* of center 0 and radius r . Then, by (3.40) and [Lee19, Corollary 6.13], for any $r > 0$ and $\eta \in (0, \bar{\eta}]$, we have

$$\bar{\nu}^\eta \left(\mathbb{T}_{\theta^*}\Theta \setminus \bar{\mathbb{B}}_r \right) = \nu^\eta \left(\mathbb{T}_{\theta^*}\Theta \setminus \bar{\mathbb{B}}_{\eta^{1/2}r} \right) = \mu^\eta \left(\Theta \setminus \bar{\mathbb{B}}(\theta^*, \eta^{1/2}r) \right). \quad (3.61)$$

However, by **A11**,

$$\begin{aligned} \mu^\eta \left(\Theta \setminus \bar{\mathbb{B}}(\theta^*, \eta^{1/2}r) \right) &\leq \phi^{-1}(\eta^{1/2}r) \int_{\Theta \setminus \{\theta^*\}} \phi(\rho_\Theta(\theta^*, \theta)) d\mu^\eta(\theta) \\ &\leq \phi^{-1}(\eta^{1/2}r) \int_{\Theta \setminus \{\theta^*\}} V(\theta) d\mu^\eta(\theta). \end{aligned} \quad (3.62)$$

Now, using **A11** and Lemma 3.18 taking $\mathbb{K}^* = \{\theta^*\}$, we have,

$$\int_{\Theta \setminus \{\theta^*\}} V(\theta) d\mu^\eta(\theta) \leq 2\eta L \left\{ \sigma_0^2 + C_1(1 + \sigma_1^2) \right\} / \lambda.$$

Combining this result and (3.62) in (3.61) implies that for any $r > 0$,

$$\begin{aligned} \bar{\nu}^\eta \left(\mathbb{T}_{\theta^*}\Theta \setminus \bar{\mathbb{B}}_r \right) &\leq 2\eta L \left\{ \sigma_0^2 + C_1(1 + \sigma_1^2) \right\} / [\lambda \phi(\eta^{1/2}r)] \\ &\leq \sup_{\eta \leq \bar{\eta}} \{ \eta / \phi(\eta^{1/2}r) \} (2L/\lambda) \left\{ \sigma_0^2 + C_1(1 + \sigma_1^2) \right\}, \end{aligned}$$

where $\lim_{r \rightarrow +\infty} \{ \sup_{\eta \leq \bar{\eta}} \eta / \phi(\eta^{1/2}r) \} = 0$ using **A11**. Therefore, for any $\varepsilon > 0$, there exists $r > 0$ such that for any $\eta \in (0, \bar{\eta}]$, $\bar{\nu}^\eta(\mathbb{T}_{\theta^*}\Theta \setminus \bar{\mathbb{B}}_r) \leq \varepsilon$. This concludes the proof that $(\bar{\nu}^\eta)_{\eta \in (0, \bar{\eta}]}$ is tight. \square

Proof of Theorem 3.7. Consider normal coordinates $(u^i)_{i \in \{1, \dots, d\}}$ centered at θ^* with respect to the orthonormal basis $(\mathbf{e}_i)_{i \in \{1, \dots, d\}}$ of $\mathbb{T}_{\theta^*}\Theta$. Define for any $i, j \in \{1, \dots, d\}$, $h^i : \Theta \rightarrow \mathbb{R}$, $\Sigma_{ij} : \Theta \rightarrow \mathbb{R}$ by $h^i = du^i(h)$ and $\Sigma_{ij} = [du^i \otimes du^j]\{\Sigma\}$. Let $g : \mathbb{T}_{\theta^*}\Theta \rightarrow \mathbb{R}$ be a smooth function with compact support. Applying Lemma 3.21 to g gives (3.43). Using **MD5**, Σ is continuous, which implies that for any $\bar{u}_0 \in \mathbb{T}_{\theta^*}\Theta$,

$$\Sigma \left(\text{Exp}_{\theta^*}(\eta^{1/2}\bar{u}_0) \right) = \sum_{i,j=1}^d \left\{ \Sigma_{ij}^* + \mathcal{R}_{\Sigma}^{ij} \left(\eta^{1/2}\bar{u}_0 \right) \right\} \partial u_i \otimes \partial u_j, \quad (3.63)$$

where for any $i, j \in \{1, \dots, d\}$, $\Sigma_{ij}^* = \Sigma_{ij}(\theta^*)$, $\mathcal{R}_{\Sigma}^{ij}$ is continuous over $\mathbb{T}_{\theta^*}\Theta$ and $\mathcal{R}_{\Sigma}^{ij}(0) = 0$. Using Lemma 3.22, replacing Σ_{ij} and h^i in (3.43) with (3.52) and (3.63) gives for any $\bar{u}_0 \in \mathbb{T}_{\theta^*}\Theta$,

$$\begin{aligned} \bar{R}_{\eta} g(\bar{u}_0) &= g(\bar{u}_0) + \eta \sum_{i=1}^d \partial_i g(\bar{u}_0) \sum_{k=1}^d \mathbf{A}_k^i \bar{u}_0^k + (\eta/2) \sum_{i,j=1}^d \partial_{ij}^2 g(\bar{u}_0) \Sigma_{ij}^* + \eta \mathcal{R}_{g, \eta, \Sigma, h}(\bar{u}_0) \\ &\quad + (\eta/6) \bar{\mathcal{R}}_{g, \eta}(\bar{u}_0), \end{aligned} \quad (3.64)$$

where \bar{u}_0^k are the components of \bar{u}_0 in $(\mathbf{e}_i)_{i \in \{1, \dots, d\}}$,

$$\begin{aligned} \mathcal{R}_{g, \eta, \Sigma, h}(\bar{u}_0) &= \eta^{-1/2} \sum_{i=1}^d \mathcal{R}_h^i(\eta^{1/2} \bar{u}_0) \partial_i g(\bar{u}_0) \\ &+ (1/2) \sum_{i, j=1}^d \left\{ \partial_{ij}^2 g(\bar{u}_0) - \eta^{1/2} \sum_{k=1}^d \Gamma_{ij}^k(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0)) \partial_k g(\bar{u}_0) \right\} \left[\mathcal{R}_{\Sigma}^{ij}(\eta^{1/2} \bar{u}_0) \right] \\ &+ (1/2) \sum_{i, j=1}^d \left\{ \partial_{ij}^2 g(\bar{u}_0) - \eta^{1/2} \sum_{k=1}^d \Gamma_{ij}^k(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0)) \partial_k g(\bar{u}_0) \right\} \left[h^i h^j (\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0)) \right] \\ &\quad - (\eta^{1/2}/2) \sum_{i, j, k=1}^d \Gamma_{ij}^k(\text{Exp}_{\theta^*}(\eta^{1/2} \bar{u}_0)) \partial_k g(\bar{u}_0) \Sigma_{\star}^{ij}. \end{aligned}$$

By Lemma 3.23, $(\bar{\nu}^n)_{n \in (0, \bar{\eta}]}$ is tight and therefore relatively compact. Therefore, it is enough that for any limit point $\bar{\nu}^*$, $\bar{\nu}^* = \mathbf{N}(0, \mathbf{V})$ where $\mathbf{V} \in \mathbb{R}^{d \times d}$ is the solution of the Lyapunov equation $\mathbf{A}\mathbf{V} + \mathbf{V}\mathbf{A}^\top = \Sigma(\theta^*)$. Let $(\eta_n)_{n \in \mathbb{N}^*}$ be a sequence with values in $(0, \bar{\eta}]$, such that $\lim_{n \rightarrow +\infty} \eta_n = 0$, and $(\bar{\nu}^{\eta_n})_{n \in \mathbb{N}^*}$ weakly converges to $\bar{\nu}^*$.

First by (3.64), we have

$$\begin{aligned} &\int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) \int_{\mathbb{T}_{\theta^* \Theta}} \bar{R}_{\eta_n}(\bar{u}_0, \mathrm{d}\bar{u}_1) g(\bar{u}_1) \\ &= \int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) g(\bar{u}_0) + \eta_n \int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) \sum_{i=1}^d \partial_i g(\bar{u}_0) \sum_{k=1}^d \mathbf{A}_k^i \bar{u}_0^k \\ &\quad + (\eta_n/2) \int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) \sum_{i, j=1}^d \partial_{ij}^2 g(\bar{u}_0) \Sigma_{\star}^{ij} + \eta_n \int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) \mathcal{R}_{g, \eta_n, \Sigma, h}(\bar{u}_0) \\ &\quad + (\eta_n/6) \int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) \bar{\mathcal{R}}_{g, \eta_n}(\bar{u}_0). \end{aligned}$$

Therefore using that $\bar{\nu}^{\eta_n}$ is stationary with respect to \bar{R}_{η_n} , we obtain that

$$\begin{aligned} &\limsup_{n \rightarrow +\infty} \left| \int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) \left\{ \sum_{i=1}^d \partial_i g(\bar{u}_0) \sum_{k=1}^d \mathbf{A}_k^i \bar{u}_0^k + \sum_{i, j=1}^d \partial_{ij}^2 g(\bar{u}_0) \Sigma_{\star}^{ij} \right\} \right| \\ &\leq \limsup_{n \rightarrow +\infty} \left| \int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) \mathcal{R}_{g, \eta_n, \Sigma, h}(\bar{u}_0) \right| + \left| \int_{\mathbb{T}_{\theta^* \Theta}} \bar{\nu}^{\eta_n}(\mathrm{d}\bar{u}_0) \bar{\mathcal{R}}_{g, \eta_n}(\bar{u}_0) \right|. \quad (3.65) \end{aligned}$$

Consider a sequence of independent random variables $(Y_n)_{n \in \mathbb{N}}$ such that for any $n \in \mathbb{N}$, the law of Y_n is $\bar{\nu}^{\eta_n}$. By Slutsky's theorem, since $(Y_n)_{n \in \mathbb{N}}$ converges in distribution and $\lim_{n \rightarrow +\infty} \eta_n = 0$, we obtain that $\eta_n^{1/2} Y_n$ converges in distribution towards 0. Moreover, using the continuous mapping theorem, we have

$$\limsup_{n \rightarrow +\infty} |\mathbb{E}[\mathcal{R}_{g, \eta_n, \Sigma, h}(Y_n)]| = 0. \quad (3.66)$$

3.10. Proofs for Section 3.4

Similarly, we use (3.44) to obtain, for any $n \in \mathbb{N}$ and $K > 0$,

$$\begin{aligned} \left| \overline{\mathcal{R}}_{g, \eta_n}(Y_n) \right| &\leq 8\eta_n^{1/2} \mathbb{1}_{\mathcal{K}_K}(\theta_n) \mathbb{E} \left[\|\mathbf{C}^3(g, \eta_n)\|_{\mathcal{Y}} \mathbb{1}_{\mathcal{A}_{\theta_n}^c} \|H_K\|_{\theta_n}^3 \mid \theta_n \right] \\ &\quad + 16\|\mathbf{C}^2(g, \eta_n)\| \mathbb{E} \left[\|Y_K\|_{\theta_n}^2 \mid \theta_n \right], \end{aligned}$$

where for any $n \in \mathbb{N}$, $\theta_n = \text{Exp}_{\theta^*}(\eta_n^{1/2} Y_n)$ are independent random variables and by (3.40), the distribution of θ_n is μ^{η_n} . Thus we obtain for any $K \geq 0$, using $\mathbb{1}_{\mathcal{K}_K}(\theta_n) \|H_K\|_{\theta_n}$ is almost surely bounded by $4[K^3 + \sup_{\theta \in \mathcal{K}_K} \|h(\theta)\|_{\theta}^3]$, Markov's inequality and **MD6**,

$$\begin{aligned} \limsup_{n \rightarrow +\infty} \left| \mathbb{E}[\overline{\mathcal{R}}_{g, \eta_n}(Y_n)] \right| &\leq \limsup_{n \rightarrow +\infty} 16\|\mathbf{C}^2(g, \eta_n)\| \mathbb{E}[\|e_{\theta_n}(X_1)\|_{\theta_n}^2 \{1 - \chi_K(\|e_{\theta_n}(X_1)\|_{\theta_n})\}] \\ &\leq 16\|\mathbf{C}^2(g, 0)\| K^{-\varepsilon} \{\tilde{\sigma}_0^2 + \tilde{\sigma}_1^2 \mathbb{E}[V(\theta^*)]\}, \end{aligned} \quad (3.67)$$

using that $(\theta_n)_{n \in \mathbb{N}}$ converges in distribution to θ^* . For any smooth function with compact support $g : \mathcal{T}_{\theta^*}\Theta \rightarrow \mathbb{R}$, combining (3.65)-(3.66)-(3.67), taking $K \rightarrow +\infty$ and using the weak convergence of $(\bar{\nu}^n)_{n \in \mathbb{N}}$ to $\bar{\nu}^*$ when $n \rightarrow +\infty$ shows that

$$\int_{\mathcal{T}_{\theta^*}\Theta} \bar{\nu}^*(d\bar{u}_0) \left\{ \sum_{i=1}^d \partial_i g(\bar{u}_0) \sum_{k=1}^d \mathbf{A}_k^i \bar{u}_0^k + \sum_{i,j=1}^d \partial_{ij}^2 g(\bar{u}_0) \Sigma_{\star}^{ij} \right\} = 0. \quad (3.68)$$

Finally, by [HJ94, Theorem 2.2.1], there exists a unique matrix $\mathbf{V} \in \mathbb{R}^{d \times d}$ solution to the Lyapunov equation $\mathbf{A}\mathbf{V} + \mathbf{V}\mathbf{A}^\top = \Sigma(\theta^*)$. By [Ken78, Theorem 10.1], $\mathcal{N}(0, \mathbf{V})$ is the unique probability distribution on $\mathcal{T}_{\theta^*}\Theta$ satisfying (3.68). This concludes the proof. \square

3.10 Proofs for Section 3.4

3.10.1 Proof of Lemma 3.8

Recall that f is λ_f -strongly geodesically convex, if and only if for any $\theta_1, \theta_2 \in \Theta$,

$$f(\theta_2) \geq f(\theta_1) + \left\langle \text{Exp}_{\theta_1}^{-1}(\theta_2), \text{grad } f(\theta_1) \right\rangle_{\theta_1} + \lambda_f \rho_{\Theta}^2(\theta_1, \theta_2). \quad (3.69)$$

Put $\theta_1 = \theta^*$ and $\theta_2 = \theta$. Since θ^* is a stationary point of f , so $\text{grad } f(\theta^*) = 0$, it follows from (3.69) that

$$f(\theta) - f(\theta^*) \geq \lambda_f \rho_{\Theta}^2(\theta^*, \theta),$$

which is the second identity in (3.13). To obtain the first identity, put $\theta_1 = \theta$ and $\theta_2 = \theta^*$, in (3.69), so

$$f(\theta^*) - f(\theta) \geq \left\langle \text{Exp}_{\theta}^{-1}(\theta^*), \text{grad } f(\theta) \right\rangle_{\theta} + \lambda_f \rho_{\Theta}^2(\theta^*, \theta). \quad (3.70)$$

Since $f(\theta^*) \leq f(\theta)$, this implies

$$-\left\langle \text{Exp}_{\theta}^{-1}(\theta^*), \text{grad } f(\theta) \right\rangle_{\theta} \geq \lambda_f \rho_{\Theta}^2(\theta^*, \theta) = \lambda_f \left\| \text{Exp}_{\theta}^{-1}(\theta^*) \right\|_{\theta}^2.$$

Or, after using the Cauchy-Schwarz inequality,

$$\|\text{grad } f(\theta)\|_{\theta} \geq \lambda_f \left\| \text{Exp}_{\theta}^{-1}(\theta^*) \right\|_{\theta} . \quad (3.71)$$

Finally, using once more the Cauchy-Schwarz inequality, and (3.70) and (3.71),

$$f(\theta) - f(\theta^*) \leq - \left\langle \text{Exp}_{\theta}^{-1}(\theta^*), \text{grad } f(\theta) \right\rangle_{\theta} \leq (1/\lambda_f) \|\text{grad } f(\theta)\|_{\theta}^2 ,$$

which is equivalent to the first identity in (3.13).

3.10.2 Proof of Lemma 3.10

Without loss of generality, we assume that $f(\theta^*) = 0$. First, we show that for any $\theta \in \Theta$,

$$f(\theta) \leq M_f \rho_{\Theta}^2(\theta^*, \theta) . \quad (3.72)$$

Let $\theta \in \Theta$ and $\gamma : [0, 1] \rightarrow \Theta$ the unique geodesic such that $\gamma(0) = \theta^*$ and $\gamma(1) = \theta$. Then since f is continuously differentiable using [Lee19, Proposition 4.15-(ii) and Theorem 4.24-(iii)], we get that $f(\theta) = \int_0^1 \langle \text{grad } f(\gamma(t)), \dot{\gamma}(t) \rangle_{\gamma(t)} dt$. Therefore, using the Cauchy-Schwarz inequality and for any $t \in [0, 1]$, $\|\dot{\gamma}(t)\|_{\gamma(t)} = \rho_{\Theta}(\theta^*, \theta)$ we obtain that $|f(\theta)| \leq \rho_{\Theta}(\theta^*, \theta) \|\text{grad } f(\gamma(t))\|_{\gamma(t)}$ which shows that (3.72) holds by assumption.

We now proceed with the proof of the main statement. Since f is twice continuously differentiable, \tilde{f} has this same property. In addition, for any $\theta \in \Theta$,

$$\text{grad } \tilde{f}(\theta) = \text{grad } f(\theta) / [2(f(\theta) + 1)^{1/2}] . \quad (3.73)$$

Therefore, using the assumption that for any $\theta \in \Theta$, $\|\text{grad } f(\theta)\|_{\theta}^2 \leq M_f \rho_{\Theta}^2(\theta^*, \theta)$ and the second inequality of Lemma 3.8, we get that

$$\begin{aligned} \|\text{grad } \tilde{f}(\theta)\|_{\theta} &= \|\text{grad } f(\theta)\|_{\theta} / [2(f(\theta) + 1)^{1/2}] \leq M_f^{1/2} \rho_{\Theta}(\theta^*, \theta) / [2(\lambda_f \rho_{\Theta}^2(\theta^*, \theta) + 1)^{1/2}] \\ &\leq C_f^{1/2} [1 \wedge \rho_{\Theta}(\theta^*, \theta)] , \end{aligned}$$

with $C_f^{1/2} \leftarrow (M_f^{1/2}/2)[1 \wedge \lambda_f^{-1/2}]$.

It remains to show that for any $\theta \in \Theta$, $-\langle \text{Exp}_{\theta}^{-1}(\theta^*), \text{grad } \tilde{f}(\theta) \rangle_{\theta} \geq \tilde{\lambda}_f V_1(\theta)$, where V_1 is defined by (3.9) with $\delta = 1$ and $\tilde{\lambda}_f \leftarrow \lambda_f^{1/2}/2$. Using (3.73) again, F2 and (3.72), we obtain that for any $\theta \in \Theta$,

$$\begin{aligned} - \left\langle \text{Exp}_{\theta}^{-1}(\theta^*), \text{grad } \tilde{f}(\theta) \right\rangle_{\theta} &= - \left\langle \text{Exp}_{\theta}^{-1}(\theta^*), \text{grad } f(\theta) \right\rangle_{\theta} / [2(f(\theta) + 1)^{1/2}] \\ &\geq \lambda_f \rho_{\Theta}^2(\theta^*, \theta) / [2(f(\theta) + 1)^{1/2}] \geq \lambda_f \rho_{\Theta}^2(\theta^*, \theta) / [2(M_f \rho_{\Theta}^2(\theta^*, \theta) + 1)^{1/2}] . \end{aligned}$$

Using that for any $\theta \in \Theta$, $V_1(\theta) = \{\rho_{\Theta}^2(\theta^*, \theta) + 1\}^{1/2} - 1 \leq \rho_{\Theta}(\theta^*, \theta)$, we get that

$$- \left\langle \text{Exp}_{\theta}^{-1}(\theta^*), \text{grad } \tilde{f}(\theta) \right\rangle_{\theta} \geq \lambda_f V_1(\theta) \rho_{\Theta}(\theta^*, \theta) / [2(M_f \rho_{\Theta}^2(\theta^*, \theta) + 1)^{1/2}] \geq \lambda_f V_1(\theta) / (2M_f^{1/2}) .$$

3.10.3 Proof of Proposition 3.11

The proof consists in an application of Theorem 3.1-(b). First, by Proposition 3.5, V_1 defined by (3.9) with $\delta = 1$, satisfies **A6**. In addition, by [Dur+20, Lemma 16], V_1 is continuously differentiable with gradient given for any $\theta \in \Theta$ by

$$\text{grad } V_1(\theta) = -\text{Exp}_{\theta}^{-1}(\theta^*)/\{1 + \rho_{\Theta}^2(\theta^*, \theta)\}^{1/2}.$$

Therefore, for any $\theta \in \Theta$, by **F3** we get

$$\begin{aligned} \langle \text{grad } V_1(\theta), \text{grad } f(\theta) \rangle_{\theta} &= -\left\langle \text{Exp}_{\theta}^{-1}(\theta^*), \text{grad } f(\theta) \right\rangle_{\theta} / \{1 + \rho_{\Theta}^2(\theta^*, \theta)\}^{1/2} \\ &\geq \tilde{\lambda}_f V_1(\theta) / \{1 + \rho_{\Theta}^2(\theta^*, \theta)\}^{1/2}. \end{aligned} \quad (3.74)$$

In addition, $t^2 \wedge 1 - ab\{(t^2 + 1)^{1/2} - 1\}/(1 + t^2)^{1/2} \leq 0$ for any $t \geq 0$, $b > 0$ and $a = 4b^{-1}$ using that $(t^2 + 1)^{1/2} - 1 \geq t^2/[2(1 + t^2)^{1/2}]$. As a result, using **F3** for any $t \geq 0$, $b > 0$ and $a = 4b^{-1}$, it follows that **A7** is satisfied with $C_1 \leftarrow 0$, $C_2 \leftarrow 4C_f/\tilde{\lambda}_f$ for $h = -\text{grad } f$ and $V \leftarrow V_1$. Therefore, we obtain using Theorem 3.1-(b) that for any $\eta \in (0, \bar{\eta}]$,

$$n^{-1} \sum_{k=0}^{n-1} \mathbb{E} \left[\langle \text{grad } V_1(\theta_k), \text{grad } f(\theta_k) \rangle_{\theta_k} \right] \leq 2V_1(\theta_0)/(n\bar{\eta}) + 2\bar{\eta}(1 + \kappa)\sigma_0^2,$$

where $\bar{\eta} = [(8C_f/\tilde{\lambda}_f)(1 + \kappa)(1 + \sigma_1^2)]^{-1}$. Using (3.74), we have

$$(\tilde{\lambda}_f/n) \sum_{k=0}^{n-1} \mathbb{E} \left[V_1(\theta_k) / \{1 + \rho_{\Theta}^2(\theta^*, \theta_k)\}^{1/2} \right] \leq 2V_1(\theta_0)/(n\bar{\eta}) + 2\bar{\eta}(1 + \kappa)\sigma_0^2,$$

which concludes the proof since $(t^2 + 1)^{1/2} - 1 \geq t^2/[2(1 + t^2)^{1/2}]$ for any $t \geq 0$ implying $V_1(\theta)/\{1 + \rho_{\Theta}^2(\theta^*, \theta)\}^{1/2} \geq D_{\Theta}^2(\theta^*, \theta)/2$ for any $\theta \in \Theta$.

3.10.4 Proof of Proposition 3.12

Define $\mathbf{X} = \{\bar{\theta}_i : i \in \{1, \dots, M_{\pi}\}\}$ and recall that $\mathbf{D} = \sup\{\rho_{\Theta}(\theta_0, \bar{\theta}) : \bar{\theta} \in \mathbf{X}\}$. Set $\mathbf{S} = \bar{\mathbf{B}}(\theta_0, \mathbf{D})$. Note that the closed ball \mathbf{S} , is compact by [Jos05, Theorem 1.7.1], geodesically convex, and $\mathbf{X} \subset \mathbf{S}$, as well as $\theta_0 \in \mathbf{S}$. We consider in this section, for any $\theta \in \Theta$ and $x \in \mathbf{X}$, $H_{\theta}(x) = \text{Exp}_{\theta}^{-1}(x)$.

First note that $\theta_n \in \mathbf{S}$, for all $n \in \mathbb{N}$ by a straightforward induction using that \mathbf{S} is geodesically convex and $\theta_0 \in \mathbf{S}$. Indeed, $\theta_0 \in \mathbf{S}$, and, if $\theta_n \in \mathbf{S}$, then θ_{n+1} lies on the geodesic segment connecting θ_n and X_{n+1} , two points which belong to \mathbf{S} , and therefore $\theta_{n+1} \in \mathbf{S}$. This means that the SGD scheme used here is equivalent to

$$\theta_{n+1} = \text{proj}_{\mathbf{S}}(\text{Exp}_{\theta_n}(\eta H_{\theta_n}(X_{n+1}))) .$$

Define \mathbf{H} and V_2 as in Proposition 3.4. It is possible to show that $\mathbf{H} = \mathbf{S}$. Indeed, for $\theta \in \mathbf{S}$, and $x \in \mathbf{X}$, since $x \in \mathbf{S}$, and \mathbf{S} is convex, the geodesic segment connecting θ to x is entirely contained in \mathbf{S} . However, by definition, this geodesic segment is the set of points

$\text{Exp}_\theta(tH_\theta(x))$, where $t \in [0, 1]$. Now, since $\eta \leq \bar{\eta} \leq 1$, Proposition 3.4 implies that V_2 verifies **A6-(i)-(ii)** where $L \leftarrow CL_\pi$, $L_\pi = (D + 1)(1 + \kappa \coth(\kappa D))$ and C is a universal constant.

The objective function f satisfies **F2** with $\lambda_f = 1/2$ (that is, f is 1/2-strongly convex), since by [Jos05, Theorem 5.6.1] $f_i(\theta) = \rho_\Theta^2(\theta, \bar{\theta}_i)/2$ is 1-strongly geodesically convex for any $i \in \{1, \dots, M_\pi\}$. Thus, by (3.69) for all $\theta \in \mathbb{S}$

$$\left\langle \text{Exp}_\theta^{-1}(\theta^*), \text{grad } f(\theta) \right\rangle_\theta \leq -(1/2)\rho_\Theta^2(\theta^*, \theta). \quad (3.75)$$

Now, for any $\theta \in \mathbb{S}$, $v \in \mathbb{T}_\theta\Theta$, using [Jos05, Theorem 5.6.1], we have,

$$\begin{aligned} \|\text{Hess } f_\theta(v, v)\|_\theta &\leq M_\pi^{-1} \sum_{i=1}^{M_\pi} \|(\text{Hess } f_i)_\theta(v, v)\|_\theta \\ &\leq M_\pi^{-1} \sum_{i=1}^{M_\pi} \kappa \rho_\Theta(\theta, \bar{\theta}_i) \coth(\kappa \rho_\Theta(\theta, \bar{\theta}_i)) \|v\|_\theta^2 \leq \tilde{L}_\pi \|v\|_\theta^2, \end{aligned}$$

where $\tilde{L}_\pi = 2D\kappa \coth(2\kappa D)$, since $t \mapsto t \coth(t)$ is non-decreasing over \mathbb{R}_+ . Therefore, by [Dur+20, Lemma 10], $\text{grad } f$ is geodesically \tilde{L}_π -Lipschitz continuous on \mathbb{S} . In particular, for any $\theta \in \mathbb{S}$,

$$\|\text{grad } f(\theta)\|_\theta \leq \tilde{L}_\pi \rho_\Theta(\theta^*, \theta). \quad (3.76)$$

By (3.75) and (3.76), it is straightforward that $V = V_2$ and $h = -\text{grad } f$ satisfy **A7**, with $C_1 = 0$ and $C_2 = 2\tilde{L}_\pi^2 \leq 2^5 L_\pi^2$. In addition, by Proposition 3.4, (3.75) implies V_2 verifies **A8-(\emptyset)**, with $\lambda = 1/2$.

Finally, **MD3** holds with $\sigma_0^2 = D^2$ and $\sigma_1^2 = 0$ since for any $\theta \in \mathbb{S}$ and $x \in \mathbb{X}$,

$$\|H_\theta(x)\|_\theta = \left\| \text{Exp}_\theta^{-1}(x) \right\|_\theta \leq 2D.$$

Therefore, we can apply Theorem 3.1-(c) which implies that for any $\eta \leq \bar{\eta}$,

$$\mathbb{E}[V_2(\theta_n)] \leq \{1 - \eta/4\}^n V_2(\theta_0) + 4\eta L_\pi D^2.$$

To conclude, it only remains to note that $V_2(\theta_n) = \rho_\Theta^2(\theta^*, \theta_n)$ and $V_2(\theta_0) = \rho_\Theta^2(\theta^*, \theta_0)$, since $(\theta_n)_{n \in \mathbb{N}}$ and θ^* belong to $\mathbb{H} = \mathbb{S}$.

3.10.5 Proof of Theorem 3.13

We consider in this section the recursion

$$\begin{aligned} \theta_{n+1} &= \text{Exp}_{\theta_n}[\eta H_{\theta_n}(X_{n+1})] \\ H_{\theta_n}(X_{n+1}) &= \text{Exp}_{\theta_n}^{-1}\left(X_{n+1}^{(1)}\right) / \left(2\{\rho_\Theta^2(\theta_n, X_{n+1}^{(2)})/2 + 1\}^{1/2}\right), \end{aligned} \quad (3.77)$$

where $X_{n+1} = (X_{n+1}^{(1)}, X_{n+1}^{(2)})$ and $(X_n^{(1)}, X_n^{(2)})_{n \in \mathbb{N}^*}$ is an i.i.d. sequence of pairs of independent random variables with distribution π . Denote by Q_η the Markov kernel corresponding to (3.77).

We give first some additional intuition and motivation behind the scheme (3.77). It can be interpreted as a stochastic optimization method to minimize

$$\tilde{f}_\pi = (f_\pi + 1)^{1/2},$$

in place of f_π . First note that f_π and \tilde{f}_π have the same minimizer, but compared to f_π it may be shown that $\text{grad } \tilde{f}_\pi$, given for any $\theta \in \Theta$ by

$$\text{grad } \tilde{f}_\pi(\theta) = (1/2)\text{grad } f_\pi(\theta)(f_\pi(\theta) + 1)^{-1/2},$$

is geodesically Lipschitz. However, note that (3.77) is not an unbiased stochastic optimization scheme for the function \tilde{f}_π since

$$\mathbb{E}[H_{\theta_n}(X_{n+1})] = (1/2)\{\text{grad } f_\pi(\theta_n)\}\mathbb{E}\left[\{\rho_\Theta^2(\theta_n, X_{n+1}^{(2)})/2 + 1\}^{-1/2}\right].$$

The proof of Theorem 3.13 then consists in adapting the proof of Theorem 3.1 to deal with this additional difficulty taking for the Lyapunov function V, V_1 defined by (3.9) with $\delta = 1$. A general theory could be derived but we believe that this is out the scope of the present document and leave it for future work. We start by preliminary technical results which are needed to establish Theorem 3.13.

Lemma 3.24. *Assume A 7 and MD 7. Let θ_π^* be the Riemannian barycenter of the probability measure π , i.e. $\theta_\pi^* = \text{argmin}_\Theta f_\pi$ where f_π is defined by (3.16). Then, for any $\theta \in \Theta$,*

$$-\int_\Theta \left\langle \text{Exp}_\theta^{-1}(\theta_\pi^*), \text{Exp}_\theta^{-1}(\nu) \right\rangle_\theta \pi(d\nu) \leq -\rho_\Theta^2(\theta, \theta_\pi^*)/2.$$

Proof. Using A 7 and [Jos05, Theorem 5.6.1], we have that for any $\nu \in \Theta$, the operator norm of the Riemannian Hessian of $\theta \mapsto \rho_\Theta^2(\theta, \nu)/2$ is lower bounded by 1. Therefore, by [Bou20, Theorem 11.19], $\theta \mapsto \rho_\Theta^2(\theta, \nu)/2$ is 1/2-strongly convex. Applying this to θ and $\theta_\pi^* \in \Theta$, we have for any $\nu \in \Theta$,

$$\rho_\Theta^2(\theta_\pi^*, \nu)/2 - \rho_\Theta^2(\theta, \nu)/2 \geq -\left\langle \text{Exp}_\theta^{-1}(\theta_\pi^*), \text{Exp}_\theta^{-1}(\nu) \right\rangle_\theta + \rho_\Theta^2(\theta, \theta_\pi^*)/2.$$

Using MD 7, we can integrate this inequality with respect to π , bringing

$$f_\pi(\theta_\pi^*) - f_\pi(\theta) \geq -\int_\Theta \left\langle \text{Exp}_\theta^{-1}(\theta_\pi^*), \text{Exp}_\theta^{-1}(\nu) \right\rangle_\theta \pi(d\nu) + \rho_\Theta^2(\theta, \theta_\pi^*)/2.$$

Since by definition of θ_π^* , $0 \geq f_\pi(\theta_\pi^*) - f_\pi(\theta)$, this completes the proof. \square

Lemma 3.25. *Assume A 7 and MD 7. Let θ_π^* be the Riemannian barycenter of the probability measure π , i.e. $\theta_\pi^* = \text{argmin}_\Theta f_\pi$ where f_π is defined by (3.16). Then, for any $\theta \in \Theta$,*

$$\int_\Theta \{\rho_\Theta^2(\theta, \nu)/2 + 1\}^{-1/2} \pi(d\nu) \geq \{\rho_\Theta^2(\theta, \theta_\pi^*) + 2f_\pi(\theta_\pi^*) + 1\}^{-1/2}.$$

Proof. Let $\theta \in \Theta$. Using Jensen's inequality with the convex function $t \mapsto (t+1)^{-1/2}$ on \mathbb{R}_+ , we have

$$\int_{\Theta} \{\rho_{\Theta}^2(\theta, \nu)/2 + 1\}^{-1/2} \pi(d\nu) \geq \{f_{\pi}(\theta) + 1\}^{-1/2}. \quad (3.78)$$

However, using the triangle and Hölder's inequalities, we have for any θ and $\nu \in \Theta$, $\rho_{\Theta}^2(\theta, \nu)/2 \leq \rho_{\Theta}^2(\theta, \theta_{\pi}^*) + \rho_{\Theta}^2(\theta_{\pi}^*, \nu)$. Taking the integral with respect to π , by **MD7** we get $f_{\pi}(\theta) \leq \rho_{\Theta}^2(\theta, \theta_{\pi}^*) + 2f_{\pi}(\theta_{\pi}^*)$. Lastly, combining this result with (3.78) and using that the function $t \mapsto (t+1)^{-1/2}$ is non-increasing on \mathbb{R}_+ completes the proof. \square

Lemma 3.26. *Assume **A 7** and **MD 7**. Let θ_{π}^* be the Riemannian barycenter of the probability measure π , i.e. $\theta_{\pi}^* = \operatorname{argmin}_{\Theta} f_{\pi}$ where f_{π} is defined by (3.16). Then, for any $\theta_0 \in \Theta$,*

$$\mathcal{Q}_{\eta} V_1(\theta_0) \leq V_1(\theta_0) - [\eta/(4C_{\pi}^{1/2})] D_{\Theta}^2(\theta_0, \theta_{\pi}^*) + 2\eta^2(1 + \kappa)\{1 + f_{\pi}(\theta_{\pi}^*)\} (f_{\pi}(\theta_{\pi}^*) + 2),$$

where V_1 is defined in (3.9) with $\delta \leftarrow 1$, $\theta^* \leftarrow \theta_{\pi}^*$, $C_{\pi} = 1 + 2f_{\pi}(\theta_{\pi}^*)$ and $D_{\Theta}^2 : \Theta^2 \rightarrow [0, 1]$ is defined by (3.14).

Proof. Let $\theta_0 \in \Theta$, and consider

$$H_{\theta_0}(X) = (1/2) \operatorname{Exp}_{\theta_0}^{-1}(X^{(1)}) / \left\{ \rho_{\Theta}^2(\theta_0, X^{(2)}) / 2 + 1 \right\}^{1/2},$$

where $X^{(1)}, X^{(2)}$ are independent random variables with distribution π .

Let $\gamma : [0, 1] \rightarrow \Theta$ be the geodesic curve defined by $\gamma : t \mapsto \operatorname{Exp}_{\theta_0}[t\eta H_{\theta_0}(X)]$. Using [Dur+20, Lemma 1] with γ and V_1 , we get

$$\begin{aligned} V_1(\gamma(1)) &\leq V_1(\theta_0) + \langle \operatorname{grad} V_1(\theta_0), \dot{\gamma}(0) \rangle_{\theta_0} + (L/2) \|\dot{\gamma}(0)\|_{\theta_0}^2 \\ &= V_1(\theta_0) + \eta \langle \operatorname{grad} V_1(\theta_0), H_{\theta_0}(X) \rangle_{\theta_0} + ((1 + \kappa)\eta^2/2) \|H_{\theta_0}(X)\|_{\theta_0}^2, \end{aligned} \quad (3.79)$$

by Proposition 3.5. We now compute the expectation of the terms in (3.79). Using that $(X^{(1)}, X^{(2)})$ are independent, we obtain

$$\mathbb{E} \left[\langle \operatorname{grad} V_1(\theta_0), H_{\theta_0}(X) \rangle_{\theta_0} \right] = (1/2) \left\langle \operatorname{grad} V_1(\theta_0), \mathbb{E} \left[\operatorname{Exp}_{\theta_0}^{-1}(X^{(1)}) \right] \mathbb{E} \left[\left\{ \rho_{\Theta}^2(\theta_0, X^{(2)}) / 2 + 1 \right\}^{-1/2} \right] \right\rangle_{\theta_0}$$

Moreover, using (3.27) and Lemmas 3.24 and 3.25 yields

$$\begin{aligned} &\mathbb{E} \left[\langle \operatorname{grad} V_1(\theta_0), H_{\theta_0}(X) \rangle_{\theta_0} \right] \\ &= -(1/2) \left\{ \rho_{\Theta}^2(\theta_0, \theta_{\pi}^*) + 1 \right\}^{-1/2} \mathbb{E} \left[\left\langle \operatorname{Exp}_{\theta_0}^{-1}(\theta_{\pi}^*), \operatorname{Exp}_{\theta_0}^{-1}(X^{(1)}) \right\rangle_{\theta_0} \right] \mathbb{E} \left[\left\{ \rho_{\Theta}^2(\theta_0, X^{(2)}) / 2 + 1 \right\}^{-1/2} \right] \\ &\leq -(1/4) \rho_{\Theta}^2(\theta_0, \theta_{\pi}^*) \left[\left\{ \rho_{\Theta}^2(\theta_0, \theta_{\pi}^*) + 1 \right\} \left\{ \rho_{\Theta}^2(\theta_0, \theta_{\pi}^*) + 2f_{\pi}(\theta_{\pi}^*) + 1 \right\} \right]^{-1/2} \\ &\leq -(16C_{\pi})^{-1/2} D_{\Theta}^2(\theta_0, \theta_{\pi}^*), \end{aligned} \quad (3.80)$$

where $C_\pi = 1 + 2f_\pi(\theta_\pi^*)$ and $D_\Theta^2 : \Theta^2 \rightarrow [0, 1]$ is defined by (3.14). Looking to bound the expectation of the last term in (3.79), we use that $\|\text{Exp}_{\theta_0}^{-1}(X^{(1)})\|_{\theta_0} = \rho_\Theta(\theta_0, X^{(1)})$ and that $X^{(1)}$ has distribution π to obtain,

$$\begin{aligned} \mathbb{E} \left[\|H_{\theta_0}(X)\|_{\theta_0}^2 \right] &= (1/4) \mathbb{E} \left[\rho_\Theta^2(\theta_0, X^{(1)}) \right] \mathbb{E} \left[\left\{ \rho_\Theta^2(\theta_0, X^{(2)})/2 + 1 \right\}^{-1} \right] \\ &= (f_\pi(\theta_0)/2) \mathbb{E} \left[\left\{ \rho_\Theta^2(\theta_0, X^{(2)})/2 + 1 \right\}^{-1} \right]. \end{aligned} \quad (3.81)$$

Denote by $M = \rho_\Theta(\theta_\pi^*, \theta_0)/2$. We bound the expectation in (3.81) using the event $\{\rho_\Theta(\theta_\pi^*, X^{(2)}) \geq M\}$ and its complement. On $\{\rho_\Theta(\theta_\pi^*, X^{(2)}) \geq M\}$, we use Markov's inequality with the increasing map $t \mapsto t^2/2 + 1$,

$$\begin{aligned} \mathbb{E} \left[\mathbb{1}_{[M, +\infty)}(\rho_\Theta(\theta_\pi^*, X^{(2)})) / [\rho_\Theta^2(\theta_0, X^{(2)})/2 + 1] \right] &\leq \mathbb{P} \left(\rho_\Theta(\theta_\pi^*, X^{(2)}) \geq M \right) \\ &\leq \left(\mathbb{E} \left[\rho_\Theta^2(\theta_\pi^*, X^{(2)}) \right] / 2 + 1 \right) / (M^2/2 + 1). \end{aligned} \quad (3.82)$$

On $\{\rho_\Theta(\theta_\pi^*, X^{(2)}) < M\}$, using the triangle inequality, we have

$$\rho_\Theta(\theta_0, X^{(2)}) \geq |\rho_\Theta(\theta_0, \theta_\pi^*) - \rho_\Theta(\theta_\pi^*, X^{(2)})| \geq \rho_\Theta(\theta_0, \theta_\pi^*) - M = M.$$

Then, we obtain

$$\mathbb{E} \left[\mathbb{1}_{[0, M)}(\rho_\Theta(\theta, X^{(2)})) / \left\{ \rho_\Theta^2(\theta_0, X^{(2)})/2 + 1 \right\} \right] \leq 1/[M^2/2 + 1]. \quad (3.83)$$

Adding (3.82) and (3.83) together and using the definition of M we obtain,

$$\mathbb{E} \left[\left\{ \rho_\Theta^2(\theta_0, X^{(2)})/2 + 1 \right\}^{-1} \right] \leq (f_\pi(\theta_\pi^*) + 2) / \left[\rho_\Theta^2(\theta_\pi^*, \theta_0)/8 + 1 \right]. \quad (3.84)$$

Plugging (3.84) in (3.81), we get

$$\mathbb{E} \left[\|H_{\theta_0}(X)\|_{\theta_0}^2 \right] \leq (f_\pi(\theta_0)/2) (f_\pi(\theta_\pi^*) + 2) / \left[\rho_\Theta^2(\theta_\pi^*, \theta_0)/8 + 1 \right]. \quad (3.85)$$

Using the triangle and Hölder's inequalities, we have for any θ and $\nu \in \Theta$, $\rho_\Theta^2(\theta, \nu)/2 \leq \rho_\Theta^2(\theta, \theta_\pi^*) + \rho_\Theta^2(\theta_\pi^*, \nu)$. Taking the integral with respect to π , by **MD 7** we get $f_\pi(\theta) \leq \rho_\Theta^2(\theta, \theta_\pi^*) + 2f_\pi(\theta_\pi^*)$. Combining this result and (3.85), we obtain

$$\mathbb{E} \left[\|H_{\theta_0}(X)\|_{\theta_0}^2 \right] \leq \left\{ \rho_\Theta^2(\theta_\pi^*, \theta_0)/2 + f_\pi(\theta_\pi^*) \right\} (f_\pi(\theta_\pi^*) + 2) / \left[\rho_\Theta^2(\theta_\pi^*, \theta_0)/8 + 1 \right] \leq 4\{1 + f_\pi(\theta_\pi^*)\} (f_\pi(\theta_\pi^*) + 2).$$

Combining this result and (3.80) in (3.79) concludes the proof. \square

Proof of Theorem 3.13. Let $\theta_0 \in \Theta, \eta > 0$ and $n \in \mathbb{N}$. Then, for any $k \in \{1, \dots, n\}$, using Markov's property and Lemma 3.26 we have,

$$\begin{aligned} [\eta/(4C_\pi^{1/2})] \mathbb{E} \left[D_\Theta^2(\theta_{k-1}, \theta_\pi^*) \right] &= [\eta/(4C_\pi^{1/2})] \int_\Theta D_\Theta^2(\theta, \theta_\pi^*) Q_\eta^{k-1}(\theta_0, d\theta) \\ &\leq Q_\eta^{k-1} V_1(\theta_0) - Q_\eta^k V_1(\theta_0) + 2\eta^2(1 + \kappa)(1 + f_\pi(\theta_\pi^*))(f_\pi(\theta_\pi^*) + 2). \end{aligned}$$




Figure 3.4: Monte Carlo approximations of the mean distance at convergence in Theorem 3.13

Summing these inequalities for $k \in \{1, \dots, n\}$ implies that

$$[\eta/(4C_{\pi}^{1/2})] \sum_{k=0}^{n-1} \mathbb{E} \left[D_{\Theta}^2(\theta_k, \theta_{\pi}^*) \right] \leq V_1(\theta_0) - Q_{\eta}^n V_1(\theta_0) + 2n\eta^2(1 + \kappa)(1 + f_{\pi}(\theta_{\pi}^*))(f_{\pi}(\theta_{\pi}^*) + 2) .$$

Finally, dividing both sides by $[n\eta/(4C_{\pi}^{1/2})]$ and using that V_1 is a non-negative function, we obtain

$$n^{-1} \sum_{k=0}^{n-1} \mathbb{E} \left[D_{\Theta}^2(\theta_k, \theta_{\pi}^*) \right] \leq 2V_1(\theta_0)C_{\pi}^{1/2} / (\eta n) + 2\eta(1 + \kappa)(f_{\pi}(\theta_{\pi}^*) + 1)(f_{\pi}(\theta_{\pi}^*) + 2)(2f_{\pi}(\theta_{\pi}^*) + 1)^{-1/2} .$$

Which concludes the proof by setting $B_{\pi} = (1 + \kappa)(f_{\pi}(\theta_{\pi}^*) + 1)(f_{\pi}(\theta_{\pi}^*) + 2)(2f_{\pi}(\theta_{\pi}^*) + 1)^{-1/2}$. \square

Similarly to Figure 3.2, Figure 3.4 illustrates Theorem 3.7. To this end, 1000 replications of the experiment derived for Figure 3.3 are performed, obtaining $\{(\theta_n^{(i)}) : i \in \{1, \dots, 1000\}\}$ for $n = \lceil 50/\eta \rceil$ and $\eta \in \{1, 2.8, 4.6, 6.4, 8.2, 10\} \times 10^{-2}$. We estimate, with these samples, the mean and the variance of $D_{\Theta}^2(\theta, \theta_{\pi}^*)$, for θ following the stationary distribution μ^{π} . We observe that the mean and variance are both linear with respect to the step-size η , indicating that the iterates of the SA scheme remain in a neighborhood of diameter $\mathcal{O}(\eta^{1/2})$ to the ground truth.

Even though the setting of this experiment goes beyond the assumptions of Theorem 3.7, it suggests that such a result may be applicable also in the setting of Theorem 3.13. The proof of such a result is left for future work.

3.11 Background on Markov chain theory and Riemannian geometry

We give here some useful definitions and results that are used throughout the paper.

3.11.1 Markov chain notions

We refer to [MT09] for a general introduction to Markov chains in general state space. Let (Y, \mathcal{Y}) be a measurable state space and P be a Markov kernel on $Y \times \mathcal{Y}$. Consider for any $y \in Y$, the distribution \mathbb{P}_y of the canonical Markov chain $(Y_n)_{n \in \mathbb{N}}$ corresponding to P and starting from y on the canonical space $(Y^{\mathbb{N}}, \mathcal{Y}^{\otimes \mathbb{N}})$. Denote by \mathbb{E}_y the corresponding expectation.

Denote for any $A \in \mathcal{Y}$, $\tau_A = \inf\{l \geq 1 : Y_l \in A\}$ and $N_A = \sum_{l=1}^{+\infty} \mathbb{1}_{\{A\}}(Y_l)$.

We say that $(Y_n)_{n \in \mathbb{N}}$ is ψ -irreducible if there exists a measure ψ on \mathcal{Y} such that whenever $\psi(A) > 0$, we have $\mathbb{P}_y(\tau_A < \infty) > 0$ for any $y \in Y$. Moreover, a set $A \in \mathcal{Y}$ is called Harris-recurrent if $\mathbb{P}_y(N_A = \infty) = 1$ for any $y \in A$. Finally, a chain $(Y_n)_{n \in \mathbb{N}}$ is called Harris-recurrent if it is ψ -irreducible and every set $A \in \mathcal{Y}$ such that $\psi(A) > 0$ is Harris-recurrent.

Let $\bar{V} : Y \rightarrow [1, +\infty)$. We say that P is \bar{V} -uniformly geometrically ergodic if there exist $\rho \in [0, 1)$ and $C \geq 0$ such that for any $y \in Y$ and $k \in \mathbb{N}$, $\|\delta_y P^k - \mu\|_{\bar{V}} \leq C\rho^k \bar{V}(y)$, where $\|\cdot\|_{\bar{V}}$ is defined for two probability measures ν_1, ν_2 on (Y, \mathcal{Y}) by $\|\nu_1 - \nu_2\|_{\bar{V}} = \sup\{|\nu_1(g) - \nu_2(g)| : \sup_Y\{|g|/\bar{V}\} \leq 1\}$.

3.11.2 Useful results from Riemannian geometry

We now give definitions and auxiliary results related to tensor fields along curves, their derivatives, and Taylor expansions on Riemannian manifolds.

Let M be a smooth manifold with or without boundary. Given a smooth curve $\gamma : I \rightarrow M$ defined on an interval I , and any $k, l \in \mathbb{N}$, a (k, l) -tensor field along γ is a continuous map $F : I \rightarrow T^{(k, l)}\text{TM}$, such that $F(t) \in T^{(k, l)}(T_{\gamma(t)}M)$ for any $t \in I$, where $T^{(k, l)}\text{TM}$ is the bundle of (k, l) -tensors on M , see e.g. [Lee19, Appendix B]. A vector field Y along γ is a $(1, 0)$ -tensor field, in which case for any $t \in I$, $Y(t)$ is just a tangent vector in $T_{\gamma(t)}M$. We say that a tensor field F along γ is extendible if there exists a tensor field \tilde{F} defined on a neighborhood of $\gamma(I)$ such that $F = \tilde{F} \circ \gamma$.

We let $\mathfrak{X}^{k, l}(\gamma)$ denote the set of smooth (k, l) -tensor fields along γ , and $\mathfrak{X}(\gamma) = \mathfrak{X}^{1, 0}(\gamma)$ denote the set of smooth vector fields along γ . In particular, $\mathfrak{X}^{0, 0}(\gamma)$ is the set of smooth functions $g : I \rightarrow \mathbb{R}$ such that for any $t \in I$, $g(t) = (\gamma(t), f(t))$ for some smooth function $f : I \rightarrow \mathbb{R}$ and therefore can be identified with the set of smooth functions $f : I \rightarrow \mathbb{R}$. In the sequel, we adopt if no confusion is possible this identification. We extend to tensor fields along γ the following definition of the trace on tensors. For any (k, l) -tensor T , we denote by $\text{Tr}_{\square, \Delta}(T)$ the $(k-1, l-1)$ -tensor with component of index $(i_1, \dots, i_{k-1}, j_1, \dots, j_{l-1})$, given by $\sum_{m=1}^d T_{i_1, \dots, i_{k-1}, m, i_{\square}, \dots, i_{k-1}}^{j_1, \dots, j_{\Delta-1}, m, j_{\Delta}, \dots, j_{l-1}}$. In particular, for any $\omega \in \mathfrak{X}^{0, 1}(\gamma), Y \in \mathfrak{X}(\gamma)$,

$$\text{Tr}_{(1, 1)}(\omega \otimes Y) = \omega(Y) .$$

Also, for any $F \in \mathfrak{X}^{k, l}(\gamma)$, any $\omega^1, \dots, \omega^{k_0} \in \mathfrak{X}^{0, 1}(\gamma)$ and $Y_1, \dots, Y_{l_0} \in \mathfrak{X}(\gamma)$, with $k_0 \leq k, l_0 \leq l$, denote by $[F : \omega^1 \otimes \dots \otimes \omega^{k_0} \otimes Y_1 \otimes \dots \otimes Y_{l_0}]$, the $(k - k_0, l - l_0)$ smooth tensor

field along γ defined by the induction:

$$[F : \omega^{\otimes 1:i}] = \text{Tr}_{(1,l+1)}([F : \omega^{\otimes 1:(i-1)}] \otimes \omega^i) \quad (3.86)$$

$$[F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:j}] = \text{Tr}_{(k-k_0+1,1)}([F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:(j-1)}] \otimes Y_j), \quad (3.87)$$

setting $\omega^{\otimes 1:i} = \omega^1 \otimes \cdots \otimes \omega^i$, $Y_{\otimes 1:j} = Y_1 \otimes \cdots \otimes Y_j$. Note that for any $\omega^{k-k_0+1}, \dots, \omega^k \in \mathfrak{X}^{0,1}(\gamma)$ and $Y_{l-l_0+1}, \dots, Y_l \in \mathfrak{X}(\gamma)$,

$$[F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:l_0}](\omega^{k-k_0+1}, \dots, \omega^k, Y_{l-l_0+1}, \dots, Y_l) = F(\omega^1, \dots, \omega^k, Y_1, \dots, Y_l). \quad (3.88)$$

Proposition 3.27. *Let M be a smooth manifold with or without border, ∇ be a connection on TM and $\gamma : I \rightarrow M$ a smooth curve defined on an interval I . Then, for any $k, l \in \mathbb{N}$, ∇ determines an operator $D_t : \mathfrak{X}^{k,l}(\gamma) \rightarrow \mathfrak{X}^{k,l}(\gamma)$, satisfying the following conditions.*

- (a) *On $\mathfrak{X}(\gamma)$, D_t is the usual covariant derivative along γ , see [Lee19, Theorem 4.24].*
- (b) *On $\mathfrak{X}^{0,0}(\gamma)$, D_t is the usual derivative for real functions, i.e. for any $f \in \mathfrak{X}^{0,0}(\gamma)$, $D_t f = df/dt$.*
- (c) *For any $F \in \mathfrak{X}^{k,l}(\gamma)$, any $\omega^1, \dots, \omega^k \in \mathfrak{X}^{0,1}(\gamma)$ and any $Y_1, \dots, Y_l \in \mathfrak{X}(\gamma)$,*

$$\begin{aligned} (D_t F) (\omega^1, \dots, \omega^k, Y_1, \dots, Y_l) &= \frac{d}{dt} \left[F (\omega^1, \dots, \omega^k, Y_1, \dots, Y_l) \right] \\ &\quad - \sum_{i=1}^k F (\omega^1, \dots, \omega^{i-1}, D_t \omega^i, \omega^{i+1}, \dots, \omega^k, Y_1, \dots, Y_l) \\ &\quad - \sum_{j=1}^l F (\omega^1, \dots, \omega^k, Y_1, \dots, Y_{j-1}, D_t Y_j, Y_{j+1}, \dots, Y_l). \end{aligned} \quad (3.89)$$

In particular, D_t satisfies these additional properties.

- (i) *D_t satisfies the product rule, i.e. for any $f \in \mathfrak{X}^{0,0}(\gamma)$, $F \in \mathfrak{X}^{k,l}(\gamma)$,*

$$D_t (fF) = \left(\frac{d}{dt} f \right) F + f D_t F.$$

- (ii) *For any $k_1, l_1, k_2, l_2 \in \mathbb{N}$, and any $F \in \mathfrak{X}^{k_1, l_1}(\gamma)$, $G \in \mathfrak{X}^{k_2, l_2}(\gamma)$,*

$$D_t (F \otimes G) = D_t F \otimes G + F \otimes D_t G.$$

- (iii) *For any positive integers $k_0 \leq k, l_0 \leq l$, $F \in \mathfrak{X}^{k,l}(\gamma)$,*

$$D_t \left\{ \text{Tr}_{(k_0, l_0)}(F) \right\} = \text{Tr}_{(k_0, l_0)}(D_t F).$$

- (iv) *Let $F \in \mathfrak{X}^{k,l}$ be an extendible tensor field, i.e., such that there exists a (k, l) -tensor field \tilde{F} defined on a neighborhood of $\gamma(I)$ satisfying for any $t \in I$, $F(t) = \tilde{F}(\gamma(t))$. Then, for any $t \in I$,*

$$D_t F(t) = \nabla_{\dot{\gamma}(t)} \tilde{F}(\gamma(t)).$$

Finally, if $\tilde{D}_t : \mathfrak{X}^{k,l}(\gamma) \rightarrow \mathfrak{X}^{k,l}(\gamma)$ is another operator satisfying (a), (b), (i), (ii) and (iii), then $D_t = \tilde{D}_t$.

Proof. Let $k, l \in \mathbb{N}$. Note first that (a)-(b) and (3.89) define $D_t F$ for any $F \in \mathfrak{X}^{k,l}(\gamma)$, setting for any $\omega \in \mathfrak{X}^{0,1}(\gamma)$ and $Y \in \mathfrak{X}(\gamma)$,

$$[D_t \omega](Y) = d[\omega(Y)]/dt - \omega(D_t Y). \quad (3.90)$$

We now show that $D_t F \in \mathfrak{X}^{k,l}$, which will imply that $D_t : \mathfrak{X}^{k,l} \rightarrow \mathfrak{X}^{k,l}$. Second, we establish that (i)-(ii)-(iii)-(iv) are satisfied. We conclude the proof by proving uniqueness of D_t .

Using [Lee19, Lemma B.6], to show that $D_t F \in \mathfrak{X}^{k,l}$ it is enough to prove that $D_t F$ is multilinear over $\mathfrak{X}^{0,0}(\gamma)$. For that, we start proving (i) on $\mathfrak{X}^{0,1}(\gamma)$. Let $\omega \in \mathfrak{X}^{0,1}(\gamma)$, $f \in \mathfrak{X}^{0,0}(\gamma)$ and $Y \in \mathfrak{X}(\gamma)$, then by (3.90),

$$[D_t(f\omega)](Y) = d[f\omega(Y)]/dt - f\omega(D_t Y) = [df/dt]\omega(Y) + f[D_t\omega](Y), \quad (3.91)$$

which proves (i) on $\mathfrak{X}^{0,1}(\gamma)$. Now, let $k, l \in \mathbb{N}$, $F \in \mathfrak{X}^{k,l}(\gamma)$, $\omega^1, \dots, \omega^k \in \mathfrak{X}^{0,1}(\gamma)$, $Y_1, \dots, Y_l \in \mathfrak{X}(\gamma)$. Let $f \in \mathfrak{X}^{0,0}(\gamma)$ and $k_0 \in \mathbb{N}^*$, $k_0 \leq k$. We have, using the multilinearity of F over $\mathfrak{X}^{0,0}(\gamma)$, the definition of D_t (3.89), and (3.91)

$$\begin{aligned} & [D_t F](\omega^1, \dots, \omega^{k_0-1}, f\omega^{k_0}, \omega^{k_0+1}, \dots, \omega^k, Y_1, \dots, Y_l) \\ &= \frac{d}{dt} \left[F(\omega^1, \dots, \omega^{k_0-1}, f\omega^{k_0}, \omega^{k_0+1}, \dots, \omega^k, Y_1, \dots, Y_l) \right] \\ &\quad - \sum_{i=1, i \neq k_0}^k f F(\omega^1, \dots, \omega^{i-1}, D_t \omega^i, \omega^{i+1}, \dots, \omega^k, Y_1, \dots, Y_l) \\ &\quad - F(\omega^1, \dots, \omega^{k_0-1}, D_t(f\omega^{k_0}), \omega^{k_0+1}, \dots, \omega^k, Y_1, \dots, Y_l) \\ &\quad - \sum_{j=1}^l f F(\omega^1, \dots, \omega^k, Y_1, \dots, Y_{j-1}, D_t Y_j, Y_{j+1}, \dots, Y_l) \\ &= \left[\frac{d}{dt} f \right] \left\{ F(\omega^1, \dots, \omega^k, Y_1, \dots, Y_k) - F(\omega^1, \dots, \omega^k, Y_1, \dots, Y_k) \right\} \\ &\quad + f [D_t F](\omega^1, \dots, \omega^k, Y_1, \dots, Y_l) \\ &= f [D_t F](\omega^1, \dots, \omega^k, Y_1, \dots, Y_l). \end{aligned}$$

The same arguments apply if we replace Y_{l_0} with fY_{l_0} , for some $l_0 \leq l$. Thus, using [Lee19, Lemma B.6], $D_t F \in \mathfrak{X}^{k,l}$.

Next, regarding (i), using the definition of D_t ,

$$\begin{aligned} [D_t f F](\omega^1, \dots, \omega^k, Y_1, \dots, Y_l) &= \left[\frac{d}{dt} f \right] F(\omega^1, \dots, \omega^k, Y_1, \dots, Y_l) \\ &\quad + f [D_t F](\omega^1, \dots, \omega^k, Y_1, \dots, Y_l), \end{aligned}$$

thus proving (i). Moreover, we prove (ii). Let $k_1, l_1, k_2, l_2 \in \mathbb{N}$ and $F \in \mathfrak{X}^{k_1, l_1}(\gamma), G \in \mathfrak{X}^{k_2, l_2}(\gamma), \omega^1, \dots, \omega^{k_1+k_2} \in \mathfrak{X}^{0,1}(\gamma), Y_1, \dots, Y_{l_1+l_2} \in \mathfrak{X}(\gamma)$. Setting

$$f = F(\omega^1, \dots, \omega^{k_1}, Y_1, \dots, Y_{l_1}) \text{ and } g = G(\omega^{k_1+1}, \dots, \omega^{k_1+k_2}, Y_{l_1+1}, \dots, Y_{l_1+l_2}),$$

we have

$$\begin{aligned} & [\mathbf{D}_t(F \otimes G)](\omega^1, \dots, \omega^{k_1+k_2}, Y_1, \dots, Y_{l_1+l_2}) \\ &= \frac{d}{dt} [fg] - \left[\sum_{i=1}^{k_1} F(\omega^1, \dots, \omega^{i-1}, \mathbf{D}_t \omega^i, \omega^{i+1}, \dots, \omega^{k_1}, Y_1, \dots, Y_{l_1}) \right. \\ & \quad \left. + \sum_{j=1}^{l_1} F(\omega^1, \dots, \omega^{k_1}, Y_1, \dots, Y_{j-1}, \mathbf{D}_t Y_j, Y_{j+1}, \dots, Y_{l_1}) \right] g \\ & \quad - f \left[\sum_{i=1}^{k_2} G(\omega^{k_1+1}, \dots, \omega^{k_1+i-1}, \mathbf{D}_t \omega^{k_1+i}, \omega^{k_1+i+1}, \dots, \omega^{k_1+k_2}, Y_{l_1+1}, \dots, Y_{l_1+l_2}) \right. \\ & \quad \left. + \sum_{j=1}^{l_2} G(\omega^{k_1+1}, \dots, \omega^{k_1+k_2}, Y_{l_1+1}, \dots, Y_{l_1+j-1}, \mathbf{D}_t Y_{l_1+j}, Y_{l_1+j+1}, \dots, Y_{l_1+l_2}) \right] \\ &= [\mathbf{D}_t F](\omega^1, \dots, \omega^{k_1}, Y_1, \dots, Y_{l_1}) g + f [\mathbf{D}_t G](\omega^{k_1+1}, \dots, \omega^{k_1+k_2}, Y_{l_1+1}, \dots, Y_{l_1+l_2}) \\ &= [\mathbf{D}_t F \otimes G + F \otimes \mathbf{D}_t G](\omega^1, \dots, \omega^{k_1+k_2}, Y_1, \dots, Y_{l_1+l_2}), \end{aligned}$$

which proves (ii). Furthermore, to prove (iii), let $t_0 \in I$ and $(\mathbf{b}_i)_{i \in \{1, \dots, d\}}$ be a basis of $\mathbf{T}_{\gamma(t_0)}\Theta$. Using (a) and [Lee19, Theorem 4.32], define for any $i \in \{1, \dots, d\}$ and $t \in I$,

$$e_i(t) = \mathbf{T}_{t_0, t}^\gamma \mathbf{b}_i,$$

where $\mathbf{T}_{t_0, t}^\gamma$ denotes the parallel transport map along γ from $\mathbf{T}_{\gamma(t_0)}\Theta$ to $\mathbf{T}_{\gamma(t)}\Theta$. As the parallel transport map is an isomorphism, $(e_i(t))_{i \in \{1, \dots, d\}}$ is a basis of $\mathbf{T}_{\gamma(t)}\Theta$, for any $t \in I$. Therefore the family of smooth vector fields $(e_i)_{i \in \{1, \dots, d\}}$ is a parallel frame along γ (with respect to ∇). Denote $(\varepsilon^j)_{j \in \{1, \dots, d\}}$ its dual coframe. Using (3.90) on $Y = e_i, \omega = \varepsilon^j$, for any $i, j \in \{1, \dots, d\}$, shows that the coframe $(\varepsilon^j)_{j \in \{1, \dots, d\}}$ is parallel along γ . Note that for $(e_i)_{i \in \{1, \dots, d\}}$ and $(\varepsilon^j)_{j \in \{1, \dots, d\}}$ to be well defined, we have used ∇ , as well as the operator \mathbf{D}_t on $\mathfrak{X}(\gamma)$ and $\mathfrak{X}^{0,1}(\gamma)$.

Let $k, l \in \mathbb{N}^*$ such that $k_0 \leq k, l_0 \leq l$, and let $F \in \mathfrak{X}^{k, l}(\gamma)$. There exist a family of functions $\{F_{i_1, \dots, i_k}^{j_1, \dots, j_l} \in \mathfrak{X}^{0,0}(\gamma) : i_1, \dots, i_k, j_1, \dots, j_l \in \{1, \dots, d\}\}$ such that

$$F = \sum_{i_1, \dots, i_k=1}^d \sum_{j_1, \dots, j_l=1}^d F_{i_1, \dots, i_k}^{j_1, \dots, j_l} \bigotimes_{\Delta=1}^k e_{i_\Delta} \bigotimes_{\square=1}^l \varepsilon^{j_\square}.$$

Since the frame and its dual coframe are parallel along γ , for any $i \in \{1, \dots, d\}$ $\mathbf{D}_t e_i = 0$ and $\mathbf{D}_t \varepsilon^i = 0$. Combining this fact with (i) and (ii) gives

$$\mathbf{D}_t F = \sum_{i_1, \dots, i_k=1}^d \sum_{j_1, \dots, j_l=1}^d \left[\frac{d}{dt} F_{i_1, \dots, i_k}^{j_1, \dots, j_l} \right] \bigotimes_{\Delta=1}^k e_{i_\Delta} \bigotimes_{\square=1}^l \varepsilon^{j_\square}. \quad (3.92)$$

Let $k_0, l_0 \in \mathbb{N}^*$ such that $k_0 \leq k, l_0 \leq l$, then by definition of $\text{Tr}_{(k_0, l_0)}$, for any $i_1, \dots, i_{k-1}, j_1, \dots, j_{l-1} \in \{1, \dots, d\}$,

$$\text{Tr}_{(k_0, l_0)}(F)_{i_1, \dots, i_{k-1}}^{j_1, \dots, j_{l-1}} = \sum_{m=1}^d F_{i_1, \dots, i_{k_0-1}, m, i_{k_0}, \dots, i_{k-1}}^{j_1, \dots, j_{l_0-1}, m, j_{l_0}, \dots, j_{l-1}}. \quad (3.93)$$

We remind the reader that $\text{Tr}_{(k_0, l_0)}(F)$ does not depend on the choice of coordinates [Lee19, Appendix B]. Thus, using (3.92) and (3.93), we have

$$\begin{aligned} D_t \left[\text{Tr}_{(k_0, l_0)}(F) \right] &= \sum_{i_1, \dots, i_{k-1}=1}^d \sum_{j_1, \dots, j_{l-1}=1}^d \frac{d}{dt} \left[\text{Tr}_{(k_0, l_0)}(F)_{i_1, \dots, i_{k-1}}^{j_1, \dots, j_{l-1}} \right] \bigotimes_{\Delta=1}^{k-1} e_{i_\Delta} \bigotimes_{\square=1}^{l-1} \varepsilon^{j_\square} \\ &= \sum_{i_1, \dots, i_{k-1}=1}^d \sum_{j_1, \dots, j_{l-1}=1}^d \sum_{m=1}^d \frac{d}{dt} F_{i_1, \dots, i_{k_0-1}, m, i_{k_0}, \dots, i_{k-1}}^{j_1, \dots, j_{l_0-1}, m, j_{l_0}, \dots, j_{l-1}} \bigotimes_{\Delta=1}^{k-1} e_{i_\Delta} \bigotimes_{\square=1}^{l-1} \varepsilon^{j_\square} \\ &= \text{Tr}_{(k_0, l_0)}(D_t F), \end{aligned}$$

thus proving (iii).

To prove (iv), first for any $f \in \mathfrak{X}^{(0,0)}(\gamma)$, extendible in \tilde{f} , we have by composition and definition of the covariant derivative, that for any $t \in [0, 1]$,

$$(df/dt)(t) = d\tilde{f}_{\gamma(t)}(\dot{\gamma}(t)) = \nabla_{\dot{\gamma}(t)} \tilde{f}(\gamma(t)). \quad (3.94)$$

Also, using [Lee19, Theorem 4.24-(iii)] gives (iv) for any $Y \in \mathfrak{X}(\gamma)$. Combining (3.94), (3.90), its counterpart for tensor fields defined over a manifold [Lee19, Proposition 4.15-(a)] and (iv) over $\mathfrak{X}(\gamma)$, proves (iv) over $\mathfrak{X}^{(0,1)}(\gamma)$. Now, for any $k, l \in \mathbb{N}$, using (iv) over $\mathfrak{X}(\gamma)$ and $\mathfrak{X}^{(0,1)}(\gamma)$ combined with (3.89) and its counterpart for tensor fields defined over a manifold [Lee19, Equation (4.12)] gives (iv) over $\mathfrak{X}^{(k,l)}(\gamma)$.

Finally, we address uniqueness. Suppose now that \tilde{D}_t is an operator on $\mathfrak{X}^{k,l}(\gamma)$ that satisfies (a), (b), (i), (ii) and (iii). First, (a) and (b) show that D_t and \tilde{D}_t coincide on $\mathfrak{X}^{0,0}(\gamma)$ and $\mathfrak{X}(\gamma)$. Second, for any $Y \in \mathfrak{X}(\gamma), \omega \in \mathfrak{X}^{0,1}(\gamma)$, writing $\omega(Y) = \text{Tr}_{(1,1)}(Y \otimes \omega)$ and using (iii) gives

$$\tilde{D}_t \omega = d[\omega(Y)]/dt - \omega(\tilde{D}_t Y) = D_t \omega,$$

using (3.90). Thus, \tilde{D}_t and D_t also agree on $\mathfrak{X}^{0,1}(\gamma)$. Therefore, the frame $(e_i)_{i \in \{1, \dots, d\}}$ and its dual coframe $(\varepsilon^j)_{j \in \{1, \dots, d\}}$ are also parallel with respect to \tilde{D}_t along γ . Let $F \in \mathfrak{X}^{k,l}(\gamma)$, then using (i) and (ii) shows that (3.92) holds for the operator \tilde{D}_t , proving that $D_t F = \tilde{D}_t F$. This concludes the proof. \square

Lemma 3.28. *Let M be a smooth manifold and ∇ be a connection on TM. Let $\gamma : [0, 1] \rightarrow M$ be a smooth curve and denote D_t the covariant derivative operator along γ associated with ∇ , defined in Proposition 3.27. Let $F \in \mathfrak{X}^{k,l}(\gamma), \omega^1, \dots, \omega^{k_0} \in \mathfrak{X}^{0,1}(\gamma)$ and*

$Y_1, \dots, Y_{l_0} \in \mathfrak{X}(\gamma)$, with $k_0 \leq k, l_0 \leq l$. Then, we have

$$\begin{aligned} \mathbb{D}_t \left([F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:l_0}] \right) &= [\mathbb{D}_t F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:l_0}] \\ &+ \sum_{i=1}^{k_0} [F : \omega^{\otimes 1:(i-1)} \otimes \mathbb{D}_t \omega^i \otimes \omega^{\otimes (i+1):k_0} \otimes Y_{\otimes 1:l_0}] \\ &+ \sum_{j=1}^{l_0} [F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:(j-1)} \otimes \mathbb{D}_t Y_j \otimes Y_{\otimes (j+1):l_0}]. \end{aligned} \quad (3.95)$$

Proof. Let F be a smooth (k, l) -tensor field along γ . We show (3.95) by induction. Following the recursive definition of the contraction in (3.86), we prove it by induction on $k_0 \in \mathbb{N}^*, k_0 \leq k$, for any $\omega^1, \dots, \omega^{k_0} \in \mathfrak{X}^{0,1}(\gamma)$.

The case $k_0 = 1$ follows from Proposition 3.27-(ii) and (iii), combined with the definition in (3.86),

$$\begin{aligned} \mathbb{D}_t [F : \omega^1] &= \mathbb{D}_t \operatorname{Tr}_{(1,l+1)}(F \otimes \omega^1) \\ &= \operatorname{Tr}_{(1,l+1)}(\mathbb{D}_t [F \otimes \omega^1]) \\ &= \operatorname{Tr}_{(1,l+1)}(\mathbb{D}_t F \otimes \omega^1 + F \otimes \mathbb{D}_t \omega^1) \\ &= [\mathbb{D}_t F : \omega^1] + [F : \mathbb{D}_t \omega^1], \end{aligned}$$

where we have used the linearity of Tr . Now assume there exists $k_0 \in \{1, \dots, k-1\}$ such that (3.95) holds for any smooth 1 forms $\omega^1, \dots, \omega^{k_0}$ and $l_0 = 0$. Moreover, consider any smooth 1 forms $\omega^1, \dots, \omega^{k_0+1}$. Then, using the same arguments as for the case $k_0 = 1$ and the induction hypothesis, we obtain

$$\begin{aligned} \mathbb{D}_t [F : \omega^{\otimes 1:(k_0+1)}] &= \mathbb{D}_t \operatorname{Tr}_{(1,l+1)} \left([F : \omega^{\otimes 1:k_0}] \otimes \omega^{k_0+1} \right) \\ &= \operatorname{Tr}_{(1,l+1)} \left(\mathbb{D}_t [F : \omega^{\otimes 1:k_0}] \otimes \omega^{k_0+1} \right) + \operatorname{Tr}_{(1,l+1)} \left([F : \omega^{\otimes 1:k_0}] \otimes \mathbb{D}_t \omega^{k_0+1} \right) \\ &= \operatorname{Tr}_{(1,l+1)} \left([\mathbb{D}_t F : \omega^{\otimes 1:k_0}] \otimes \omega^{k_0+1} \right) + [F : \omega^{\otimes 1:k_0} \otimes \mathbb{D}_t \omega^{k_0+1}] \\ &\quad + \sum_{i=1}^{k_0} \operatorname{Tr}_{(1,l+1)} \left([F : \omega^{\otimes 1:(i-1)} \otimes \mathbb{D}_t \omega^i \otimes \omega^{\otimes (i+1):k_0}] \otimes \omega^{k_0+1} \right) \\ &= [\mathbb{D}_t F : \omega^{\otimes 1:(k_0+1)}] + \sum_{i=1}^{k_0+1} [F : \omega^{\otimes 1:(i-1)} \otimes \mathbb{D}_t \omega^i \otimes \omega^{\otimes (i+1):(k_0+1)}]. \end{aligned}$$

Subsequently, using the recursive definition of the contraction in (3.87), we prove (3.95) by induction on $l_0 \in \mathbb{N}^*, l_0 \leq l$ for any $k_0 \leq k$ and any $\omega^1, \dots, \omega^{k_0} \in \mathfrak{X}^{0,1}(\gamma)$. Let $Y_1 \in \mathfrak{X}(\gamma)$. Then, using once again Proposition 3.27-(ii) and (iii), (3.87), and (3.95) in the case $l_0 = 0$

justified above, the case $l_0 = 1$ is proven as follows,

$$\begin{aligned}
 D_t \left[F : \omega^{\otimes 1:k_0} \otimes Y_1 \right] &= \text{Tr}_{(k-k_0+1,1)} \left(D_t \left\{ \left[F : \omega^{\otimes 1:k_0} \right] \otimes Y_1 \right\} \right) \\
 &= \text{Tr}_{(k-k_0+1,1)} \left(\left[D_t F : \omega^{\otimes 1:k_0} \right] \otimes Y_1 \right) + \left[F : \omega^{\otimes 1:k_0} \otimes D_t Y_1 \right] \\
 &\quad + \sum_{i=1}^{k_0} \text{Tr}_{(k-k_0+1,1)} \left(\left[F : \omega^{\otimes 1:(i-1)} \otimes D_t \omega^i \otimes \omega^{\otimes (i+1):k_0} \right] \otimes Y_1 \right) \\
 &= \left[D_t F : \omega^{\otimes 1:k_0} \otimes Y_1 \right] + \sum_{i=1}^{k_0} \left[F : \omega^{\otimes 1:(i-1)} \otimes D_t \omega^i \otimes \omega^{\otimes (i+1):k_0} \otimes Y_1 \right] \\
 &\quad + \left[F : \omega^{\otimes 1:k_0} \otimes D_t Y_1 \right] .
 \end{aligned}$$

Furthermore, assume there exists $l_0 \in \{1, \dots, l-1\}$ such that (3.95) holds for any $k_0 \leq k$, any $\omega^1, \dots, \omega^{k_0} \in \mathfrak{X}^{0,1}(\gamma)$ and any $Y_1, \dots, Y_{l_0} \in \mathfrak{X}(\gamma)$. Let $Y_1, \dots, Y_{l_0+1} \in \mathfrak{X}(\gamma)$. Then using the same arguments as for the case $l_0 = 1$ and the induction hypothesis, we obtain

$$\begin{aligned}
 D_t \left[F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:(l_0+1)} \right] &= \text{Tr}_{(k-k_0+1,1)} \left(D_t \left\{ \left[F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:l_0} \right] \otimes Y_{l_0+1} \right\} \right) \\
 &= \text{Tr}_{(k-k_0+1,1)} \left(\left[D_t F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:l_0} \right] \otimes Y_{l_0+1} \right) \\
 &\quad + \sum_{i=1}^{k_0} \text{Tr}_{(k-k_0+1,1)} \left(\left[F : \omega^{\otimes 1:(i-1)} \otimes D_t \omega^i \otimes \omega^{\otimes (i+1):k_0} \otimes Y_{\otimes 1:l_0} \right] \otimes Y_{l_0+1} \right) \\
 &\quad + \sum_{j=1}^{l_0} \text{Tr}_{(k-k_0+1,1)} \left(\left[F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:(j-1)} \otimes D_t Y_j \otimes Y_{\otimes (j+1):l_0} \right] \otimes Y_{l_0+1} \right) \\
 &\quad + \text{Tr}_{(k-k_0+1,1)} \left(\left[F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:l_0} \right] \otimes D_t Y_{l_0+1} \right) \\
 &= \left[D_t F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:(l_0+1)} \right] + \sum_{i=1}^{k_0} \left[F : \omega^{\otimes 1:(i-1)} \otimes D_t \omega^i \otimes \omega^{\otimes (i+1):k_0} \otimes Y_{\otimes 1:(l_0+1)} \right] \\
 &\quad + \sum_{j=1}^{l_0+1} \left[F : \omega^{\otimes 1:k_0} \otimes Y_{\otimes 1:(j-1)} \otimes D_t Y_j \otimes Y_{\otimes (j+1):(l_0+1)} \right] ,
 \end{aligned}$$

which concludes the proof. \square

Theorem 3.29. *Let M be a smooth manifold and ∇ be a connection on TM . Let $\gamma : [0, 1] \rightarrow M$ be a geodesic and $Y : M \rightarrow TM$ a smooth vector field. Then, for any $t \in [0, 1], n \in \mathbb{N}$,*

$$\begin{aligned}
 T_{t_0}^Y Y(\gamma(t)) &= \sum_{k_0=0}^n (t^{k_0}/k_0!) \nabla^{k_0} Y_{\gamma(0)} (\dot{\gamma}(0), \dots, \dot{\gamma}(0)) \\
 &\quad + \int_0^t [(t-s)^n/n!] T_{s_0}^Y \nabla^{n+1} Y_{\gamma(s)} (\dot{\gamma}(s), \dots, \dot{\gamma}(s)) ds ,
 \end{aligned} \tag{3.96}$$

where $T_{t_0}^\gamma : T_{\gamma(t)}\mathbf{M} \rightarrow T_{\gamma(0)}\mathbf{M}$ is the parallel transport map along γ , and the $(1, k_0)$ -tensor field $\nabla^{k_0}Y$ is the total derivative of order k_0 of the $(1, 0)$ -tensor field Y .

For a definition of the total covariant derivative, see [Lee19, Proposition 4.15]. Also, in (3.96), remark that even though $\dot{\gamma}$ is only a vector field along γ , and not a vector field, the value of a vector field $\nabla_X Y$ evaluated at $\theta \in \mathbf{M}$ only depends on $X(\theta)$ and on values of Y along smooth curves $c : [0, 1] \rightarrow \mathbf{M}$ satisfying $c(0) = \theta$ and $\dot{c}(0) = X(\theta)$; by [Lee19, Proposition 4.26]. Therefore the expression $\nabla^{k_0}Y_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t))$ in Theorem 3.29 is well defined for any $k_0 \in \mathbb{N}, t \in [0, 1]$.

Proof. Consider $\mathcal{V} : [0, 1] \rightarrow \mathbf{M}$ the smooth vector field along γ and the function $\varphi : [0, 1] \rightarrow T_{\gamma(0)}\mathbf{M}$ defined by

$$\mathcal{V} = Y \circ \gamma \text{ and } \varphi : t \mapsto T_{t_0}^\gamma \mathcal{V}(t). \quad (3.97)$$

Then we check by induction on $n \in \mathbb{N}^*$ that φ is n -times differentiable with derivative of order n given for any $t \in [0, 1]$ by $\varphi^{(n)}(t) = T_{t_0}^\gamma [D_t^n \mathcal{V}(t)]$ and $D_t^n \mathcal{V}(t) = \nabla^n Y_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t))$, where D_t is the covariant derivative operator along γ with respect to the connection ∇ , defined in Proposition 3.27.

First, the case $n = 1$ is a direct application of [Lee19, Theorem 4.34, Theorem 4.24] since Y is an extension of \mathcal{V} . Assume now that the property holds for $n \in \mathbb{N}^*$. Then, for any $t_0, t \in [0, 1], t \neq t_0$, we have

$$\left[\varphi^{(n)}(t) - \varphi^{(n)}(t_0) \right] / (t - t_0) = T_{t_0}^\gamma \left[T_{t_0}^\gamma D_t^n \mathcal{V}(t) - D_t^n \mathcal{V}(t_0) \right] / (t - t_0).$$

Now [Lee19, Theorem 4.34] ensures that the limit of the quantity above exists when $t \rightarrow t_0$ and in addition this limit is

$$\varphi^{(n+1)}(t_0) = T_{t_0}^\gamma D_t^{n+1} \mathcal{V}(t_0),$$

which shows that φ is $n + 1$ times differentiable on $[0, 1]$. We now show that for any $t \in [0, 1]$, $D_t^{n+1} \mathcal{V}(t) = \nabla^{n+1} Y_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t))$. Using Lemma 3.28 on the smooth $(1, n)$ -tensor field along γ $F = (\nabla^n Y) \circ \gamma$, taking $k_0 = 0$ and n times the vector field $\dot{\gamma}$, we have

$$D_t [F : \dot{\gamma} \otimes \dots \otimes \dot{\gamma}] = [D_t F : \dot{\gamma} \otimes \dots \otimes \dot{\gamma}],$$

since $D_t \dot{\gamma} = 0$ because γ is a geodesic. Also, by (3.88), $[D_t F : \dot{\gamma} \otimes \dots \otimes \dot{\gamma}] = D_t F(\dot{\gamma}, \dots, \dot{\gamma})$. Finally, as $\nabla^n Y$ is an extension of F , using the induction hypothesis and the definition of the total derivative give for any $t \in [0, 1]$,

$$\begin{aligned} D_t^{n+1} \mathcal{V}(t) &= D_t F(\dot{\gamma}, \dots, \dot{\gamma})(t) = \nabla_{\dot{\gamma}(t)} (\nabla^n Y)_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t)) \\ &= (\nabla^{n+1} Y)_{\gamma(t)}(\dot{\gamma}(t), \dots, \dot{\gamma}(t)), \end{aligned}$$

concluding the induction.

Finally, (3.96) is simply a consequence of Taylor's formula with integral remainder of the vectorial valued function φ identifying $T_{\gamma(0)}\mathbf{M}$ with \mathbb{R}^d . \square

Proposition 3.30. *Let M be a smooth manifold, ∇ be a symmetric connection defined over the smooth vector fields of M . For any smooth function $f : M \rightarrow \mathbb{R}$ and any local coordinates $(u_i)_{i \in \{1, \dots, d\}}$, we have*

$$\begin{aligned} \nabla \text{Hess } f = \sum_{i,j,k=1}^d \left\{ \partial_{kij}^3 f - \sum_{l=1}^d \left[\Gamma_{ij}^l \partial_{kl}^2 f + \Gamma_{ki}^l \partial_{jl}^2 f + \Gamma_{kj}^l \partial_{il}^2 f \right] - \sum_{m=1}^d \partial_k \Gamma_{ij}^m \partial_m f \right. \\ \left. + \sum_{l,m=1}^d \left[\Gamma_{kj}^l \Gamma_{il}^m + \Gamma_{ki}^l \Gamma_{lj}^m \right] \partial_m f \right\} du^i \otimes du^j \otimes du^k, \end{aligned}$$

where $(\Gamma_{ij}^k)_{i,j,k \in \{1, \dots, d\}}$ are the Christoffel symbols in these local coordinates, the local frame and its dual coframe are denoted by $(\partial u_i)_{i \in \{1, \dots, d\}}$ and $(du^j)_{j \in \{1, \dots, d\}}$.

Proof. Let $(u_i)_{i \in \{1, \dots, d\}}$ be local coordinates. By [Lee19, Example 4.22], in this chart, we have

$$\text{Hess } f = \sum_{i,j=1}^d F_{ij} du^i \otimes du^j, \text{ where for any } i, j \in \{1, \dots, d\}, F_{ij} = \partial_{ij}^2 f - \sum_{m=1}^d \Gamma_{ij}^m \partial_m f. \quad (3.98)$$

Applying [Lee19, Proposition 4.18] on $\text{Hess } f$, we obtain that $\nabla \text{Hess } f = \sum_{i,j,k=1}^d G_{ijk} du^i \otimes du^j \otimes du^k$, where for any $i, j, k \in \{1, \dots, d\}$,

$$G_{ijk} = \partial_k F_{ij} - \sum_{l=1}^d \left(\Gamma_{kj}^l F_{il} + \Gamma_{ki}^l F_{lj} \right).$$

Expanding the expression above using (3.98) gives for any $i, j, k \in \{1, \dots, d\}$,

$$\begin{aligned} G_{ijk} = \partial_{ijk}^3 f - \sum_{m=1}^d \left(\partial_k \Gamma_{ij}^m \partial_m f + \Gamma_{ij}^m \partial_{km}^2 f \right) - \sum_{l=1}^d \Gamma_{kj}^l \left(\partial_{il}^2 f - \sum_{m=1}^d \Gamma_{il}^m \partial_m f \right) \\ - \sum_{l=1}^d \Gamma_{ki}^l \left(\partial_{lj}^2 f - \sum_{m=1}^d \Gamma_{lj}^m \partial_m f \right). \end{aligned}$$

The desired result is obtained by reordering this equation, which concludes the proof. \square

CHAPTER 4

Optimal Scaling Results for a Wide Class of Proximal MALA Algorithms

FRANCESCA R. CRUCINIO¹, ALAIN DURMUS², PABLO JIMÉNEZ², GARETH O. ROBERTS³

We consider a recently proposed class of MCMC methods which uses proximity maps instead of gradients to build proposal mechanisms which can be employed for both differentiable and non-differentiable targets. These methods have been shown to be stable for a wide class of targets, making them a valuable alternative to Metropolis-adjusted Langevin algorithms (MALA); and have found wide application in imaging contexts. The wider stability properties are obtained by building the Moreau-Yoshida envelope for the target of interest, which depends on a parameter λ . In this work, we investigate the optimal scaling problem for this class of algorithms, which encompasses MALA, and provide practical guidelines for the implementation of these methods.

Contents

4.1	Introduction	169
4.2	Proximal MALA Algorithms	171
4.3	Optimal scaling of Proximal MALA	174
4.3.1	Differentiable targets	175
4.3.2	Laplace target	178
4.4	Practical Implications and Numerical Simulations	181
4.4.1	Numerical Experiments	182
4.5	Discussion	184
4.6	Proof of the Result for the Laplace distribution	185

¹CREST, ENSAE Paris

²Centre de Mathématiques Appliquées, Ecole Polytechnique, France, Institut Polytechnique de Paris

³Department of Statistics, University of Warwick

4.6.1	Proof of Theorem 4.4	186
4.6.2	Proof of Proposition 4.5	188
4.6.3	Proof of Proposition 4.6	195
4.6.4	Proof of Theorem 4.7	198
4.7	Proof of Theorem 4.1	206
4.7.1	Auxiliary Results for the Proof of Case (a)	208
4.7.2	Auxiliary Results for the Proof of Case (b)	216
4.7.3	Auxiliary Results for the Proof of Case (c)	219
4.7.4	Proof of Theorem 4.1	222
4.8	Numerical Experiments	223
4.8.1	Differentiable Targets	223
4.8.2	Laplace Target	223
4.8.3	Mix of a Laplace and differentiable target	228
4.9	Taylor Expansions for the Results on Differentiable Targets	231
4.9.1	Coefficients of the Taylor Expansion	231
4.9.2	Taylor Expansions of the Log-acceptance Ratio	237
4.9.3	Derivatives of the Proximity Map for Differentiable Targets	242
4.10	Moments and Integrals for the Laplace Distribution	244
4.10.1	Moments of Acceptance Ratio for the Laplace Distribution	244
4.10.2	Bound on Second Moment of Acceptance Ratio for the Laplace Distribution	250
4.10.3	Additional Integrals for the Laplace Distribution	254
4.10.4	Integrals for Moment Computations	255

4.1 Introduction

Gradient-based Markov chain Monte Carlo (MCMC) methods have proved to be very successful at sampling from high-dimensional target distributions [Bro+11]. The key to their success is that in many cases their mixing time appears to be better than their competitor algorithms which do not use gradient information (see for example [RR98b]), while their implementation has similar computational cost. Indeed, gradients of target densities can often be computed with computational complexity (in dimension d) which scales no worse than evaluation of the target density itself.

Gradient-based MCMC methods are mainly motivated from stochastic processes constructed to have the target density as limiting distribution [BFR19; BVD18; MCF15; VSS21]. Our analysis will concentrate on Metropolis Adjusted Langevin Algorithm (MALA) and its proximal variants which are based on the Langevin diffusion

$$d\mathbf{L}_t = d\mathbf{B}_t + \frac{\nabla \log \pi(\mathbf{L}_t)}{2} dt, \quad (4.1)$$

where π denotes the target density with respect to the Lebesgue measure and $(\mathbf{B}_t)_{t \geq 0}$ a standard Brownian motion. It is well-known that under appropriate conditions, (4.1) defines a continuous-time Markov process associated with a Markov semigroup which is reversible with respect to π . From this observation, it has been suggested to use a Euler-Maruyama (EM) approximation of (4.1). This scheme has been popularized in statistics by [GM94b] and referred to as the *Unadjusted Langevin Algorithm (ULA)* in [RT96c]. Due to the time-discretization, ULA does not have π as stationary distribution. To address this problem, [RDF78] and independently Besag in his contribution to [GM94b] proposed to add a Metropolis acceptance step at each iteration of the EM scheme, leading to the Metropolis Adjusted Langevin Algorithm (MALA) following [RT96c] who also derive basic stability analysis. The accept/reject step in this algorithm confers two significant advantages: it ensures that the resulting algorithm has exactly the correct invariant distribution, while step sizes can be chosen larger than in the unadjusted case as there is not need to make the step size small to reduce discretization error. On the other hand, MALA algorithms are typically hard to analyze theoretically (see e.g. [BH13; DM22; Dwi+18]). However, [RR98b] (see also [BRS09; PST12]) have established that MALA has better convergence properties than the Random Walk Metropolis (RWM) algorithm with respect to the dimension d from an optimal scaling perspective (see also [RGG97b]).

Whereas gradient-based methods have been successively applied and offer interesting features, they are typically less robust than their vanilla alternatives (for example see [RT96c]); while intuition suggests, and existing underpinning theory requires, that target densities need to be sufficiently smooth for the gradients to be aiding Markov chain convergence. Moreover, while gradient-based MCMC have been successful for smooth densities, there is no reason to believe that they should be effective for densities which are not differentiable at a subset $X \subseteq \mathbb{R}^d$. For non-smooth densities, [Per16] proposes modified gradient-based algorithms. Their proposed P-MALA algorithm is inspired by the proximal algorithms popular in the optimization literature (e.g. [PB14]). The main idea is to approximate the (possibly non differentiable but) log-concave target density $\pi \propto \exp(-G)$

by substituting the potential G with its Moreau-Yoshida envelope G^λ (see (4.3) below for its definition), to obtain a distribution π^λ whose level of smoothness is controlled by the proximal parameter $\lambda > 0$. Given this smooth approximation to π one can then build proposals based on time discretizations of the Langevin diffusion targeting π^λ [DMP18; Per16]:

$$\xi_{k+1} = \xi_k - \frac{\sigma^2}{2} \nabla G^\lambda(\xi_k) + \sigma Z_{k+1}, \quad (4.2)$$

where $\sigma^2 > 0$ is a fixed stepsize and $(Z_k)_{k \in \mathbb{N}^*}$ is a sequence of i.i.d. zero-mean Gaussian random variables with identity covariance matrix. Our aims in this paper are broadly to provide theoretical underpinning for a slightly larger family of *proximal MALA* algorithms, analyze how these methods scale with dimension, and to give insights and practical guidance into how they should be implemented supported by the theory we establish.

Proximal optimization and MCMC methods proved to be particularly well-suited for image estimation, where penalties involving the sparsity inducing norms are common [DMP18; Per16; VDC19]. Similar targets are also common in sparse regression contexts [Atc15; GSS21; ZCZ22]. In these situations, the set of non-differentiability points for the target density π is a null set for the Lebesgue measure, and, following [Dur+17], we shall focus on this case. However, in contrast to the conclusions of [Dur+17] for RWM, we shall demonstrate that optimal scaling of proximal MALA may be affected by non-smoothness.

More precisely, in this work, we first extend the results of [Pil22] and consider a wider range of proximal MALA algorithms, as well as a wider class of finite dimensional target distributions. We let both the steps size σ^2 and the regularization parameter λ depend on the dimension d of the target and find that the scaling properties of proximal MALA depend on the relative speed at which λ and σ converge to 0 as $d \rightarrow \infty$. We start by considering a class of sufficiently differentiable target distributions π to which MALA can also be applied, to allow direct comparison between MALA and proximal MALA and thus between a gradient-based method and one which approximates the gradient through proximal operators. When λ goes to 0 at least as fast as σ^2 , we find that the scaling properties of proximal MALA are equivalent to those of MALA (i.e. σ^2 should decay as $d^{-1/3}$; see Theorem 4.1–(b), Theorem 4.1–(c)); when λ converges to 0 more slowly than σ^2 , proximal MALA is less efficient than MALA with σ^2 decaying as $d^{-1/2}$ (Theorem 4.1–(a)). As in some cases the proximal operator for a given distribution π is cheaper to compute than $\nabla \log \pi$ [CP11; PB14; Per16], we anticipate that proximal MALA with an appropriately tuned λ might provide a cheaper alternative to MALA retaining similar scaling properties.

We then turn to the optimal scaling of proximal MALA applied to the Laplace distribution $\pi(x) \propto e^{-|x|}$. We focus on this particular non-smooth target since it is the most widely used in applications of proximal MALA, including image deconvolution [DMP18; Per16; VDC19], LASSO, and sparse regression [Atc15; GSS21; ZCZ22]. We establish that non-differentiability of the target even at one point leads to a different optimal scaling than MALA. In particular, the step size has to scale as $d^{-2/3}$ and not as $d^{-1/3}$ (Theorem 4.4).

This appears to be a new optimal scaling for Metropolis MCMC algorithms which is between the one of RWM and MALA. From the conclusion for smooth target distributions,

we restrict our study to the choice of λ going to 0 at least as fast as σ^2 .

The proof of the result for the differentiable case extends that of [RR98b] for MALA, while the structure of the proof for the Laplace target is similar to that of [Dur+17] and constitutes the main element of novelty in this paper. As a special case of the result for the Laplace distribution, we also obtain the optimal scaling for MALA on Laplace targets. We point out that the strategy adopted in the proof of this result is not unique to the Laplace distribution, and could be applied to other distributions provided that the required integrals can be obtained.

To sum up, our main contributions are:

- 1) We extend the result of [Pil22] beyond the Gaussian case, covering all finite dimensional (sufficiently) differentiable targets, and show that, in some cases, proximal MALA affords the same scaling properties of MALA if the proximal parameter λ is chosen appropriately.
- 2) Motivated by applications in imaging and sparse regression applications, we study the scaling of proximal MALA methods for the Laplace target, and show that for values of λ decaying sufficiently fast, the optimal scaling of proximal MALA, i.e. the choice for σ^2 , is different from the one for MALA on differentiable targets and is of order $d^{-2/3}$.
- 3) We use the insights obtained with the aforementioned results to provide practical guidelines for the selection of the proximal parameter λ .

Organization of the paper The paper is structured as follows. In Section 4.2, we rigorously introduce the class of proximal MALA algorithms that are studied and discuss related works on optimal scaling for MCMC algorithms. In Section 4.3.1 we state the main result for differentiable targets, showing that the scaling properties of proximal MALA depend on the relative speed at which λ goes to 0 with respect to σ . In Section 4.3.2 we obtain a scaling limit for proximal MALA when π is a Laplace distribution, as a special case of our result we also obtain the scaling properties of a sub-gradient version of MALA for this target. We collect in Section 4.4 the main practical takeaways from these results and discuss possible extensions in Section 4.5. Finally, in Section 4.6 we prove the result for the Laplace distribution. The proof of the result for differentiable targets is postponed to Appendix 4.7.

4.2 Proximal MALA Algorithms

We now introduce the general class of proximal MALA algorithms, first studied in [Per16]. This class of algorithms aims at sampling from a density with respect to the Lebesgue measure on \mathbb{R}^d of the form $\pi(\mathbf{x}) = \exp(-G(\mathbf{x})) / \int_{\mathbb{R}^d} \exp(-G(\tilde{\mathbf{x}})) d\tilde{\mathbf{x}}$, with G satisfying the following assumption

A18. *The function $G : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex, proper and lower semi-continuous.*

The main idea behind proximal MALA is to approximate the (possibly non differentiable) target density π by approximating the potential G with its Moreau-Yoshida envelope $G^\lambda : \mathbb{R}^d \rightarrow \mathbb{R}$ defined for $\lambda > 0$ by

$$G^\lambda(\mathbf{x}) = \min_{\mathbf{u} \in \mathbb{R}^d} [G(\mathbf{u}) + \|\mathbf{u} - \mathbf{x}\|^2 / (2\lambda)] . \quad (4.3)$$

Since G is supposed to be convex, by [RW98, Theorem 2.26], the Moreau-Yoshida envelope is well-defined, convex and continuously differentiable with

$$\nabla G^\lambda(\mathbf{x}) = \lambda^{-1}(\mathbf{x} - \text{prox}_G^\lambda(\mathbf{x})), \quad \text{prox}_G^\lambda(\mathbf{x}) = \arg \min_{\mathbf{u} \in \mathbb{R}^d} [G(\mathbf{u}) + \|\mathbf{u} - \mathbf{x}\|^2 / (2\lambda)]. \quad (4.4)$$

The proximity operator $\mathbf{x} \mapsto \text{prox}_G^\lambda(\mathbf{x})$ behaves similarly to a gradient mapping and moves points in the direction of the minimizers of G . In the limit $\lambda \rightarrow 0$ the quadratic penalty dominates (4.4) and the proximity operator coincides with the identity operator, i.e. $\text{prox}_G^\lambda(\mathbf{x}) = \mathbf{x}$; in the limit $\lambda \rightarrow \infty$, the quadratic penalty term vanishes and (4.4) maps all points to the set of minimizers of G .

It was shown in [DMP18, Proposition 1] that, under **A18**, $\int_{\mathbb{R}^d} \exp(-G^\lambda(\mathbf{x})) d\mathbf{x} < \infty$, and therefore the probability density $\pi^\lambda \propto \exp(-G^\lambda)$ is well-defined. In addition, it has been shown that $\|\pi - \pi^\lambda\|_{\text{TV}} \rightarrow 0$ as $\lambda \rightarrow 0$. Based on this observation and since as we have emphasized π^λ is now continuously differentiable, it has been suggested in [DMP18; Per16] to use the discretization of the Langevin diffusion associated with π^λ given by (4.2), which can be rewritten using (4.4) as

$$\xi_{k+1} = \left(1 - \frac{\sigma^2}{2\lambda}\right) \xi_k + \frac{\sigma^2}{2\lambda} \text{prox}_G^\lambda(\xi_k) + \sigma Z_{k+1}. \quad (4.5)$$

Similarly to other MCMC methods based on discretizations of the Langevin diffusion (e.g. [RT96c]), one can build unadjusted schemes which target π^λ , expecting draws from these schemes to be close to draws from π for small enough λ , or add a Metropolis-Hastings step to ensure that the resulting algorithm targets π . Unadjusted proximal MCMC methods have been analyzed in [DMP18]; in this paper we focus on Metropolis adjusted proximal MCMC methods and study their scaling properties. More precisely, at each step k and given the current state of the Markov chain X_k , a candidate Y_{k+1} is generated from the transition density associated to (4.5), $(\mathbf{x}, \mathbf{y}) \mapsto q(\mathbf{x}, \mathbf{y}) = \varphi(\mathbf{y}; [1 - \sigma^2/(2\lambda)]\mathbf{x} + \sigma^2 \text{prox}_G^\lambda(\mathbf{x})/2\lambda, \sigma^2 \mathbf{I}_d)$, where $\varphi(\cdot; \mathbf{u}, \mathbf{\Sigma})$ stands for the d -dimension Gaussian density with mean \mathbf{u} and covariance matrix $\mathbf{\Sigma}$. Given X_k and Y_{k+1} , Then, the next state is set as:

$$X_{k+1} = Y_{k+1} b_{k+1} + X_k (1 - b_{k+1}), \quad b_{k+1} = \mathbb{1}_{\mathbb{R}_+} \left(\frac{\pi(Y_{k+1})q(Y_{k+1}, X_k)}{\pi(X_k)q(X_k, Y_{k+1})} \wedge 1 - U_{k+1} \right), \quad (4.6)$$

where $(U_i)_{i \in \mathbb{N}^*}$ is a sequence of i.i.d. uniform random variables on $[0, 1]$.

The value of λ characterizes how close the distribution π^λ is to the original target π and therefore how good the proposal is. Small values of λ provide better approximations to π and therefore better proposals (see [DMP18, Proposition 1]), while larger values of λ provide higher levels of smoothing for non-differentiable distributions (see [Per16, Figure 1]). In the case $\lambda = \sigma^2/2$ we obtain the special case of proximal MALA referred to as P-MALA in [Per16].

The main contribution of this paper is to analyze the optimal scaling for proximal MALA defined by (4.6).

Optimal scaling and related works We briefly summarize here some examples of MCMC algorithms and their optimal scaling results; a full review is out of the scope of this paper and we only mention algorithms to which we will compare proximal MALA in the development of this work.

Popular examples of Metropolis MCMC are RWM and MALA. RWM uses as a proposal the transition density $(\mathbf{x}, \mathbf{y}) \mapsto \varphi(\mathbf{y} ; \mathbf{x}, \sigma^2 \mathbf{I}_d)$, where $\sigma^2 > 0$. The MALA scheme uses as proposal $(\mathbf{x}, \mathbf{y}) \mapsto \varphi(\mathbf{y} ; \mathbf{x} + (\sigma^2/2)\nabla \log \pi(\mathbf{x}), \sigma^2 \mathbf{I}_d)$. As we will show in Section 4.3.1, proximal MALA can be considered as an extension of MALA.

A natural question to address when implementing Metropolis adjusted algorithms is how to set the parameter σ^2 (variance parameter for RWM, step size parameter for MALA) to maximize the efficiency of the algorithm. Small values of σ^2 result in higher acceptance probability and cause sticky behaviour, while large values of σ^2 result in a high number of rejections with the chain $(X_k)_{k \geq 0}$ moving slowly [RR+01]. Optimal scaling studies aim to address this question by investigating how σ^2 should behave with respect to the dimension d of the support of π in the high dimensional setting $d \rightarrow \infty$, to obtain the best compromise.

The standard optimal scaling set-up considers the case of d -dimensional targets π_d which are product form, i.e.

$$\pi_d(\mathbf{x}^d) = \prod_{i=1}^d \pi(x_i^d), \quad (4.7)$$

where x_i^d stands for the i -th component of \mathbf{x}^d and π is a one-dimensional probability density with respect to the Lebesgue measure. Under appropriate assumptions on the regularity of π , and assuming that the MCMC algorithm is initialized at stationarity, the optimal value of σ^2 scales as $\ell^2/d^{2\alpha}$ with $\ell > 0$, $2\alpha = 1$ for RWM [RGG97b] and $2\alpha = 1/3$ for MALA [RR98b].

By setting α to these values, it is then possible to show that each as $d \rightarrow \infty$ each 1-dimensional component of the Markov chain defined by RWM and MALA, appropriately rescaled in time, converges to the Langevin diffusion

$$dL_t = h(\ell)^{1/2} dB_t - \frac{h(\ell)}{2} [\log \pi]'(x) dt,$$

where $(B_t)_{t \geq 0}$ is a standard Brownian motion and $h(\ell)$, referred to as speed function of the diffusion, is a function of the parameter $\ell > 0$ that we may tune. Indeed, it is well-known that $(L_{h(\ell)t})_{t \geq 0}$ is a solution of the Langevin diffusion (4.1). As a result, we may identify the values of ℓ maximizing $h(\ell)$ for the algorithms at hand to approximate the fastest version of the Langevin diffusion. The optimal values for ℓ results in an optimal average acceptance probability of 0.234 for RWM and 0.574 for MALA.

The scaling properties allow to get an intuition of the efficiency of the corresponding algorithms: RWM requires $\mathcal{O}(d)$ steps to achieve convergence on a d -dimensional target, i.e. its efficiency is $\mathcal{O}(d^{-1})$, while MALA has efficiency $\mathcal{O}(d^{-1/3})$. While these results are asymptotic in d , the insights obtained by considering the limit case $d \rightarrow \infty$ prove to be useful in practice [RR+01].

In the context of non-smooth and even discontinuous target distributions, studying the simpler RWM algorithm applied to a class of distributions on compact intervals, [NR08; NRY+12] show that the lack of smoothness affects the optimal scaling of RWM with respect to dimension d . More precisely, they show that for a class of discontinuous densities which includes the uniform distribution on $[0, 1]$, the optimal scaling of RWM is of order $\mathcal{O}(d^{-2})$. On the other hand, in the case where the set of non-differentiability X of π is a null set with respect to the Lebesgue measure, [Dur+17] shows that under appropriate conditions, including L^p differentiability, the optimal scaling of RWM is of order $\mathcal{O}(d^{-1})$ still.

The scaling properties of proximal MALA have been partially investigated in [Pil22], which shows that P-MALA, obtained when $\lambda = \sigma^2/2$, has the same scaling properties of MALA for the finite dimensional Gaussian density and for a class of infinite dimensional target measures (Theorem 2.1 and Theorem 5.1 therein, respectively).

4.3 Optimal scaling of Proximal MALA

We consider the same set up of [RR98b] and briefly recalled above. Given a real-valued function $g : \mathbb{R} \rightarrow \mathbb{R}$ satisfying **A18** we consider the i.i.d. d -dimensional target specified by (4.7) with

$$\pi(x) \propto \exp(-g(x)). \quad (4.8)$$

Since for any \mathbf{x}^d , $G(\mathbf{x}^d) = \sum_{i=1}^d g(x_i^d)$, we have by [PB14, Section 2.1]

$$\text{prox}_G^\lambda(\mathbf{x}^d) = (\text{prox}_g^\lambda(x_1^d), \dots, \text{prox}_g^\lambda(x_d^d))^\top.$$

It follows that the distribution of the proposal with target π_d in (4.7)-(4.8) is also product form $q_d(\mathbf{x}^d, \mathbf{y}^d) = \prod_{i=1}^d q(x_i^d, y_i^d)$ with

$$q(x_i^d, y_i^d) = \frac{1}{(2\pi\sigma^2)^{1/2}} \exp\left(-\frac{(y_i^d - (1 - \sigma^2/(2\lambda))x_i^d - \sigma^2 \text{prox}_g^\lambda(x_i^d)/(2\lambda))^2}{2\sigma^2}\right),$$

and $\lambda > 0$. For any dimension $d \in \mathbb{N}^*$, we denote by $(X_k^d)_{k \in \mathbb{N}}$ the Markov chain defined by the Metropolis recursion (4.6) with target distribution π_d and proposal density q_d and associated to the sequence of candidate moves

$$Y_{k+1}^d = \left(1 - \frac{\sigma^2}{2\lambda}\right) X_k^d + \frac{\sigma^2}{2\lambda} \text{prox}_G^\lambda(X_k^d) + \sigma Z_{k+1}^d. \quad (4.9)$$

As mentioned in the introduction, the focus of this work is on investigating the optimal dependence of the proposal variance σ^2 on the dimension d of the target π . In this section, we make the dependence of the proposal variance on the dimension explicit and let $\sigma_d^2 = \ell^2/d^{2\alpha}$ and $\lambda_d = c^2/2d^{2\beta}$ for some $\alpha, \beta > 0$ and some constants c, ℓ independent on d . Thus, we can write λ_d as a function of σ_d , $\lambda_d = \sigma_d^{2m} r/2$, where we defined $r = c^2/\ell^{2m} > 0$ and $m = \beta/\alpha$. By writing λ_d as a function of σ_d we can decouple the effect of the constants

c, ℓ from that of the dependence on d (i.e. α, β). The value of m controls the relative speed at which σ_d and λ_d converge to 0 as $d \rightarrow \infty$, when $m = 1$, σ_d and λ_d decay to 0 at the same rate, for $m > 1$ the decay of λ_d is faster than that of σ_d and for $m < 1$ the decay of λ_d is slower than that of σ_d . The parameter r allows to refine the comparison between σ_d and λ_d as $\beta = \alpha$. In the case $m = 1, r = 1$ we get the P-MALA algorithm studied in [Per16; Pil22], while for all other values of r, m we have a family of proposals whose behaviour depends on r and m .

4.3.1 Differentiable targets

We start with the case where π is continuously differentiable. Since MALA can be applied to this class of targets, the results obtained in this section allow direct comparison of proximal MALA algorithms with MALA and thus between gradient-based algorithms (MALA) and algorithms that use proximal operator-based approximations of the gradient (proximal MALA). If $G = -\log \pi$ is continuously differentiable, using [BC11, Corollary 17.6], $\text{prox}_G^\lambda(\mathbf{x}) = -\lambda \nabla G(\text{prox}_G^\lambda(\mathbf{x})) + \mathbf{x}$, and (4.5) reduces to

$$\xi_{k+1} = \xi_k - \frac{\sigma^2}{2} \nabla G(\text{prox}_G^\lambda(\xi_k)) + \sigma Z_{k+1}. \quad (4.10)$$

Hence, the value of λ controls how close to ξ_k is the point at which the gradient is evaluated. For $\lambda \rightarrow 0$, the proximal MALA proposal becomes arbitrarily close to that of MALA, while, as λ increases (4.10) moves away from MALA.

Our main result, Theorem 4.1 below, shows that the relative speed of decay (i.e. m) influences the optimal scaling of the resulting proximal MALA algorithm, while the constant r influences the speed function of the limiting diffusion.

We make the following assumptions on the regularity of g .

A 19. *g is a C^8 -function whose derivatives are bounded by some polynomial: there exists $k_0 \in \mathbb{N}$ such that*

$$\sup_{x \in \mathbb{R}} \max_{i \in \{0, \dots, 8\}} [g^{(i)}(x)/(1 + |x|^{k_0})] < \infty.$$

Note that under **A 18** and **A 19**, [DMP18, Lemma A.1] implies that $\int_{\mathbb{R}} x^k \exp(-g(x)) dx < \infty$ for any $k \in \mathbb{N}$. We also assume that the sequence of proximal MALA algorithms is initialized at stationarity.

A 20. *For any $d \in \mathbb{N}^*$, X_0^d has distribution π_d .*

The assumptions above closely resemble those of [RR98b] used to obtain the optimal scaling results for MALA. In particular, **A 19** ensures that we can approximate the log-acceptance ratio in (4.6) with a Taylor expansion, while **A 20** avoids technical complications due to the transient phase of the algorithm. We discuss how the latter assumption could be relaxed in Section 4.5.

For technical reasons, and to allow direct comparisons with the results established in [RR98b] for MALA, we will also consider the following regularity assumption

A 21. *The function g' is Lipschitz continuous.*

We denote by L_t^d the linear interpolation of the first component of the discrete time Markov chain $(X_k^d)_{k \geq 0}$ obtained with the generic proximal MALA algorithm described above

$$L_t^d = (\lceil d^{2\alpha} t \rceil - d^{2\alpha} t) X_{\lfloor d^{2\alpha} t \rfloor, 1}^d + (d^{2\alpha} t - \lfloor d^{2\alpha} t \rfloor) X_{\lceil d^{2\alpha} t \rceil, 1}^d, \quad (4.11)$$

where $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ denote the lower and upper integer part functions, respectively, and denote by $X_{k,1}^d$ the first component of X_k^d . The following result shows that in the limit $d \rightarrow \infty$ the properties of proximal MALA depend on the relative speed at which $\sigma_d^2 = \ell^2/d^{2\alpha}$ and $\lambda_d = c^2/2d^{2\beta}$ converge to 0. Recall that we set $r = c^2/\ell^{2m} > 0$ and under **A 20**, consider for any $d \in \mathbb{N}^*$,

$$a_d(\ell, r) = \mathbb{E} \left[\frac{\pi_d(Y_1^d) q_d(Y_1^d, X_0^d)}{\pi_d(X_0^d) q_d(X_0^d, Y_1^d)} \wedge 1 \right]. \quad (4.12)$$

Theorem 4.1. *Assume **A 18**, **A 19** and **A 20**. For any $d \in \mathbb{N}^*$, let $\sigma_d^2 = \ell^2/d^{2\alpha}$ and $\lambda_d = c^2/2d^{2\beta}$ with $\alpha, \beta > 0$. Then, the following statements hold.*

- (a) *If $\alpha = 1/4$, $\beta = 1/8$ and $r > 0$, we have $\lim_{d \rightarrow +\infty} a_d(\ell, r) = 2\Phi(-\ell^2 K_1(r)/2)$, where Φ is the distribution function of a standard normal and*

$$K_1^2(r) = \frac{r^2}{4} \mathbb{E} \left[\left\{ g''(X_{0,1}^d) g'(X_{0,1}^d) \right\}^2 \right].$$

*If in addition, **A 21** holds.*

- (b) *If $\alpha = 1/6$, $\beta = 1/6$ and $r > 0$, we have $\lim_{d \rightarrow +\infty} a_d(\ell, r) = 2\Phi(-\ell^3 K_2(r)/2)$, where Φ is the distribution function of a standard normal and*

$$\begin{aligned} K_2^2(r) &= \left(\frac{r}{8} + \frac{r^2}{4} \right) \mathbb{E} \left[\left\{ g''(X_{0,1}^d) g'(X_{0,1}^d) \right\}^2 \right] + \left(\frac{1}{16} + \frac{r}{8} \right) \mathbb{E} \left[g''(X_{0,1}^d)^3 \right] \\ &\quad + \frac{5}{48} \mathbb{E} \left[g'''(X_{0,1}^d)^2 \right]. \end{aligned}$$

- (c) *If $\alpha = 1/6$, $\beta > 1/6$ and $r > 0$, we have $\lim_{d \rightarrow +\infty} a_d(\ell, r) = 2\Phi(-\ell^3 K_2(0)/2)$, where Φ is the distribution function of a standard normal.*

In addition, in all these cases, as $d \rightarrow \infty$ the process $(L_t^d)_{t \geq 0}$ converges weakly to the Langevin diffusion

$$dL_t = h(\ell, r)^{1/2} dB_t - \frac{h(\ell, r)}{2} g'(x) dt, \quad (4.13)$$

where $(B_t)_{t \geq 0}$ denotes standard Brownian motion and $h(\ell, r) = \ell^2 a(\ell, r)$ is the speed of the diffusion, setting $a(\ell, r) = \lim_{d \rightarrow \infty} a_d(\ell, r)$. If $\alpha = 1/4$, $\beta = 1/8$, for any $r > 0$, $\ell \mapsto h(\ell, r)$ is maximized at the unique value of ℓ such that $a(\ell, r) = 0.452$; while if $\alpha = 1/6$, $\beta = m/6$ with $m \geq 1$ and $r > 0$, $\ell \mapsto h(\ell, r)$ is maximized at the unique value of ℓ such that $a(\ell, r) = 0.574$.

4.3. Optimal scaling of Proximal MALA

Proof. The proof follows that of [RR98b, Theorem 1, Theorem 2] and is postponed to Appendix 4.7. \square

The theorem above shows that the relative speed at which λ_d converges to 0 influences the scaling of the resulting proximal algorithm. In case (c), $m > 1$ and λ_d decays with d at a faster rate than σ_d^2 . This causes the proximity map (4.4) to collapse onto the identity and therefore the proposal (4.10) is arbitrarily close to that of MALA. The resulting scaling limit also coincides with that of MALA established in [RR98b, Theorem 1, Theorem 2].

If λ_d and σ_d^2 decay at the same rate (case (b)), the amount of gradient information provided by the proximity map is controlled by r . Comparing our result for case (b) with [RR98b, Theorem 1] we find that

$$K_2^2(0) = \frac{1}{16} \mathbb{E} [g''(X_{0,1}^d)^3] + \frac{5}{48} \mathbb{E} [g'''(X_{0,1}^d)^2] = K_{\text{MALA}}^2;$$

thus, we have

$$\begin{aligned} K_2^2(r) &= K_2^2(0) + \left(\frac{r}{8} + \frac{r^2}{4}\right) \mathbb{E} [\{g''(X_{0,1}^d)g'(X_{0,1}^d)\}^2] + \frac{r}{8} \mathbb{E} [g''(X_{0,1}^d)^3] \\ &= K_{\text{MALA}}^2 + \left(\frac{r}{8} + \frac{r^2}{4}\right) \mathbb{E} [\{g''(X_{0,1}^d)g'(X_{0,1}^d)\}^2] + \frac{r}{8} \mathbb{E} [g''(X_{0,1}^d)^3] \geq K_{\text{MALA}}^2, \end{aligned}$$

since the convexity of g implies that $g'' \geq 0$. In particular, $K_2^2(r)$ is an increasing function of r achieving its minimum when $r \rightarrow 0$ (i.e. MALA), see Figure 4.1(a).

In case (a), $m = 1/2$ and λ_d decays more slowly than σ_d^2 . As a consequence, the gradient information provided by the proximity map is smaller than in cases (b)–(c), and the resulting scaling differs from that of MALA. The value of $K_1^2(r)$ is increasing in r and the speed of the corresponding diffusion also depends on r (see Figure 4.1(a) gray lines and Figure 4.1(b)).

Example 4.2 (Gaussian target). Take $g(x) = x^2/2$, $\text{prox}_\lambda^g(x) = x/(1 + \lambda)$. In this case, g' is Lipschitz continuous and we have $K_1^2(r) = r^2/4$, $K_2^2(r) = (1 + 4r + 4r^2)/16$ and $K_2^2(0) = K_{\text{MALA}}^2 = 1/16$. The corresponding speeds are given in Figure 4.1(a). Optimizing for $m = 1, r = 0$ (MALA) and $m = 1, r = 1$ (P-MALA) we obtain

$$h^{\text{MALA}}(\ell, r) = 1.5639, \quad h^{\text{P-MALA}}(\ell, r) = 0.7519,$$

achieved with $\ell^{\text{MALA}} = 1.6503$ and $\ell^{\text{P-MALA}} = 1.1443$, respectively. For Gaussian targets, MALA is geometrically ergodic [DM22], and therefore the optimal choice in terms of speed of convergence is MALA which is obtained for $r = 0$. The result for $r = 1$ and $m = 1$ are also given in [Pil22, Theorem 2.1].

Example 4.3 (Target with light tails). Take $g(x) = x^4$, which gives a normalized distribution with normalizing constant $2\Gamma(5/4)$. The proximity map is

$$\text{prox}_g^\lambda(x) = \frac{1}{2} \left[\frac{\sqrt[3]{9\lambda^2x + \sqrt{54\lambda^4x^2 + 3\lambda^3}}}{3^{2/3}\lambda} - \frac{1}{\sqrt[3]{27\lambda^2x + 3\sqrt{54\lambda^4x^2 + 3\lambda^3}}} \right].$$

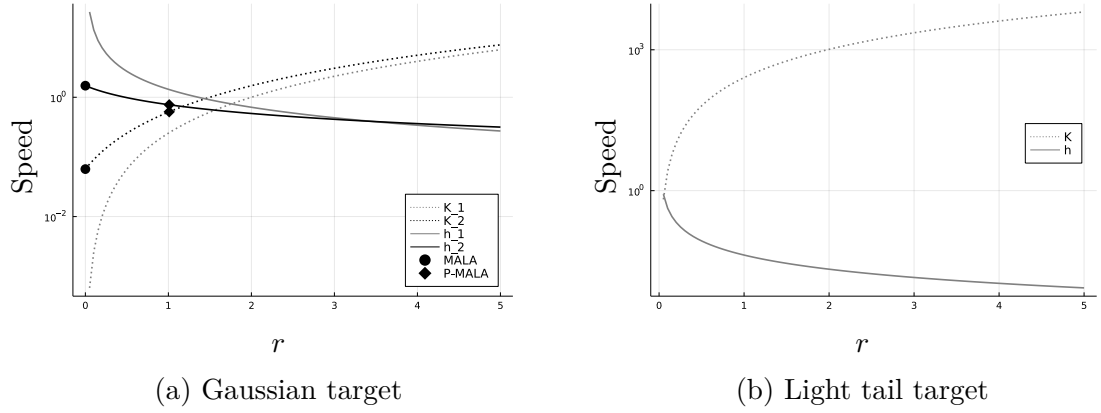


Figure 4.1: Value of K for $i = 1, 2$ and speed of the corresponding Langevin diffusion as a function of r for a Gaussian target and a light tail target. We denote by h_1 the speed obtained in case (a), by h_2 that obtained in (b). In case (c) both K_3 and the speed h_3 are constant w.r.t. r and coincide with that of MALA. For the Gaussian target we report the results for case (a)–(c) while for the light tail target we only report (a).

In this case g' is not Lipschitz continuous and therefore we only consider (a), for which we have $K_1^2(r) = 144r^2\Gamma(11/4)/\Gamma(5/4)$. The corresponding speed is given in Figure 4.1(b).

4.3.2 Laplace target

As discussed in the introduction, proximal MALA has been widely used to quantify uncertainty in imaging applications, in which target distributions involving the ℓ^1 norm are particularly common [Agr+21; DMP18; Per16; ZCZ22].

Here, we consider π_d^L to be the product of d i.i.d. Laplace distributions as in (4.7),

$$\pi_d^L(\mathbf{x}^d) = \prod_{i=1}^d \pi^L(x_i^d), \text{ for } \mathbf{x}^d \in \mathbb{R}^d, \text{ where } \pi^L(x) = 2^{-1} \exp(-|x|) . \quad (4.14)$$

For this particular choice of one-dimensional target distribution, the corresponding potential G is $x \mapsto |x|$ and satisfies A18. Then, the proximity map is given by the soft thresholding operator [PB14, Section 6.1.3]

$$\text{prox}_G^\lambda(x) = (x - \text{sgn}(x)\lambda) \mathbb{1}\{|x| \geq \lambda\} , \quad (4.15)$$

where $\text{sgn} : \mathbb{R} \rightarrow \{-1, 1\}$ is the sign function, given by $\text{sgn}(x) = -1$ if $x < 0$, $\text{sgn}(0) = 0$, and $\text{sgn}(x) = 1$ otherwise. This operator is a continuous but not continuously differentiable map whose non-differentiability points are the extrema of the interval $[-\lambda, \lambda]$ and are controlled by the value of the proximity parameter λ .

4.3. Optimal scaling of Proximal MALA

Plugging (4.15) in (4.9), the proximal MALA algorithm applied to π_d^L proposes component-wise for $i = 1, \dots, d$

$$Y_{k+1,i}^d = X_{k,i}^d - \frac{\sigma_d^2}{2} \operatorname{sgn}(X_{k,i}^d) \mathbb{1}\{|X_{k,i}^d| \geq \lambda_d\} - \frac{\sigma_d^2}{2\lambda_d} X_{k,i}^d \mathbb{1}\{|X_{k,i}^d| < \lambda_d\} + \sigma_d Z_{k+1,i}^d. \quad (4.16)$$

For $X_{k,i}^d$ close to 0 (i.e. the point of non-differentiability) the proximal MALA proposal is a biased random walk around $X_{k,i}^d$, while outside the region $[-\lambda_d, \lambda_d]$ the proposal coincides with that of MALA. As $\lambda_d \rightarrow 0$ the region in which the proximal MALA proposal coincide with that of MALA increases and when $\lambda_d \approx 0$ the region $[-\lambda_d, \lambda_d]$ in which the proposal corresponds to a biased random walk is negligible, as confirmed by the asymptotic acceptance rate in Theorem 4.4.

We also consider the case $\lambda_d = 0$ for any d . Then, the proposal (4.16) becomes the proposal for the subgradient version of MALA: $Y_{k+1,i}^d = X_{k,i}^d - (\sigma_d^2/2) \operatorname{sgn}(X_{k,i}^d) + \sigma_d Z_{k+1,i}^d$, referred to as sG-MALA.

The proof of the optimal scaling for the Laplace distribution follows the structure of that of [Dur+17] for L^p -mean differentiable distributions. We start by characterizing the asymptotic acceptance ratio of a generic proximal MALA algorithm; contrary to Theorem 4.1 for differentiable targets, in the limit $d \rightarrow \infty$ the properties of proximal MALA do not depend on the relative speed at which $\sigma_d^2 = \ell^2/d^{2\alpha}$ and $\lambda_d = c^2/2d^{2\beta}$ converge to 0, as long as λ_d decays at least at the same rate as σ_d^2 . In this regime, the region in which the proposal (4.16) corresponds to a biased random walk proposal is negligible, and therefore we obtain the same scaling obtained with $\lambda_d = 0$ and corresponding to sG-MALA.

Theorem 4.4. *Assume A20 and consider the sequence of target distributions $\{\pi_d^L\}_{d \in \mathbb{N}^*}$ given in (4.14). For any $d \in \mathbb{N}^*$, let $\sigma_d^2 = \ell^2/d^{2\alpha}$ and $\lambda_d = c^2/2d^{2\beta}$ with $\alpha = 1/3$ and $\beta = m/3$ for $m \geq 1$. Then, we have $\lim_{d \rightarrow \infty} a_d(\ell, r) = a^L(\ell) = 2\Phi(-\ell^{3/2}/(72\pi)^{1/4})$, where $(a_d(\ell, r))_{d \in \mathbb{N}^*}$ is defined in (4.12), with $r = c^2/\ell^{2m}$, and Φ is the distribution function of a standard normal.*

Proof. The proof is postponed to Section 4.6.1. □

Theorem 4.4 shows that the asymptotic average acceptance rate $a^L(\ell)$ does not depend on r and as a result on c .

Having identified the possible scaling for proximal MALA with Laplace target, we are now ready to show weak convergence to the appropriate Langevin diffusion. To this end, we adapt the proof strategy followed in [JLM14] and [Dur+17].

As for the differentiable case, consider the linear interpolation $(L_t^d)_{t \geq 0}$ of the first component of the Markov chain $(X_k^d)_{k \geq 0}$ given in (4.11). For any $d \in \mathbb{N}^*$, denote by ν_d the law of the process $(L_t^d)_{t \geq 0}$ on the space of continuous functions from \mathbb{R}_+ to \mathbb{R} , $C(\mathbb{R}^+, \mathbb{R})$, endowed with the topology of uniform convergence over compact sets and its corresponding σ -field. We first show that the sequence $(\nu_d)_{d \in \mathbb{N}^*}$, admits a weak limit point as $d \rightarrow \infty$.

Proposition 4.5. *Assume A20 and consider the sequence of target distributions $\{\pi_d^L\}_{d \in \mathbb{N}^*}$ given in (4.14). For any $d \in \mathbb{N}^*$, let $\sigma_d^2 = \ell^2/d^{2\alpha}$ and $\lambda_d = c^2/2d^{2\beta}$ with $\alpha = 1/3$ and*

$\beta = m/3$. The sequence $(\nu_d)_{d \in \mathbb{N}^*}$ is tight in $\mathbf{M}^1(C(\mathbb{R}^+, \mathbb{R}))$, the set of probability measures acting on $C(\mathbb{R}^+, \mathbb{R})$.

Proof. See Section 4.6.2. \square

By Prokhorov's theorem, the tightness of $(\nu_d)_{d \in \mathbb{N}^*}$ implies existence of a weak limit point ν . In our next result, we give a sufficient condition to show that any limit point of $(\nu_d)_{d \in \mathbb{N}^*}$ coincides with the law of a solution of:

$$dL_t = [h^L(\ell)]^{1/2} dB_t - \frac{h^L(\ell)}{2} \operatorname{sgn}(L_t) dt. \quad (4.17)$$

To this end, we consider the martingale problem (see [SV79]) associated with (4.17), that we now present. Let us denote by $C_c^\infty(\mathbb{R}, \mathbb{R})$ the subset of functions of $C(\mathbb{R}, \mathbb{R})$ which are infinitely many times differentiable and with compact support, and define the generator of (4.17) for $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$ by

$$LV(x) = \frac{h^L(\ell)}{2} [V''(x) - \operatorname{sgn}(x)V'(x)]. \quad (4.18)$$

Denote by $(W_t)_{t \geq 0}$ the canonical process on $C(\mathbb{R}_+, \mathbb{R})$, $W_t : \{w_s\}_{s \geq 0} \mapsto w_t$ and the corresponding filtration by $(\mathfrak{F}_t)_{t \geq 0}$. A probability measure ν is said to solve the martingale problem associated with (4.17) with initial distribution π^L , if the pushforward of ν by W_0 is π^L and if for all $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$, the process

$$\left(V(W_t) - V(W_0) - \int_0^t LV(W_u) du \right)_{t \geq 0}$$

is a martingale with respect to ν and the filtration $(\mathfrak{F}_t)_{t \geq 0}$. The following proposition gives a sufficient condition to prove that ν is a solution of the martingale problem:

Proposition 4.6. *Suppose that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$, $m \in \mathbb{N}$, $\rho : \mathbb{R}^m \rightarrow \mathbb{R}$ bounded and continuous, and for any $0 \leq t_1 \leq \dots \leq t_m \leq s \leq t$:*

$$\lim_{d \rightarrow +\infty} \mathbb{E}^{\nu_d} \left[\left(V(W_t) - V(W_s) - \int_s^t LV(W_u) du \right) \rho(W_{t_1}, \dots, W_{t_m}) \right] = 0.$$

Then any limit point of $(\nu_d)_{d \in \mathbb{N}^}$ on $\mathbf{M}^1(C(\mathbb{R}^+, \mathbb{R}))$ is a solution to the martingale problem associated with (4.17).*

Proof. See Section 4.6.3. \square

Finally, we use this sufficient condition to establish that any limit point of $(\nu_d)_{d \in \mathbb{N}^*}$ is a solution of the martingale problem for (4.17). Uniqueness in law of solutions of (4.17) allows to conclude that $(L_t^d)_{t \geq 0}$ converges weakly to the Langevin diffusion (4.17), which establishes our main result.

Theorem 4.7. *The sequence of processes $\{(L_t^d)_{t \geq 0} : d \in \mathbb{N}^*\}$ converges in distribution towards $(L_t)_{t \geq 0}$, solution of (4.17) as $d \rightarrow \infty$, with $h^L(\ell) = \ell^2 a^L(\ell)$ and a^L defined in Theorem 4.4. In addition, h^L is maximized at the unique value of ℓ such that $a^L(\ell) = 0.360$.*

Proof. See Section 4.6.4. \square

4.4 Practical Implications and Numerical Simulations

The optimal scaling results in Sections 4.3.1 and 4.3.2 provide some guidance on the choice of the parameters σ and λ of proximal MALA algorithms, suggesting that smaller values of λ provide better efficiency in terms of number of steps necessary to convergence (Theorem 4.1).

However, a number of other factors must be taken into account. First, as shown in [JR07; MT96; RT96b; RT96c] the convergence properties of Metropolis adjusted algorithms are influenced by the shape of the target distribution and, in particular, by its tail behavior. Secondly, when comparing proximal MALA algorithms with gradient-based methods (e.g. MALA) one must take into account the cost of obtaining the gradients, whether this comes from automatic differentiation algorithms or from evaluating a potentially complicated gradient function. On the other hand, proximity mappings can be quickly found or approximated solving convex optimization problems which have been widely studied in the convex optimization literature (e.g. [PB14, Chapter 6], [CP11] and [Per16, Section 3.2.3]).

In terms of convergence properties, we are usually interested in the family of distributions for which the discrete time Markov chain produced by our algorithm is geometrically ergodic, together with the optimal scaling results briefly recalled in Section 4.2. Normally, the ergodicity results are given by considering the one-dimensional class of distributions $\mathcal{E}(\beta, \gamma)$ introduced in [RT96c] and defined for $\gamma > 0$ and $0 < \beta < \infty$ by

$$\mathcal{E}(\beta, \gamma) : \left\{ \pi : \mathbb{R} \rightarrow [0, +\infty) : \pi(x) \propto \exp\left(-\gamma|x|^\beta\right), |x| > x_0 \text{ for some } x_0 > 0 \right\}.$$

As observed by [LZ22], there usually is a trade-off between ergodicity and optimal scaling results, algorithms providing better optimal scaling results tend to be geometrically ergodic for a smaller set of targets (e.g. MALA w.r.t. RWM).

As suggested by Theorem 4.1, the scaling properties of proximal MALA on differentiable targets are close to those of MALA. This leads to a natural comparison between the two algorithms. First, we observe that A18 rules out targets for which G is not convex and therefore restricts the families $\mathcal{E}(\beta, \gamma)$ to $\beta \geq 1$. To compare MALA with proximal MALA we therefore focus on distributions with $\beta \geq 1$.

It is shown in [RT96c] that MALA is geometrically ergodic for targets in $\mathcal{E}(\beta, \gamma)$ with $1 \leq \beta \leq 2$ (with some caveat for $\beta = 2$). Theorem 4.1–(b) and (c) show that in this case proximal MALA has the same scaling properties of MALA but in case (b) the asymptotic speed of convergence decays as the constant r increases (Figure 4.1(a)), with the maximum achieved for $r \rightarrow 0$, for which proximal MALA collapses onto MALA. Since MALA is geometrically ergodic, and achieves better (or equivalent) scaling properties than proximal MALA, it would be natural to prefer MALA to proximal MALA for this set of targets. However, if the gradient is costly to obtain, one might instead consider to use proximal MALA with a small λ , to retain scaling properties as close as possible to that of MALA but to reduce the computational cost of evaluating the gradient.

In the case of differentiable targets with light-tails (i.e. $\beta > 2$), MALA is known not to be geometrically ergodic [RT96c, Section 4.2] while the ergodicity properties of

proximal MALA have only been partially studied in [Per16, Section 3.2.2] for the case $\lambda = \sigma^2/2$ (P-MALA). As shown in [Per16, Section 2.1], given a distribution $\pi \in \mathcal{E}(\beta, \gamma)$ with $\beta \geq 1$, the distribution π_λ obtained using the potential (4.3) belongs to $\mathcal{E}(\beta', \gamma')$, where $\beta' = \min(\beta, 2)$ and γ' depending on λ . This suggests that proximal MALA is likely to be geometrically ergodic for appropriate choices of λ ; a first result in this direction is given in [Per16, Corollary 3.2] for the P-MALA case $\lambda = \sigma^2/2$. Theorem 4.1–(c) restricts the sets of available λ s showing that for light-tail distributions (for which A21 does not hold) λ should decay at half the speed of σ^2 . Studying the ergodicity properties of proximal MALA in function of the parameter λ is, of course, an interesting problem that we leave for future work.

For the Laplace distribution, Theorem 4.4 shows that the value of λ does not influence the asymptotic acceptance ratio of proximal MALA, as long as λ decays with d at least as fast as σ^2 . The scaling properties and the asymptotic speed $h(\ell)$ in Theorem 4.7 do not depend on λ and coincide with that of the sG-MALA (obtained for $\lambda = 0$). Hence, in terms of optimal scaling, there does not seem to be a difference between proximal MALA and sG-MALA for the Laplace distribution.

4.4.1 Numerical Experiments

To illustrate the results established in Section 4.3.1 and 4.3.2 we consider here a small collection of simulation studies. The aim of these studies is to empirically confirm the optimal scalings identified in Theorem 4.1 and 4.4, investigate the dimension d at which the asymptotic acceptance ratio $\lim_{d \rightarrow \infty} a_d(\ell, r)$ well approximates the empirical average acceptance ratio and, consequently, for which dimensions d we can expect the optimal asymptotic acceptances in Theorem 4.1 and 4.4 to guarantee maximal speed $h(\ell, r)$ (approximated by the expected squared jumping distance, see, e.g. [GRG96]) for the corresponding diffusion. We summarize here our findings, a more detailed discussion can be found in Section 4.8.

For the differentiable case, we consider the Gaussian distribution in Example 4.2 and four algorithmic settings which correspond to the three cases identified in Theorem 4.1 and MALA. The different values of r and m influence the dimension required to observe convergence to the theoretical limit in Theorem 4.1: for $r \rightarrow 0$ and $m = 1$ (MALA) and $m = 1/2, r = 1$ (corresponding to Theorem 4.1–(a)) the theoretical limit is already achieved for d of order 10^2 , while in the cases $m = 3, r = 2$ and $m = r = 1$ (corresponding to Theorem 4.1–(c) and (b), respectively) our simulation result match the theoretical limit only for d of order 10^5 or higher.

The results for the Laplace case are similar, with the case $m > 1$ requiring a higher d to observe convergence to the theoretical limit. Figure 4.2 and Figure 4.3 provide numerical simulations of the behavior, as d increases, of the mean acceptance ratio $(a_d(\ell, r))_{d \in \mathbb{N}^*}$ as a function of ℓ and $(\text{ESJD}_d)_{d \in \mathbb{N}^*}$ as a function of $(a_d(\ell, r))_{d \in \mathbb{N}^*}$, for sG-MALA ($r = 0$) and P-MALA ($r = 1$) respectively. These confirm our theoretical founding Theorem 4.4 and Theorem 4.7.

In general, we find that the optimal average acceptance ratios in Theorem 4.1 and 4.7 guarantee maximal speed $h(\ell, r)$ for d sufficiently large (for small d the optimal acceptance ratio often differs from the optimal asymptotic one, see, e.g. [SR+09, Section 2.1]).

4.4. Practical Implications and Numerical Simulations




Figure 4.2: Proximal MALA with Laplace target and $m = 1, r = 0$ (sG-MALA). Left: acceptance rate as a function of ℓ for increasing dimension d ; Right: $ESJD_d$ as a function of the acceptance rate $a_d(\ell, r)$.

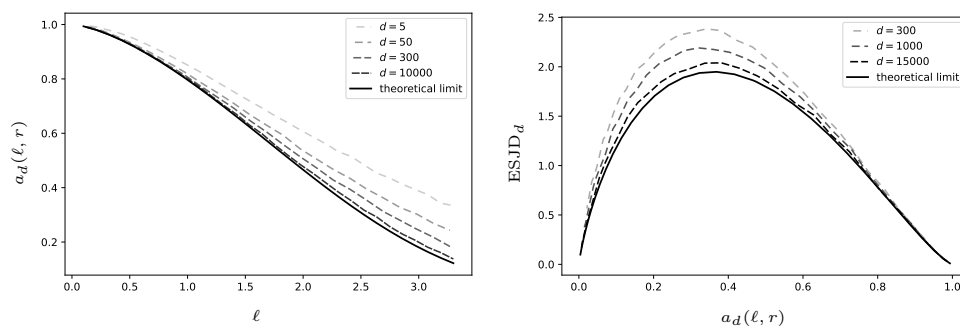


Figure 4.3: Proximal MALA with Laplace target and $m = 1, r = 1$ (P-MALA). Left: acceptance rate as a function of ℓ for increasing dimension d ; Right: $ESJD_d$ as a function of the acceptance rate $a_d(\ell, r)$.




Figure 4.4: Proximal MALA for the target (4.19) and $m = 3, r = 2$, with $\sigma^2 = \ell/d^{2\alpha}$ and $\alpha = 1/3$. Left: acceptance rate as a function of ℓ for increasing dimension d ; Right: $ESJD_d$ as a function of the acceptance rate $a_d(\ell, r)$.

To further investigate the scaling of proximal MALA to other non-differentiable densities, we empirically study the case where the sequence of targets is given by: $x^d \in \mathbb{R}^d$,

$$\pi_d^{\text{GL}}(x^d) = \prod_{i=1}^d \exp(-g(x_i^d)), \quad g(x) = |x| + x^2/2, \quad (4.19)$$

which, like the Laplace distribution in 0, is non-differentiable but convex. The study of such a potential is motivated by Bayesian inverse problems considered in [DMP18; Per16], for which the posterior distribution arises from Gaussian observations and sparsity-induced priors like the Laplace distribution. The posterior then has the form (up to a multiplicative constant) $x^d \mapsto \exp(-\|y^d - \mathbf{A}x^d\| - c_r \sum_{i=1}^d |x_i|)$.

For this choice of target (4.19), the proposal of proximal MALA is given for any $d \in \mathbb{N}^*$, $k \in \mathbb{N}$ and $i \in \{1, \dots, d\}$,

$$Y_{k+1,i}^d = \left(1 - \frac{\sigma_d^2}{2\lambda_d}\right) X_{k,i}^d + \frac{\sigma_d^2}{2\lambda_d} \left((X_{k,i}^d - \text{sgn}(X_{k,i}^d)) \mathbb{1}\{|X_{k,i}^d| \geq \lambda_d\} - \lambda_d X_{k,i}^d \right) + \sigma_d Z_{k+1,i}^d,$$

where $(Z_{k+1,i}^d)_{k \in \mathbb{N}}$ is a sequence of standard normal random variables. We then repeated the same experiments as for the Laplace distribution. The results are shown in Figures 4.4, 4.12 and 4.13. From these figures it is clear that the same scaling holds, i.e. choosing $\sigma_d^2 = \ell/d^{2\alpha}$, $\lambda = \sigma_d^{2m} r/2$ with $\alpha = 1/3$, $r \geq 0$.

4.5 Discussion

In this work we analyze the scaling properties of a wide class of proximal MALA algorithms introduced in [DMP18; Per16] for smooth targets and for the Laplace distribution. We show that the scaling properties of proximal MALA are influenced by the relative speed at which the proximal parameter λ_d and the proposal variance σ_d decay to 0 as $d \rightarrow \infty$ and suggest practical ways to choose λ_d as a function of σ_d to guarantee good results.

In the case of smooth targets, we provide a detailed comparison between proximal MALA and MALA, showing that proximal MALA scales no better than MALA (Theorem 4.1). In particular, Theorem 4.1–(a) shows that if λ_d is too large w.r.t. σ_d then the efficiency of proximal MALA is of order $\mathcal{O}(d^{-1/2})$ and therefore worse than the $\mathcal{O}(d^{-1/3})$ of MALA, suggesting that λ_d should be chosen to decay approximately as σ_d , if possible. If λ_d decays sufficiently fast, then MALA and proximal MALA have similar scaling properties and, in the case in which the proximity map is cheaper to compute than the gradient, one can build proximal MALA algorithms which are as efficient as MALA in terms of scaling but more computationally efficient.

In the case of the Laplace distribution, we show that the scaling of proximal MALA is $\mathcal{O}(d^{-2/3})$ for any λ_d decaying sufficiently fast w.r.t. σ_d and, in the case $\lambda_d = 0$, we obtain a novel optimal scaling result for sG-MALA on Laplace targets.

As discussed in Section 4.4, our analysis provides some guidance on the choice of the parameters that need to be specified to implement proximal MALA, but this analysis should be complemented by an exploration of the ergodicity properties of proximal MALA to obtain a comprehensive description of the algorithms. We conjecture that for sufficiently large values of λ , proximal MALA applied to light tail distributions will be exponentially ergodic; establishing exactly how large should λ be to guarantee fast convergence is an interesting question that we leave for future work. Obtaining these results would open the doors to adaptive tuning strategies for proximal MALA, which are likely to produce better results than those given by the strategies currently used.

The set up under which we carried out our analysis closely resembles that of [RR98b]; we anticipate that A20 could be relaxed following similar ideas as those in [CRR05; JLM14] and that our analysis could be extended to d -dimensional targets π_d possessing some dependence structure following the approach of [BR08; SR+09; YRR20]. Finally, the analysis carried out for the Laplace distribution could be extended to other piecewise smooth distributions provided that the moments necessary for the proof in Section 4.6 can be computed.

4.6 Proof of the Result for the Laplace distribution

In this section we prove the results in Section 4.3.2 which give the scaling properties of proximal MALA (and sG-MALA) for the Laplace distribution. We collect technical results (e.g. moment computations, bounds, etc.) in Appendix 4.10.

We recall that $\sigma_d^2 = \ell^2/d^{2\alpha}$ and $\lambda_d = c^2/2d^{2\beta}$ for some $\alpha, \beta > 0$ and some constants c, ℓ independent on d . Thus, we can write λ_d as a function of σ_d , $\lambda_d = \sigma_d^{2m}r/2$, where we define $r = c^2/\ell^{2m} \geq 0$ and $m = \beta/\alpha$.

In order to study the scaling limit of proximal MALA with Laplace target, consider the mapping $b_d : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$b_d : (x, z) \mapsto z - \frac{\sigma_d}{2} \operatorname{sgn}(x) \mathbb{1} \left\{ |x| \geq \sigma_d^{2m} r / 2 \right\} - \frac{1}{\sigma_d^{2m-1} r} x \mathbb{1} \left\{ |x| < \sigma_d^{2m} r / 2 \right\}, \quad (4.20)$$

which allows us to write the proposal as $Y_{1,i}^d = X_{0,i}^d + \sigma_d b_d(X_{0,i}^d, Z_{1,i}^d)$, for any $i \in \{1, \dots, d\}$.

We consider also the function $\phi_d : \mathbb{R}^2 \rightarrow \mathbb{R}$, given by

$$\begin{aligned}
 \phi_d : (x, z) &\mapsto \log \frac{\pi(x + \sigma_d b_d(x, z)) q(x + \sigma_d b_d(x, z), x)}{\pi(x) q(x, x + \sigma_d b_d(x, z))} \\
 &= |x| - |x + \sigma_d b_d(x, z)| + \frac{z^2}{2} \\
 &\quad - \frac{1}{2\sigma_d^2} \left\{ \frac{\sigma_d^2}{2} \operatorname{sgn}[x + \sigma_d b_d(x, z)] \mathbb{1} \left\{ |x + \sigma_d b_d(x, z)| \geq \frac{\sigma_d^{2m} r}{2} \right\} \right. \\
 &\quad \left. - \sigma_d b_d(x, z) \right. \\
 &\quad \left. + \frac{1}{\sigma_d^{2(m-1)} r} [x + \sigma_d b_d(x, z)] \mathbb{1} \left\{ |x + \sigma_d b_d(x, z)| < \frac{\sigma_d^{2m} r}{2} \right\} \right\}^2.
 \end{aligned} \tag{4.21}$$

4.6.1 Proof of Theorem 4.4

The proof of Theorem 4.4 uses the first three moments of $\phi_{d,1}$, whose computation is postponed to Appendix 4.10.1, and is an application of Lindeberg's central limit theorem. We introduce, for $i \in \{1, \dots, d\}$, $\phi_{d,i} = \phi_d(X_{0,i}^d, Z_{1,i}^d)$ for the sake of conciseness. This allows us to rewrite $a_d(\ell, r)$, defined in (4.12), in the following way,

$$a_d(\ell, r) = \mathbb{E} \left[\exp \left(\sum_{i=1}^d \phi_{d,i} \right) \wedge 1 \right].$$

Remark 4.8. Under A20, the families of random variables $(b_d(X_{0,i}^d, Z_{1,i}^d))_{i \in \{1, \dots, d\}}$ and $(\phi_{d,i})_{i \in \{1, \dots, d\}}$ are i.i.d.

To identify the optimal scaling for the Laplace distribution, we look for those values of α such that $\sum_{i=1}^d \mathbb{E}[\phi_{d,i}]$ and $\operatorname{Var}(\sum_{i=1}^d \phi_{d,i})$ converge to a finite value. Using Remark 4.8, we have that,

$$\sum_{i=1}^d \mathbb{E}[\phi_{d,i}] = d \mathbb{E}[\phi_{d,1}] \quad \text{and} \quad \operatorname{Var} \left(\sum_{i=1}^d \phi_{d,i} \right) = d \operatorname{Var}(\phi_{d,1}). \tag{4.22}$$

Then, using the integrals in Appendix 4.10.1, we find that the only value of α for which (4.22) converge to a finite value with the variance strictly positive is $\alpha = 1/3$ as confirmed empirically in Appendix 4.8.2.

Having identified $\alpha = 1/3$, we can then proceed applying Lindeberg's CLT.

Proof of Theorem 4.4. We start by showing that the acceptance ratio converges to a Gaussian distribution. Define $\mu_d = \mathbb{E}[\phi_{d,1}]$ and $\mathcal{F}_{d,i} = \sigma((X_{0,j}^d, Z_{1,j}^d), 1 \leq j \leq i)$, the natural filtration for $(X_{0,i}^d, Z_{1,i}^d)_{d \in \mathbb{N}, 1 \leq i \leq d}$. The square-integrable martingale sequence

$$\left(\sum_{j=1}^i W_{d,j}, \mathcal{F}_{d,i} \right)_{d \in \mathbb{N}^*, 1 \leq i \leq d}$$

4.6. Proof of the Result for the Laplace distribution

where $W_{d,i} = \phi_{d,i} - \mu_d$, forms a triangular array, to which we can apply the corresponding CLT (e.g. [Shi96, Theorem 4, page 543]). In particular, we have that,

$$\lim_{d \rightarrow \infty} \sum_{i=1}^d \mathbb{E} \left[W_{d,i}^2 \mid \mathcal{F}_{d,i-1} \right] = \lim_{d \rightarrow \infty} d \operatorname{Var}(\phi_{d,1}) = \frac{2\ell^3}{3\sqrt{2\pi}},$$

as shown in Proposition 4.23 in Appendix 4.10.1. It remains to verify Lindeberg's condition: for $\varepsilon > 0$,

$$\lim_{d \rightarrow \infty} d \mathbb{E} \left[W_{d,1}^2 \mathbb{1}_{\{|W_{d,1}| > \varepsilon\}} \right] = 0.$$

In order to verify Lindeberg's condition we verify the stronger Lyapunov condition: there exists $\epsilon > 0$ such that

$$\lim_{d \rightarrow \infty} d \mathbb{E} \left[W_{d,1}^{2+\epsilon} \right] = 0.$$

Pick $\epsilon = 1$ and expand the cube using $\mu_d = \mathbb{E}[\phi_{d,i}]$,

$$\mathbb{E} \left[W_{d,1}^3 \right] = \mathbb{E} \left[\phi_{d,i}^3 \right] - 3\mu_d \mathbb{E} \left[\phi_{d,i}^2 \right] + 2\mu_d^3. \quad (4.23)$$

By Proposition 4.22 in Appendix 4.10.1, we have $\lim_{d \rightarrow \infty} d\mu_d^3 = 0$, $\lim_{d \rightarrow \infty} \mu_d = 0$, and, by Proposition 4.23 in Appendix 4.10.1,

$$\lim_{d \rightarrow \infty} d \mathbb{E} \left[\phi_{d,i}^2 \right] = \frac{2\ell^3}{3\sqrt{2\pi}}.$$

Finally, for the remaining term in (4.23) we use Proposition 4.24 in Appendix 4.10.1 to show that $\lim_{d \rightarrow \infty} d \mathbb{E}[\phi_{d,i}^3] = 0$. The above and the fact that, by Proposition 4.22 in Appendix 4.10.1,

$$\lim_{d \rightarrow \infty} d\mu_d = -\frac{\ell^3}{3\sqrt{2\pi}},$$

show, by Lindeberg's CLT, that the acceptance ratio converges in law to a normal random variable \tilde{Z} with mean $-\ell^3/(3\sqrt{2\pi})$ and variance $2\ell^3/(3\sqrt{2\pi})$.

To conclude the proof, we apply the continuous mapping theorem to the bounded and continuous function $x \mapsto e^x \wedge 1$ and obtain

$$\lim_{d \rightarrow \infty} \exp \left(\sum_{i=1}^d \phi_{d,i} \right) \wedge 1 \stackrel{d}{=} e^{\tilde{Z}} \wedge 1 \quad \text{and} \quad \lim_{d \rightarrow \infty} a_d(\ell, r) = \mathbb{E} \left[e^{\tilde{Z}} \wedge 1 \right],$$

where the limit does not depend on r . Defining $a^L(\ell) = \lim_{d \rightarrow \infty} a_d(\ell, r)$ and using [RGG97b, Proposition 2.4], we have the result. □

4.6.2 Proof of Proposition 4.5

We are interested in the law ν_d of the linear interpolant $(L_t^d)_{t \geq 0}$, defined in (4.11), of the first component of the chain $(X_k^d)_{k \in \mathbb{N}}$. Let us recall the definition of the chain: assumption **A20** gives the initial distribution π_d , then, for any $k \in \mathbb{N}$, the proposal $Y_{k+1}^d = (Y_{k+1,i}^d)_{1 \leq i \leq d}$ is defined in (4.16) with $\sigma_d^2 = \ell^2/d^{2\alpha}$, $\lambda_d = \sigma_d^{2m}r/2$ with $\alpha = 1/3$ and $m \geq 1$. The proposal (4.16) can be written as

$$Y_{k+1,i}^d = X_{k,i}^d + \sigma_d b_d(X_{k,i}^d, Z_{k+1,i}^d), \quad (4.24)$$

for any $i \in \{1, \dots, d\}$, where b_d is defined in (4.20) and $r = c^2/\ell^{2m}$. We further define the acceptance event $A_{k+1}^d = \{b_{k+1}^d = 1\}$ where b_{k+1}^d is as in (4.6).

We can now expand the expression of the linear interpolant L_t^d using (4.6), (4.11) and the definition of A_{k+1}^d ,

$$L_t^d = \begin{cases} X_{\lfloor d^{2\alpha}t \rfloor, 1}^d + (d^{2\alpha}t - \lfloor d^{2\alpha}t \rfloor) \left[\sigma_d Z_{\lfloor d^{2\alpha}t \rfloor, 1}^d - \frac{\sigma_d^2}{2} \operatorname{sgn}(X_{\lfloor d^{2\alpha}t \rfloor, 1}^d) \right] \mathbb{1}_{A_{\lfloor d^{2\alpha}t \rfloor}^d} \\ \quad \text{if } |X_{\lfloor d^{2\alpha}t \rfloor, 1}^d| \geq \frac{\sigma_d^{2m}r}{2} \\ X_{\lfloor d^{2\alpha}t \rfloor, 1}^d + (d^{2\alpha}t - \lfloor d^{2\alpha}t \rfloor) \left[\sigma_d Z_{\lfloor d^{2\alpha}t \rfloor, 1}^d - \frac{1}{\sigma_d^{2(m-1)}r} X_{\lfloor d^{2\alpha}t \rfloor, 1}^d \right] \mathbb{1}_{A_{\lfloor d^{2\alpha}t \rfloor}^d} \\ \quad \text{otherwise} \end{cases}, \quad (4.25)$$

or, equivalently,

$$L_t^d = \begin{cases} X_{\lceil d^{2\alpha}t \rceil, 1}^d - (\lceil d^{2\alpha}t \rceil - d^{2\alpha}t) \left[\sigma_d Z_{\lceil d^{2\alpha}t \rceil, 1}^d - \frac{\sigma_d^2}{2} \operatorname{sgn}(X_{\lceil d^{2\alpha}t \rceil, 1}^d) \right] \mathbb{1}_{A_{\lceil d^{2\alpha}t \rceil}^d} \\ \quad \text{if } |X_{\lceil d^{2\alpha}t \rceil, 1}^d| \geq \frac{\sigma_d^{2m}r}{2} \\ X_{\lceil d^{2\alpha}t \rceil, 1}^d - (\lceil d^{2\alpha}t \rceil - d^{2\alpha}t) \left[\sigma_d Z_{\lceil d^{2\alpha}t \rceil, 1}^d - \frac{1}{\sigma_d^{2(m-1)}r} X_{\lceil d^{2\alpha}t \rceil, 1}^d \right] \mathbb{1}_{A_{\lceil d^{2\alpha}t \rceil}^d} \\ \quad \text{otherwise} \end{cases}.$$

In order to prove Proposition 4.5, we consider Kolmogorov's criterion for tightness (see [Kal21, Theorem 23.7]): the sequence $(\nu_d)_{d \geq 1}$ is tight if the sequence $(L_0^d)_{d \in \mathbb{N}^*}$ is tight, and

$$\mathbb{E} \left[(L_t^d - L_s^d)^4 \right] \leq \gamma(t)(t-s)^2,$$

for some non-decreasing positive function γ , all $0 \leq s \leq t$ and all $d \in \mathbb{N}^*$. The condition on $(L_0^d)_{d \in \mathbb{N}^*}$ is straightforward to check, since by **A20** the distribution of $L_0^d = X_{0,1}^d$ is π^L for all $d \in \mathbb{N}^*$.

Proof of Proposition 4.5. Consider $\mathbb{E} \left[(L_t^d - L_s^d)^4 \right]$, if $\lfloor d^{2\alpha}s \rfloor = \lfloor d^{2\alpha}t \rfloor$, the inequality follows straightforwardly recalling that the moments of normal distributions are bounded: in the case $|X_{\lfloor d^{2\alpha}t \rfloor, 1}^d| = |X_{\lfloor d^{2\alpha}s \rfloor, 1}^d| \geq \sigma_d^{2m}r/2$ it follows directly from the boundedness of the sgn function, while in the case $|X_{\lfloor d^{2\alpha}t \rfloor, 1}^d| = |X_{\lfloor d^{2\alpha}s \rfloor, 1}^d| < \sigma_d^{2m}r/2$ we exploit the boundedness of $X_{\lfloor d^{2\alpha}t \rfloor, 1}^d$ itself.

For all $0 \leq s \leq t$ such that $\lceil d^{2\alpha}s \rceil \leq \lfloor d^{2\alpha}t \rfloor$, we can distinguish three cases.

4.6. Proof of the Result for the Laplace distribution

Case 1 If $|X_{[d^{2\alpha}t],1}^d| \geq \sigma_d^{2m}r/2$ and $|X_{[d^{2\alpha}s],1}^d| \geq \sigma_d^{2m}r/2$, then

$$\begin{aligned} L_t^d - L_s^d &= X_{[d^{2\alpha}t],1}^d - X_{[d^{2\alpha}s],1}^d \\ &\quad + (d^{2\alpha}t - [d^{2\alpha}t]) \left[\sigma_d Z_{[d^{2\alpha}t],1}^d - \frac{\sigma_d^2}{2} \operatorname{sgn}(X_{[d^{2\alpha}t],1}^d) \right] \mathbb{1}_{\mathcal{A}_{[d^{2\alpha}t]}^d} \\ &\quad + ([d^{2\alpha}s] - d^{2\alpha}s) \left[\sigma_d Z_{[d^{2\alpha}s],1}^d - \frac{\sigma_d^2}{2} \operatorname{sgn}(X_{[d^{2\alpha}s],1}^d) \right] \mathbb{1}_{\mathcal{A}_{[d^{2\alpha}s]}^d}. \end{aligned}$$

Using Hölder's inequality and the fact that $0 \leq d^{2\alpha}t - [d^{2\alpha}t] \leq 1$ (and similarly for s) we have

$$\begin{aligned} \mathbb{E} \left[(L_t^d - L_s^d)^4 \right] &\leq C \mathbb{E} \left[\left(X_{[d^{2\alpha}t],1}^d - X_{[d^{2\alpha}s],1}^d \right)^4 \right] \\ &\quad + C \frac{(d^{2\alpha}t - [d^{2\alpha}t])^2}{d^{4\alpha}} \mathbb{E} \left[\left(\ell Z_{[d^{2\alpha}t],1}^d \right)^4 + \frac{\ell^8}{2^4 d^{4\alpha}} \right] \\ &\quad + C \frac{([d^{2\alpha}s] - d^{2\alpha}s)^2}{d^{4\alpha}} \mathbb{E} \left[\left(\ell Z_{[d^{2\alpha}s],1}^d \right)^4 + \frac{\ell^8}{2^4 d^{4\alpha}} \right]. \end{aligned}$$

Recalling that the moments of Z^d are bounded and that $d^{2\alpha}s \leq [d^{2\alpha}s] \leq [d^{2\alpha}t] \leq d^{2\alpha}t$, it follows

$$\mathbb{E} \left[(L_t^d - L_s^d)^4 \right] \leq C \left((t-s)^2 + \mathbb{E} \left[\left(X_{[d^{2\alpha}t],1}^d - X_{[d^{2\alpha}s],1}^d \right)^4 \right] \right). \quad (4.26)$$

Case 2 If $|X_{[d^{2\alpha}t],1}^d| \geq \sigma_d^{2m}r/2$ and $|X_{[d^{2\alpha}s],1}^d| < \sigma_d^{2m}r/2$ or $|X_{[d^{2\alpha}t],1}^d| < \sigma_d^{2m}r/2$ and $|X_{[d^{2\alpha}s],1}^d| \geq \sigma_d^{2m}r/2$. We only describe the argument for the first case, the second case follows from analogous steps. Take

$$\begin{aligned} L_t^d - L_s^d &= X_{[d^{2\alpha}t],1}^d + (d^{2\alpha}t - [d^{2\alpha}t]) \left[\sigma_d Z_{[d^{2\alpha}t],1}^d - \frac{\sigma_d^2}{2} \operatorname{sgn}(X_{[d^{2\alpha}t],1}^d) \right] \mathbb{1}_{\mathcal{A}_{[d^{2\alpha}t]}^d} \\ &\quad - X_{[d^{2\alpha}s],1}^d - ([d^{2\alpha}s] - d^{2\alpha}s) \left(\sigma_d Z_{[d^{2\alpha}s],1}^d - \frac{1}{\sigma_d^{2(m-1)}r} X_{[d^{2\alpha}s],1}^d \right) \mathbb{1}_{\mathcal{A}_{[d^{2\alpha}s]}^d}. \end{aligned}$$

Proceeding as above, we find that

$$\begin{aligned} \mathbb{E} \left[(L_t^d - L_s^d)^4 \right] &\leq C \left((t-s)^2 + \mathbb{E} \left[\left(X_{[d^{2\alpha}t],1}^d - X_{[d^{2\alpha}s],1}^d \right)^4 \right] \right. \\ &\quad \left. + ([d^{2\alpha}s] - d^{2\alpha}s)^4 \mathbb{E} \left[\left(\frac{1}{\sigma_d^{2(m-1)}r} X_{[d^{2\alpha}s],1}^d \right)^4 \right] \right), \end{aligned}$$

and recalling that $|X_{[d^{2\alpha}s],1}^d| < \sigma_d^{2m}r/2$ we have that $|X_{[d^{2\alpha}s],1}^d| / (r\sigma_d^{2(m-1)}) < \sigma_d^2/2$. Using this and the same arguments as above, we have

$$\mathbb{E} \left[(L_t^d - L_s^d)^4 \right] \leq C \left((t-s)^2 + \mathbb{E} \left[\left(X_{[d^{2\alpha}t],1}^d - X_{[d^{2\alpha}s],1}^d \right)^4 \right] \right). \quad (4.27)$$

Case 3 If $|X_{[d^{2\alpha}t],1}^d| < \sigma_d^{2m}r/2$ and $|X_{[d^{2\alpha}s],1}^d| < \sigma_d^{2m}r/2$, then

$$\begin{aligned} L_t^d - L_s^d &= X_{[d^{2\alpha}t],1}^d + (d^{2\alpha}t - [d^{2\alpha}t]) \left(\sigma_d Z_{[d^{2\alpha}t],1}^d - \frac{1}{\sigma_d^{2(m-1)}r} X_{[d^{2\alpha}t],1}^d \right) \mathbb{1}_{\mathcal{A}_{[d^{2\alpha}t]}^d} \\ &\quad - X_{[d^{2\alpha}s],1}^d + ([d^{2\alpha}s] - d^{2\alpha}s) \left(\sigma_d Z_{[d^{2\alpha}s],1}^d - \frac{1}{\sigma_d^{2(m-1)}r} X_{[d^{2\alpha}s],1}^d \right) \mathbb{1}_{\mathcal{A}_{[d^{2\alpha}s]}^d}. \end{aligned}$$

Using the boundedness of moments of Gaussian distributions and of $X_{[d^{2\alpha}t],1}^d, X_{[d^{2\alpha}s],1}^d$, we have

$$\mathbb{E} \left[(L_t^d - L_s^d)^4 \right] \leq C \left((t-s)^2 + \mathbb{E} \left[\left(X_{[d^{2\alpha}t],1}^d - X_{[d^{2\alpha}s],1}^d \right)^4 \right] \right). \quad (4.28)$$

Putting (4.26), (4.27) and (4.28) together and using Lemma 1 below we obtain

$$\begin{aligned} \mathbb{E} \left[(L_t^d - L_s^d)^4 \right] &\leq C \left((t-s)^2 + \sum_{p=2}^4 \frac{(|d^{2\alpha}t| - [d^{2\alpha}s])^p}{d^{2\alpha p}} \right) \\ &\leq C(t-s)^2 + C \sum_{p=2}^4 \frac{d^{2\alpha p} (t-s)^p}{d^{2\alpha p}} \leq C(2+t+t^2)(t-s)^2, \end{aligned}$$

which concludes the proof. \square

We are now ready to state and prove Lemma 1:

Lemma 1. *There exists $C > 0$ such that for any $k_1, k_2 \in \mathbb{N}$ with $0 \leq k_1 < k_2$,*

$$\mathbb{E} \left[\left(X_{k_2,1}^d - X_{k_1,1}^d \right)^4 \right] \leq C \sum_{p=2}^4 \frac{(k_2 - k_1)^p}{d^{2\alpha p}},$$

where $\alpha = 1/3$.

Proof. Recalling the definition of the proposal in (4.24) and the definition of b_d in (4.20) we can write

$$\mathbb{E} \left[\left(X_{k_2,1}^d - X_{k_1,1}^d \right)^4 \right] = \mathbb{E} \left[\left(\sum_{k=k_1+1}^{k_2} \sigma_d b_d \left(X_{k-1,1}^d, Z_{k,1}^d \right) \mathbb{1}_{\mathcal{A}_k^d} \right)^4 \right].$$

Then, we expand all acceptance or rejection terms between k_1 and k_2 and use Hölder's inequality to obtain

$$\begin{aligned} \mathbb{E} \left[\left(X_{k_2,1}^d - X_{k_1,1}^d \right)^4 \right] &\leq C \left\{ \sigma_d^4 \mathbb{E} \left[\left(\sum_{k=k_1+1}^{k_2} b_d \left(X_{k-1,1}^d, Z_{k,1}^d \right) \right)^4 \right] \right. \\ &\quad \left. + \sigma_d^4 \mathbb{E} \left[\left(\sum_{k=k_1+1}^{k_2} b_d \left(X_{k-1,1}^d, Z_{k,1}^d \right) \mathbb{1}_{(\mathcal{A}_k^d)^c} \right)^4 \right] \right\}. \end{aligned}$$

Using again Hölder's inequality, for the first term we have

$$\begin{aligned}
 \mathbb{E} \left[\left(\sum_{k=k_1+1}^{k_2} b_d \left(X_{k-1,1}^d, Z_{k,1}^d \right) \right)^4 \right] &\leq C \left\{ \mathbb{E} \left[\left(\sum_{k=k_1+1}^{k_2} Z_{k,1}^d \right)^4 \right] \right. \\
 &+ \frac{\sigma_d^4}{2^4} \mathbb{E} \left[\left(\sum_{k=k_1+1}^{k_2} \operatorname{sgn} \left(X_{k-1,1}^d \right) \mathbb{1} \left\{ |X_{k-1,1}^d| \geq \sigma_d^{2m} r/2 \right\} \right)^4 \right] \\
 &+ \left. \mathbb{E} \left[\left(\sum_{k=k_1+1}^{k_2} \frac{1}{\sigma_d^{2m-1} r} X_{k-1,1}^d \mathbb{1} \left\{ |X_{k-1,1}^d| < \sigma_d^{2m} r/2 \right\} \right)^4 \right] \right\} \\
 &\leq C \left[3(k_2 - k_1)^2 + \frac{2\sigma_d^4}{2^4} (k_2 - k_1)^4 \right], \tag{4.29}
 \end{aligned}$$

where the last line follows using the moments of $Z_{k,1}^d$ and the boundedness of $X_{k-1,1}^d$ in the set $\{|X_{k-1,1}^d| < \sigma_d^{2m} r/2\}$.

Using a Binomial expansion of the rejection term, we obtain

$$\mathbb{E} \left[\left(\sum_{k=k_1+1}^{k_2} b_d \left(X_{k-1,1}^d, Z_{k,1}^d \right) \mathbb{1}_{(\mathbf{A}_k^d)^c} \right)^4 \right] = \sum \mathbb{E} \left[\prod_{i=1}^4 b_d \left(X_{m_i-1,1}^d, Z_{m_i,1}^d \right) \mathbb{1}_{(\mathbf{A}_{m_i}^d)^c} \right], \tag{4.30}$$

where the sum is over the quadruplets $(m_i)_{1 \leq i \leq 4}$ with $m_i \in \{k_1 + 1, \dots, k_2\}$.

We separate the terms in the sum according to their cardinality, let us denote, for $j \in \{1, \dots, 4\}$,

$$\mathcal{I}_j = \left\{ (m_1, \dots, m_4) \in \{k_1 + 1, \dots, k_2\}^4 : \#\{m_1, \dots, m_4\} = j \right\};$$

and define, for any $(m_1, \dots, m_4) \in \{k_1 + 1, \dots, k_2\}^4$, $\tilde{X}_0^d = X_0^d$ and for any $i \in \{1, \dots, d\}$,

$$\tilde{X}_{k+1,i}^d = \tilde{X}_{k,i}^d + \mathbb{1}_{\{m_1-1, \dots, m_4-1\}^c}(k) \mathbb{1}_{\tilde{\mathbf{A}}_{k+1}^d} \sigma_d b_d \left(\tilde{X}_{k,i}^d, Z_{k+1,i}^d \right),$$

where

$$\tilde{\mathbf{A}}_{k+1}^d = \left\{ U_{k+1} \leq \exp \left[\sum_{i=1}^d \phi_d \left(\tilde{X}_{k,i}^d, Z_{k+1,i}^d \right) \right] \right\}, \tag{4.31}$$

and ϕ_d in (4.21). Denote by \mathcal{F} the σ -algebra generated by the process $(\tilde{X}_k^d)_{k \geq 0}$ and observe that on the event $\bigcap_{j=1}^4 (\mathbf{A}_{m_j}^d)^c$, X_k^d is equal to \tilde{X}_k^d .

We consider now the terms in the sum (4.30).

(i) If $(m_1, \dots, m_4) \in \mathcal{I}_4$, then the m_j s are all distinct and

$$\mathbb{E} \left[\prod_{j=1}^4 b_d \left(X_{m_j-1,1}^d, Z_{m_j,1}^d \right) \mathbb{1}_{\left(A_{m_j}^d \right)^c} \middle| \mathcal{F} \right] = \mathbb{E} \left[\prod_{j=1}^4 b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right) \mathbb{1}_{\left(\tilde{A}_{m_j}^d \right)^c} \middle| \mathcal{F} \right].$$

However, $\{b_d(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d) \mathbb{1}_{(\tilde{A}_{m_j}^d)^c}\}_{j=1,\dots,4}$ are independent conditionally on \mathcal{F} . Thus,

$$\begin{aligned} \mathbb{E} \left[\prod_{j=1}^4 b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right) \mathbb{1}_{\left(\tilde{A}_{m_j}^d \right)^c} \middle| \mathcal{F} \right] &= \prod_{j=1}^4 \mathbb{E} \left[b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right) \mathbb{1}_{\left(\tilde{A}_{m_j}^d \right)^c} \middle| \mathcal{F} \right] \\ &= \prod_{j=1}^4 \mathbb{E} \left[b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right) \times \left(1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right) \middle| \mathcal{F} \right], \end{aligned}$$

by integrating the uniform variables U_{m_j} in (4.31).

Recalling the definition of b_d in (4.20), we can bound the expectation above with

$$\begin{aligned} &\left| \mathbb{E} \left[b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right) \left\{ 1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right\} \middle| \mathcal{F} \right] \right| \quad (4.32) \\ &\leq \left| \mathbb{E} \left[\left(\frac{\sigma_d}{2} \operatorname{sgn} \left(\tilde{X}_{m_j-1,1}^d \right) \mathbb{1}_{\left\{ |\tilde{X}_{m_j-1,1}^d| \geq \sigma_d^{2m} r/2 \right\}} \right. \right. \right. \\ &\quad \left. \left. - \frac{1}{\sigma_d^{2m-1} r} \tilde{X}_{m_j-1,1}^d \mathbb{1}_{\left\{ |\tilde{X}_{m_j-1,1}^d| < \sigma_d^{2m} r/2 \right\}} \right) \right. \\ &\quad \left. \times \left\{ 1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right\} \middle| \mathcal{F} \right] \right| \\ &\quad + \left| \mathbb{E} \left[Z_{m_j,1}^d \left\{ 1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right\} \middle| \mathcal{F} \right] \right|. \end{aligned}$$

For the first one, we use the boundedness of the sgn function and of $\tilde{X}_{m_j-1,1}^d$ in the set $\{|\tilde{X}_{m_j-1,1}^d| \leq \sigma_d^{2m} r/2\}$ to obtain

$$\begin{aligned} &\left| \mathbb{E} \left[\left(\frac{\sigma_d}{2} \operatorname{sgn} \left(\tilde{X}_{m_j-1,1}^d \right) \mathbb{1}_{\left\{ |\tilde{X}_{m_j-1,1}^d| \geq \sigma_d^{2m} r/2 \right\}} - \frac{\tilde{X}_{m_j-1,1}^d}{\sigma_d^{2m-1} r} \mathbb{1}_{\left\{ |\tilde{X}_{m_j-1,1}^d| < \sigma_d^{2m} r/2 \right\}} \right) \right. \right. \\ &\quad \left. \left. \times \left\{ 1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right\} \middle| \mathcal{F} \right] \right| \\ &\leq \frac{\sigma_d}{2} \mathbb{E} \left[\left\| \left\{ 1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right\} \middle| \mathcal{F} \right\| \right] \leq \frac{\sigma_d}{2}. \quad (4.33) \end{aligned}$$

We can write the second term as

$$\begin{aligned} & \mathbb{E} \left[Z_{m_j,1}^d \left(1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right) \Big| \mathcal{F} \right] \\ &= \mathbb{E} \left[\mathcal{G} \left(\tilde{X}_{m_j-1,1}^d, \sum_{i=2}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \Big| \mathcal{F} \right], \end{aligned}$$

where we define $\mathcal{G}(a, b) = \mathbb{E} \left[Z (1 - \exp(\phi_d(a, Z) + b))_+ \right]$ with Z a standard Gaussian. Because the function $x \mapsto (1 - \exp(x))_+$ is 1-Lipschitz, we have, using Cauchy-Schwarz and Lemma 3 in Appendix 4.10.2,

$$\begin{aligned} \left| \mathbb{E} \left[Z (1 - \exp(\phi_d(a, Z) + b))_+ \right] - \mathbb{E} \left[Z (1 - \exp(b))_+ \right] \right| &\leq \mathbb{E} [|Z| |\phi_d(a, Z)|] \\ &\leq \mathbb{E} [Z^2]^{1/2} \mathbb{E} [\phi_d(a, Z)^2]^{1/2} \\ &\leq \mathbb{E} [\phi_d(a, Z)^2]^{1/2} \\ &\leq Cd^{-\alpha}. \end{aligned}$$

However, $\mathbb{E} \left[Z (1 - \exp(b))_+ \right] = \mathbb{E} [Z] (1 - \exp(b))_+ = 0$, and therefore

$$\left| \mathbb{E} \left[\mathcal{G} \left(\tilde{X}_{m_j-1,1}^d, \sum_{i=2}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \Big| \mathcal{F} \right] \right| \leq Cd^{-\alpha}. \quad (4.34)$$

Combining equations (4.32), (4.33) and (4.34) and recalling that $\sigma_d = \ell d^{-\alpha}$, we have

$$\left| \mathbb{E} \left[b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right) \left\{ 1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right\}_+ \Big| \mathcal{F} \right] \right| \leq Cd^{-\alpha}, \quad (4.35)$$

from which follows that

$$\begin{aligned} \sum_{(m_1, \dots, m_4) \in \mathcal{I}_4} \left| \mathbb{E} \left[\prod_{i=1}^4 b_d \left(X_{m_i-1,1}^d, Z_{m_i,1}^d \right) \mathbb{1}_{(\mathcal{A}_{m_i}^d)^c} \right] \right| &\leq \sum_{(m_1, \dots, m_4) \in \mathcal{I}_4} \mathbb{E} \left[\prod_{j=1}^4 \frac{C}{d^\alpha} \right] \\ &\leq \binom{k_2 - k_1}{4} \frac{C}{d^{4\alpha}} \leq C \frac{(k_2 - k_1)^4}{d^{4\alpha}}, \end{aligned} \quad (4.36)$$

using that $|\mathcal{I}_4| = \binom{k_2 - k_1}{4}$.

(ii) If $(m_1, \dots, m_4) \in \mathcal{I}_3$, only three of the m_i s take distinct values; proceeding as in case

(i), we have

$$\begin{aligned} & \left| \mathbb{E} \left[\prod_{j=1}^3 b_d \left(X_{m_j-1,1}^d, Z_{m_j,1}^d \right)^{1+\delta_{1,j}} \mathbb{1}_{\left(A_{m_j}^d \right)^c} \middle| \mathcal{F} \right] \right| \\ &= \prod_{j=1}^3 \left| \mathbb{E} \left[b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right)^{1+\delta_{1,j}} \left\{ 1 - \exp \left(\sum_{i=1}^d \phi_d \left(\tilde{X}_{m_j-1,i}^d, Z_{m_j,i}^d \right) \right) \right\}_+ \middle| \mathcal{F} \right] \right|, \end{aligned}$$

where $\delta_{1,j}$ denotes a Dirac's delta. For the terms $j \neq 1$, we use (4.35), while for the term $j = 1$ we bound the indicator function by 1 to obtain

$$\begin{aligned} & \left| \mathbb{E} \left[\prod_{j=1}^3 b_d \left(X_{m_j-1,1}^d, Z_{m_j,1}^d \right)^{1+\delta_{1,j}} \mathbb{1}_{\left(A_{m_j}^d \right)^c} \middle| \mathcal{F} \right] \right| \\ & \leq \left| \mathbb{E} \left[b_d \left(\tilde{X}_{m_1-1,1}^d, Z_{m_1,1}^d \right)^2 \middle| \mathcal{F} \right] \right| \prod_{j=2}^3 \frac{C}{d^\alpha} \\ & \leq \left(3 + \frac{2\sigma_d^2}{2^2 d^{2\alpha}} \right) \frac{C^2}{d^{2\alpha}} \leq C \frac{1}{d^{2\alpha}}, \end{aligned}$$

where the second-to-last inequality follows using the same approach taken for (4.29) and recalling that $\sigma_d = \ell d^{-\alpha}$. Hence,

$$\begin{aligned} & \sum_{(m_1, \dots, m_4) \in \mathcal{I}_3} \left| \mathbb{E} \left[\prod_{i=1}^4 b_d \left(X_{m_i-1,1}^d, Z_{m_i,1}^d \right) \mathbb{1}_{\left(A_{m_i}^d \right)^c} \right] \right| \tag{4.37} \\ & \leq C \binom{k_2 - k_1}{3} \frac{1}{d^{2\alpha}} \leq C \frac{(k_2 - k_1)^3}{d^{2\alpha}}. \end{aligned}$$

(iii) If $(m_1, \dots, m_4) \in \mathcal{I}_2$, we have two different cases: the m_i s take the two values twice or three m_i s have the same value. For the first one, we have, bounding the indicator function with 1,

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E} \left[\prod_{j=1}^2 b_d \left(X_{m_j-1,1}^d, Z_{m_j,1}^d \right)^2 \mathbb{1}_{\left(A_{m_j}^d \right)^c} \middle| \mathcal{F} \right] \right] \\ & \leq \mathbb{E} \left[\prod_{j=1}^2 \mathbb{E} \left[b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right)^2 \middle| \mathcal{F} \right] \right]. \end{aligned}$$

Since, conditionally on \mathcal{F} , the random variables inside the expectation are normals with bounded mean and variance 1, we have, using the same approach taken for (4.29),

$$\mathbb{E} \left[\prod_{j=1}^2 \mathbb{E} \left[b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right)^2 \middle| \mathcal{F} \right] \right] \leq \left(1 + \frac{2\sigma_d^2}{2^2} \right)^2 \leq C.$$

The second case follows similarly

$$\begin{aligned} & \left| \mathbb{E} \left[\mathbb{E} \left[\prod_{j=1}^2 b_d \left(X_{m_j-1,1}^d, Z_{m_j,1}^d \right)^{1+2\delta_{1,j}} \mathbb{1}_{(\mathbf{A}_{m_j}^d)^c} \middle| \mathcal{F} \right] \right] \right| \\ & \leq \mathbb{E} \left[\mathbb{E} \left[\prod_{j=1}^2 \left| b_d \left(\tilde{X}_{m_j-1,1}^d, Z_{m_j,1}^d \right) \right|^{1+2\delta_{1,j}} \middle| \mathcal{F} \right] \right] \leq C, \end{aligned}$$

where $\delta_{1,j}$ denotes a Dirac's delta. Therefore,

$$\begin{aligned} & \sum_{(m_1, \dots, m_4) \in \mathcal{I}_2} \left| \mathbb{E} \left[\prod_{i=1}^4 \left\{ b_d \left(X_{m_i-1,1}^d, Z_{m_i,1}^d \right) \mathbb{1}_{(\mathbf{A}_{m_i}^d)^c} \right\} \right] \right| \tag{4.38} \\ & \leq C \left(\binom{4}{2} + \binom{4}{3} \right) \binom{k_2 - k_1}{2} \leq C(k_2 - k_1)^2. \end{aligned}$$

(iv) If $(m_1, \dots, m_4) \in \mathcal{I}_1$ (i.e. all m_i s take the same value), we bound the indicator function by 1 and, using the same approach taken for (4.29), we find

$$\mathbb{E} \left[b_d \left(X_{m_1-1,1}^d, Z_{m_1,1}^d \right)^4 \mathbb{1}_{(\mathbf{A}_{m_1}^d)^c} \right] \leq C \left(3 + \frac{2\sigma_d^4}{2^4} \right) \leq C,$$

since $\sigma_d = \ell d^{-\alpha}$ and $d \in \mathbb{N}$. Hence,

$$\sum_{(m_1, \dots, m_4) \in \mathcal{I}_1} \left| \mathbb{E} \left[\prod_{i=1}^4 b_d \left(X_{m_i-1,1}^d, Z_{m_i,1}^d \right) \mathbb{1}_{(\mathbf{A}_{m_i}^d)^c} \right] \right| \leq C \binom{k_2 - k_1}{1} = C(k_2 - k_1). \tag{4.39}$$

The result follows combining (4.36), (4.37), (4.38) and (4.39) in (4.30).

□

4.6.3 Proof of Proposition 4.6

We start by proving the following lemma.

Lemma 2. *Let ν be a limit point of the sequence of laws $(\nu_d)_{d \geq 1}$ of $\{(L_t^d)_{t \geq 0} : d \in \mathbb{N}^*\}$. Then for any $t \geq 0$, the pushforward measure of ν by W_t is $\pi^L(dx) = \exp(-|x|)dx/2$.*

Proof. Using (4.25), we have

$$\begin{aligned}
 & \mathbb{E} \left[\left| L_t^d - X_{\lfloor d^{2\alpha} t \rfloor, 1}^d \right| \right] \\
 & \leq \mathbb{E} \left[\left| (d^{2\alpha} t - \lfloor d^{2\alpha} t \rfloor) \left[\sigma_d Z_{\lfloor d^{2\alpha} t \rfloor, 1}^d - \frac{\sigma_d^2}{2} \operatorname{sgn}(X_{\lfloor d^{2\alpha} t \rfloor, 1}^d) \right] \mathbb{1} \left\{ |X_{\lfloor d^{2\alpha} t \rfloor, 1}^d| \geq \sigma_d^{2m} r / 2 \right\} \mathbb{1}_{\mathcal{A}_{\lfloor d^{2\alpha} t \rfloor}^d} \right| \right] \\
 & + \mathbb{E} \left[\left| (d^{2\alpha} t - \lfloor d^{2\alpha} t \rfloor) \left[\sigma_d Z_{\lfloor d^{2\alpha} t \rfloor, 1}^d - \frac{1}{\sigma_d^{2(m-1)} r} X_{\lfloor d^{2\alpha} t \rfloor, 1}^d \right] \mathbb{1} \left\{ |X_{\lfloor d^{2\alpha} t \rfloor, 1}^d| < \sigma_d^{2m} r / 2 \right\} \mathbb{1}_{\mathcal{A}_{\lfloor d^{2\alpha} t \rfloor}^d} \right| \right] \\
 & \leq (d^{2\alpha} t - \lfloor d^{2\alpha} t \rfloor) \left(\sigma_d \mathbb{E} \left[\left| Z_{\lfloor d^{2\alpha} t \rfloor, 1}^d \right| \right] + \frac{\sigma_d^2}{2} \mathbb{E} \left[\left| \operatorname{sgn}(X_{\lfloor d^{2\alpha} t \rfloor, 1}^d) \right| \right] \right. \\
 & \quad \left. + \frac{1}{\sigma_d^{2(m-1)} r} \mathbb{E} \left[\left| X_{\lfloor d^{2\alpha} t \rfloor, 1}^d \right| \mathbb{1} \left\{ |X_{\lfloor d^{2\alpha} t \rfloor, 1}^d| < \sigma_d^{2m} r / 2 \right\} \right] \right) \\
 & \leq (d^{2\alpha} t - \lfloor d^{2\alpha} t \rfloor) \left(\frac{\ell}{d^\alpha} \mathbb{E} \left[(Z_{\lfloor d^{2\alpha} t \rfloor, 1}^d)^2 \right]^{1/2} + \frac{\ell^2}{2d^{2\alpha}} + \frac{1}{\sigma_d^{2(m-1)} r} \mathbb{E} \left[\frac{\sigma_d^{2m} r}{2} \right] \right) \\
 & \leq (d^{2\alpha} t - \lfloor d^{2\alpha} t \rfloor) \left(\frac{\ell}{d^\alpha} + \frac{\ell^2}{2d^{2\alpha}} + \frac{\ell^2}{2d^{2\alpha}} \right) \leq \frac{C}{d^\alpha},
 \end{aligned}$$

where we used Cauchy-Schwarz inequality and the fact that the moments of $Z_{\lfloor d^{2\alpha} t \rfloor, 1}^d$ are bounded. The above guarantees that,

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\left| L_t^d - X_{\lfloor d^{2\alpha} t \rfloor, 1}^d \right| \right] = 0.$$

As $(\nu_d)_{d \geq 1}$ converges weakly towards ν , for any Lipschitz bounded function $\psi : \mathbb{R} \rightarrow \mathbb{R}$,

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\psi \left(X_{\lfloor d^{2\alpha} t \rfloor, 1}^d \right) \right] = \lim_{d \rightarrow \infty} \mathbb{E} \left[\psi \left(L_t^d \right) \right] = \mathbb{E}^\nu \left[\psi(W_t) \right].$$

The result follows since $X_{\lfloor d^{2\alpha} t \rfloor, 1}^d$ is distributed according to $\pi^L(dx) = \exp(-|x|)dx/2$ for any $t \geq 0$ and $d \in \mathbb{N}$. \square

We are now ready to prove Proposition 4.6:

Proof of Proposition 4.6. Let ν be a limit point of $(\nu_d)_{d \geq 1}$. We start by showing that if for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$, $m \in \mathbb{N}$, any bounded and continuous mapping $\rho : \mathbb{R}^m \rightarrow \mathbb{R}$ and any $0 \leq t_1 \leq \dots \leq t_m \leq s \leq t$, ν satisfies

$$\mathbb{E}^\nu \left[\left(V(W_t) - V(W_s) - \int_s^t LV(W_u) du \right) \rho(W_{t_1}, \dots, W_{t_m}) \right] = 0, \quad (4.40)$$

then ν is a solution to the martingale problem associated with L .

Let \mathfrak{F}_s denote the σ -algebra generated by

$\{\rho(W_{t_1}, \dots, W_{t_m}) : m \in \mathbb{N}, \rho : \mathbb{R}^m \rightarrow \mathbb{R} \text{ bounded and continuous, and } 0 \leq t_1 \leq \dots \leq t_m \leq s\}$.

Then,

$$\mathbb{E}^\nu \left[V(W_t) - V(W_s) - \int_s^t LV(W_u) du \middle| \mathfrak{F}_s \right] = 0 ,$$

showing that the process

$$\left(V(W_t) - V(W_0) - \int_0^t LV(W_u) du \right)_{t \geq 0}$$

is a martingale w.r.t. ν and the filtration $(\mathfrak{F}_t)_{t \geq 0}$.

To prove (4.40), it is enough to show that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$, $m \in \mathbb{N}$ and any bounded and continuous mapping $\rho : \mathbb{R}^m \rightarrow \mathbb{R}$ and any $0 \leq t_1 \leq \dots \leq t_m \leq s \leq t$, the mapping

$$\Psi_{s,t} : w \mapsto \left(V(w_t) - V(w_s) - \int_s^t LV(w_u) du \right) \rho(w_{t_1}, \dots, w_{t_m}) ,$$

is continuous on a ν -almost sure subset of $C(\mathbb{R}_+, \mathbb{R})$. Let

$$\mathbf{W} = \{w \in C(\mathbb{R}_+, \mathbb{R}) : w_u \neq 0 \text{ for almost any } u \in [s, t]\} .$$

Since $w \in \mathbf{W}^c$ if and only if $\int_s^t \mathbb{1}_{\{0\}}(w_u) du > 0$, using Lemma 2 and the Fubini–Tonelli’s theorem,

$$\mathbb{E}^\nu \left[\int_s^t \mathbb{1}_{\{0\}}(W_u) du \right] = \int_s^t \mathbb{E}^\nu \left[\mathbb{1}_{\{0\}}(W_u) \right] du = \int_s^t \pi^L(\{0\}) du = 0 ,$$

and we have that $\nu(\mathbf{W}^c) = 0$.

Since $w \mapsto w_u$ is continuous for any $u \geq 0$, so are $w \mapsto V(w_u)$ and $w \mapsto \rho(w_{t_1}, \dots, w_{t_m})$. Thus, it is enough to prove that the mapping $w \mapsto \int_s^t LV(w_u) du$ is continuous. Let $(w^n)_{n \geq 0}$ be a sequence in $C(\mathbb{R}_+, \mathbb{R})$ that converges to $w \in \mathbf{W}$ in the uniform topology on compact sets. Let u be such that $w_u \neq 0$, therefore, since the sgn function is continuous in a neighbourhood of w_u , $\lim_{n \rightarrow \infty} LV(w_u^n) = LV(w_u)$, thus $\lim_{n \rightarrow \infty} LV(w_u^n) = LV(w_u)$ for almost any $u \in [s, t]$. Finally, using the boundedness of the sequence $(LV(w_u^n))_{n \geq 0}$ and Lebesgue’s dominated convergence theorem,

$$\lim_{n \rightarrow \infty} \int_s^t LV(w_u^n) du = \int_s^t LV(w_u) du ,$$

which proves that the mappings $\Psi_{s,t}$ are continuous on \mathbf{W} . □

4.6.4 Proof of Theorem 4.7

Let us introduce, for any $n \in \mathbb{N}$, $\mathcal{F}_{n,1}^d = \sigma(\{X_{k,1}^d, 0 \leq k \leq n\})$, the σ -algebra generated by the first components of $\{X_k^d \mid 0 \leq k \leq n\}$. We also introduce for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$

$$\begin{aligned} M_n^d(V) &= \frac{\ell}{d^\alpha} \sum_{k=0}^{n-1} V'(X_{k,1}^d) \\ &\quad \times \left(b_d \left(X_{k,1}^d, Z_{k+1,1}^d \right) \mathbb{1}_{\mathbb{A}_{k+1}^d} - \mathbb{E} \left[b_d \left(X_{k,1}^d, Z_{k+1,1}^d \right) \mathbb{1}_{\mathbb{A}_{k+1}^d} \middle| \mathcal{F}_{k,1}^d \right] \right) \\ &+ \frac{\ell^2}{2d^{2\alpha}} \sum_{k=0}^{n-1} V''(X_{k,1}^d) \\ &\quad \times \left(b_d \left(X_{k,1}^d, Z_{k+1,1}^d \right)^2 \mathbb{1}_{\mathbb{A}_{k+1}^d} - \mathbb{E} \left[b_d \left(X_{k,1}^d, Z_{k+1,1}^d \right)^2 \mathbb{1}_{\mathbb{A}_{k+1}^d} \middle| \mathcal{F}_{k,1}^d \right] \right). \end{aligned} \quad (4.41)$$

where b_d is defined in (4.20).

The proof of Theorem 4.7 follows using the sufficient condition in Proposition 4.6, the tightness of the sequence $(\nu_d)_{d \geq 1}$ established in Proposition 4.5 and Proposition 4.9 below.

Proof. Using Proposition 4.5, Proposition 4.6 and Proposition 4.9 below, it is enough to show that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$, $m \geq 1$, any $0 \leq t_1 \leq \dots \leq t_m \leq s \leq t$ and any bounded and continuous mapping $\rho : \mathbb{R}^m \rightarrow \mathbb{R}$,

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\left(M_{\lceil d^{2\alpha} t \rceil}^d(V) - M_{\lceil d^{2\alpha} s \rceil}^d(V) \right) \rho(L_{t_1}^d, \dots, L_{t_m}^d) \right] = 0,$$

where, for any $n \geq 1$, $M_n^d(V)$ is given by (4.41). However, this is straightforwardly obtained by taking successively the conditional expectations with respect to $\mathcal{F}_{k,1}^d$ for $k = \lceil d^{2\alpha} t \rceil, \dots, \lceil d^{2\alpha} s \rceil$. \square

Proposition 4.9. *For any $0 \leq s \leq t$, $V \in C_c(\mathbb{R}, \mathbb{R})$ we have*

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\left[V(L_t^d) - V(L_s^d) - \int_s^t LV(L_u^d) du - \left(M_{\lceil d^{2\alpha} t \rceil}^d(V) - M_{\lceil d^{2\alpha} s \rceil}^d(V) \right) \right] \right] = 0, \quad (4.42)$$

where $(L_t^d)_{t \geq 0}$ is defined in (4.11).

Proof. The process $(L_t^d)_{t \geq 0}$ is piecewise linear, thus it has finite variation. For any $\tau \geq 0$, we define

$$dL_\tau^d = d^{2\alpha} \sigma_d b_d \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d, Z_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) \mathbb{1}_{\mathbb{A}_{\lceil d^{2\alpha} \tau \rceil}^d} d\tau.$$

Since $\sigma_d = \ell d^{-\alpha}$ with $\alpha = 1/3$ and using the fundamental theorem of calculus for piecewise C^1 maps

$$V(L_t^d) - V(L_s^d) = \ell d^\alpha \int_s^t V'(L_\tau^d) b_d \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d, Z_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) \mathbb{1}_{\mathbb{A}_{\lceil d^{2\alpha} \tau \rceil}^d} d\tau, \quad (4.43)$$

4.6. Proof of the Result for the Laplace distribution

where b_d is defined in (4.20). A Taylor expansion of V' with Lagrange remainder about $X_{[d^{2\alpha}\tau],1}^d$ gives

$$\begin{aligned} V'(L_\tau^d) &= V'(X_{[d^{2\alpha}\tau],1}^d) \\ &\quad + \frac{\ell}{d^\alpha} (d^{2\alpha}\tau - [d^{2\alpha}\tau]) V''(X_{[d^{2\alpha}\tau],1}^d) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d) \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} \\ &\quad + \frac{\ell^2}{2d^{2\alpha}} (d^{2\alpha}\tau - [d^{2\alpha}\tau])^2 V^{(3)}(\chi_\tau) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d) \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d}, \end{aligned}$$

where for any point $\tau \in [s, t]$, there exists $\chi_\tau \in [X_{[d^{2\alpha}\tau],1}^d, L_\tau^d]$. Substituting the above into (4.43) we obtain

$$\begin{aligned} V(L_t^d) - V(L_s^d) &= \ell d^\alpha \int_s^t V'(X_{[d^{2\alpha}\tau],1}^d) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d) \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau \\ &\quad + \ell^2 \int_s^t (d^{2\alpha}\tau - [d^{2\alpha}\tau]) V''(X_{[d^{2\alpha}\tau],1}^d) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau \quad (4.44) \\ &\quad + \frac{\ell^3}{2d^\alpha} \int_s^t (d^{2\alpha}\tau - [d^{2\alpha}\tau])^2 V^{(3)}(\chi_\tau) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d)^3 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau. \end{aligned}$$

Since $V^{(3)}$ is bounded, using Fubini-Tonelli's theorem and recalling the definition of b_d in (4.20), we have that

$$\begin{aligned} &\frac{\ell^3}{2d^\alpha} \mathbb{E} \left[\left| \int_s^t (d^{2\alpha}\tau - [d^{2\alpha}\tau])^2 V^{(3)}(\chi_\tau) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d)^3 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau \right| \right] \\ &\leq C \frac{\ell^3}{2d^\alpha} \int_s^t \mathbb{E} \left[\left(|Z_{[d^{2\alpha}\tau],1}^d| + \frac{\ell}{2d^\alpha} \right)^3 \right] d\tau \xrightarrow{d \rightarrow \infty} 0, \end{aligned}$$

since the moments of $Z_{[d^{2\alpha}\tau],1}^d$ are bounded.

For the second term in (4.44), we observe that most of the integrand is piecewise constant since the process $X_{[d^{2\alpha}\tau],1}^d$ evolves in discrete time. Then, for any integer $d^{2\alpha}s \leq k \leq d^{2\alpha}t - 1$,

$$\begin{aligned} &\int_{k/d^{2\alpha}}^{(k+1)/d^{2\alpha}} (d^{2\alpha}\tau - [d^{2\alpha}\tau]) V''(X_{[d^{2\alpha}\tau],1}^d) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau \\ &= \frac{1}{2d^{2\alpha}} V''(X_{k,1}^d) b_d(X_{k,1}^d, Z_{k+1,1}^d)^2 \mathbb{1}_{\mathbb{A}_{k+1}^d} \\ &= \frac{1}{2} \int_{k/d^{2\alpha}}^{(k+1)/d^{2\alpha}} V''(X_{[d^{2\alpha}\tau],1}^d) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau. \end{aligned}$$

Thus, we can write

$$\begin{aligned} I &= \int_s^t (d^{2\alpha}\tau - [d^{2\alpha}\tau]) V''(X_{[d^{2\alpha}\tau],1}^d) b_d(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau \\ &= I_1 + I_2, \end{aligned}$$

where we define

$$I_2 := \frac{1}{2} \int_s^t V'' \left(X_{[d^{2\alpha}\tau],1}^d \right) b_d \left(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d \right)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau,$$

and

$$I_1 := \left[\int_s^{\lceil d^{2\alpha}s \rceil / d^{2\alpha}} + \int_{\lfloor d^{2\alpha}t \rfloor / d^{2\alpha}}^t \right] \left(d^{2\alpha}\tau - \lfloor d^{2\alpha}\tau \rfloor - \frac{1}{2} \right) V'' \left(X_{[d^{2\alpha}\tau],1}^d \right) \\ \times b_d \left(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d \right)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} d\tau.$$

In addition, we have

$$I_1 = \frac{1}{2d^{2\alpha}} \left(d^{2\alpha}s - \lfloor d^{2\alpha}s \rfloor \right) \left(\lceil d^{2\alpha}s \rceil - d^{2\alpha}s \right) V'' \left(X_{[d^{2\alpha}s],1}^d \right) b_d \left(X_{[d^{2\alpha}s],1}^d, Z_{[d^{2\alpha}s],1}^d \right)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}s]}^d} \\ + \frac{1}{2d^{2\alpha}} \left(d^{2\alpha}t - \lfloor d^{2\alpha}t \rfloor \right) \left(\lceil d^{2\alpha}t \rceil - d^{2\alpha}t \right) V'' \left(X_{[d^{2\alpha}t],1}^d \right) b_d \left(X_{[d^{2\alpha}t],1}^d, Z_{[d^{2\alpha}t],1}^d \right)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}t]}^d},$$

and, since V'' and the moments of $Z_{[d^{2\alpha}t],1}^d$ are bounded, $\lim_{d \rightarrow \infty} \mathbb{E}[I_1] = 0$. Thus,

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\left| V \left(L_t^d \right) - V \left(L_s^d \right) - I_{s,t} \right| \right] = 0,$$

where

$$I_{s,t} = \int_s^t \left\{ \ell d^\alpha V' \left(X_{[d^{2\alpha}\tau],1}^d \right) b_d \left(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d \right) \right. \\ \left. + \frac{\ell^2}{2} V'' \left(X_{[d^{2\alpha}\tau],1}^d \right) b_d \left(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d \right)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} \right\} d\tau. \quad (4.45)$$

Next, we use (4.18) and write

$$\int_s^t \text{LV} \left(L_\tau^d \right) d\tau = \int_s^t \frac{h^L(\ell)}{2} \left[V'' \left(X_{[d^{2\alpha}\tau],1}^d \right) - \text{sgn} \left(X_{[d^{2\alpha}\tau],1}^d \right) V' \left(X_{[d^{2\alpha}\tau],1}^d \right) \right] d\tau - T_3^d, \quad (4.46)$$

where we define

$$T_3^d = \int_s^t \left(\text{LV} \left(X_{[d^{2\alpha}\tau],1}^d \right) - \text{LV} \left(L_\tau^d \right) \right) d\tau.$$

Finally, we write the difference $M_{[d^{2\alpha}t]}^d(V) - M_{[d^{2\alpha}s]}^d(V)$ as the integral of a piecewise constant function

$$M_{[d^{2\alpha}t]}^d(V) - M_{[d^{2\alpha}s]}^d(V) = I_{s,t} \\ - \int_s^t \left(\ell d^\alpha V' \left(X_{[d^{2\alpha}\tau],1}^d \right) \mathbb{E} \left[b_d \left(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d \right) \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} \middle| \mathcal{F}_{[d^{2\alpha}\tau],1}^d \right] \right. \\ \left. + \frac{\ell^2}{2} V'' \left(X_{[d^{2\alpha}\tau],1}^d \right) \mathbb{E} \left[b_d \left(X_{[d^{2\alpha}\tau],1}^d, Z_{[d^{2\alpha}\tau],1}^d \right)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha}\tau]}^d} \middle| \mathcal{F}_{[d^{2\alpha}\tau],1}^d \right] \right) d\tau \\ - T_4^d - T_5^d, \quad (4.47)$$

4.6. Proof of the Result for the Laplace distribution

where T_4^d and T_5^d account for the difference between the sum in (4.41) and the integral, and are defined as

$$\begin{aligned} T_4^d &= -\frac{\ell}{d^\alpha} \left(\lceil d^{2\alpha} t \rceil - d^{2\alpha} t \right) V' \left(X_{\lceil d^{2\alpha} t \rceil, 1}^d \right) \left\{ b_d \left(X_{\lceil d^{2\alpha} t \rceil, 1}^d, Z_{\lceil d^{2\alpha} t \rceil, 1}^d \right) \mathbb{1}_{\mathcal{A}_{\lceil d^{2\alpha} t \rceil}^d} \right. \\ &\quad \left. - \mathbb{E} \left[b_d \left(X_{\lceil d^{2\alpha} t \rceil, 1}^d, Z_{\lceil d^{2\alpha} t \rceil, 1}^d \right) \mathbb{1}_{\mathcal{A}_{\lceil d^{2\alpha} t \rceil}^d} \middle| \mathcal{F}_{\lceil d^{2\alpha} t \rceil, 1}^d \right] \right\} \\ &\quad - \frac{\ell^2}{2d^{2\alpha}} \left(\lceil d^{2\alpha} t \rceil - d^{2\alpha} t \right) V'' \left(X_{\lceil d^{2\alpha} t \rceil, 1}^d \right) \left\{ b_d \left(X_{\lceil d^{2\alpha} t \rceil, 1}^d, Z_{\lceil d^{2\alpha} t \rceil, 1}^d \right)^2 \mathbb{1}_{\mathcal{A}_{\lceil d^{2\alpha} t \rceil}^d} \right. \\ &\quad \left. - \mathbb{E} \left[b_d \left(X_{\lceil d^{2\alpha} t \rceil, 1}^d, Z_{\lceil d^{2\alpha} t \rceil, 1}^d \right)^2 \mathbb{1}_{\mathcal{A}_{\lceil d^{2\alpha} t \rceil}^d} \middle| \mathcal{F}_{\lceil d^{2\alpha} t \rceil, 1}^d \right] \right\}, \end{aligned}$$

and $T_5^d = -T_4^d$ with t substituted by s . Putting (4.45), (4.46) and (4.47) together we obtain

$$I_{s,t} - \int_s^t \text{LV} \left(L_\tau^d \right) d\tau - \left(M_{\lceil d^{2\alpha} t \rceil}^d(V) - M_{\lceil d^{2\alpha} s \rceil}^d(V) \right) = T_1^d + T_2^d + T_3^d + T_4^d + T_5^d,$$

where T_1^d takes into account all the terms involving $V'(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d)$, and T_2^d the terms involving $V''(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d)$:

$$\begin{aligned} T_1^d &= \int_s^t V' \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) \\ &\quad \times \left\{ \ell d^\alpha \mathbb{E} \left[b_d \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d, Z_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) \mathbb{1}_{\mathcal{A}_{\lceil d^{2\alpha} \tau \rceil}^d} \middle| \mathcal{F}_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right] + \frac{h^L(\ell)}{2} \text{sgn} \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) \right\} d\tau, \\ T_2^d &= \int_s^t V'' \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) \\ &\quad \times \left\{ \frac{\ell^2}{2} \mathbb{E} \left[b_d \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d, Z_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right)^2 \mathbb{1}_{\mathcal{A}_{\lceil d^{2\alpha} \tau \rceil}^d} \middle| \mathcal{F}_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right] - \frac{h^L(\ell)}{2} \right\} d\tau. \end{aligned}$$

To obtain (4.42) it is then sufficient to prove that for any $1 \leq i \leq 5$, $\lim_{d \rightarrow \infty} \mathbb{E} \left[|T_i^d| \right] = 0$.

Since V', V'' are bounded and b_d is bounded in expectation because the moments of $Z_{\lceil d^{2\alpha} \tau \rceil, 1}^d$ are bounded, it is easy to show that $\lim_{d \rightarrow \infty} \mathbb{E} \left[|T_i^d| \right] = 0$ for $i = 4, 5$. For T_3^d , we write $T_3^d = h^L(\ell)(T_{3,1}^d - T_{3,2}^d)/2$, where

$$\begin{aligned} T_{3,1}^d &= \int_s^t \left\{ V'' \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) - V'' \left(L_\tau^d \right) \right\} d\tau, \\ T_{3,2}^d &= \int_s^t \left\{ \text{sgn} \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) V' \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) - \text{sgn} \left(L_\tau^d \right) V' \left(L_\tau^d \right) \right\} d\tau. \end{aligned}$$

Using Fubini-Tonelli's theorem, the convergence of $X_{\lceil d^{2\alpha} \tau \rceil, 1}^d$ to L_τ^d in Lemma 2 and Lebesgue's dominated convergence theorem we obtain

$$\mathbb{E} \left[|T_{3,1}^d| \right] \leq \int_s^t \mathbb{E} \left[\left| V'' \left(X_{\lceil d^{2\alpha} \tau \rceil, 1}^d \right) - V'' \left(L_\tau^d \right) \right| \right] d\tau \xrightarrow{d \rightarrow \infty} 0.$$

We can further decompose $T_{3,2}^d$ as

$$\begin{aligned} T_{3,2}^d &= \int_s^t \left\{ \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) - \operatorname{sgn} \left(L_\tau^d \right) \right\} V' \left(X_{[d^{2\alpha_\tau],1}^d} \right) d\tau \\ &\quad + \int_s^t \operatorname{sgn} \left(L_\tau^d \right) \left\{ V' \left(X_{[d^{2\alpha_\tau],1}^d} \right) - V' \left(L_\tau^d \right) \right\} d\tau . \end{aligned}$$

Proceeding as for $T_{3,1}^d$, it is easy to show that the second integral converges to 0 as $d \rightarrow \infty$. We then bound the first integral by

$$\begin{aligned} &\mathbb{E} \left[\left| \int_s^t \left\{ \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) - \operatorname{sgn} \left(L_\tau^d \right) \right\} V' \left(X_{[d^{2\alpha_\tau],1}^d} \right) d\tau \right| \right] \\ &\leq C \int_s^t \mathbb{E} \left[\left| \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) - \operatorname{sgn} \left(L_\tau^d \right) \right| \right] d\tau . \end{aligned}$$

However, since $\{\operatorname{sgn}(X_{[d^{2\alpha_\tau],1}^d]) \neq \operatorname{sgn}(L_\tau^d)\} \subset \{\operatorname{sgn}(X_{[d^{2\alpha_\tau],1}^d]) \neq \operatorname{sgn}(X_{[d^{2\alpha_\tau],1}^d])\}$, using Lemma 4 in Appendix 4.10.3 we have that

$$\begin{aligned} \mathbb{E} \left[\left| \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) - \operatorname{sgn} \left(L_\tau^d \right) \right| \right] &= 2\mathbb{E} \left[\mathbb{1} \left\{ \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) \neq \operatorname{sgn} \left(L_\tau^d \right) \right\} \right] \\ &= 2\mathbb{E} \left[\mathbb{1} \left\{ \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) \neq \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) \right\} \right] \xrightarrow{d \rightarrow \infty} 0 . \end{aligned}$$

The above and the dominated converge theorem show that

$$\mathbb{E} \left[\left| \int_s^t \left\{ \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) - \operatorname{sgn} \left(L_\tau^d \right) \right\} V' \left(X_{[d^{2\alpha_\tau],1}^d} \right) d\tau \right| \right] \xrightarrow{d \rightarrow \infty} 0 .$$

Consider then T_1^d , recalling that the derivatives of V are bounded, we have

$$\begin{aligned} \mathbb{E} \left[\left| T_1^d \right| \right] &\leq \int_s^t C \mathbb{E} \left[\left| \ell d^\alpha \mathbb{E} \left[b_d \left(X_{[d^{2\alpha_\tau],1}^d}, Z_{[d^{2\alpha_\tau],1}^d} \right) \mathbb{1}_{\mathbb{A}_{[d^{2\alpha_\tau],1}^d}} \left| \mathcal{F}_{[d^{2\alpha_\tau],1}^d} \right. \right. \right. \right. \\ &\quad \left. \left. \left. + \frac{h^L(\ell)}{2} \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) \right| \right] d\tau \\ &\leq \int_s^t C \left\{ \mathbb{E} \left[\left| D_{1,\tau}^{(1)} \right| \right] + \mathbb{E} \left[\left| D_{2,\tau}^{(1)} \right| \right] \right\} d\tau , \end{aligned}$$

where we define

$$\begin{aligned} D_{1,\tau}^{(1)} &= \ell d^\alpha \mathbb{E} \left[Z_{[d^{2\alpha_\tau],1}^d} \mathbb{1}_{\mathbb{A}_{[d^{2\alpha_\tau],1}^d}} \left| \mathcal{F}_{[d^{2\alpha_\tau],1}^d} \right. \right] , \\ D_{2,\tau}^{(1)} &= \frac{h^L(\ell)}{2} \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) \\ &\quad - \ell d^\alpha \left(\frac{\sigma_d}{2} \operatorname{sgn} \left(X_{[d^{2\alpha_\tau],1}^d} \right) \mathbb{1}_{|X_{[d^{2\alpha_\tau],1}^d}| \geq \sigma_d^{2m} r/2} + \frac{1}{\sigma_d^{2m-1} r} X_{[d^{2\alpha_\tau],1}^d} \mathbb{1}_{|X_{[d^{2\alpha_\tau],1}^d}| < \sigma_d^{2m} r/2} \right) \\ &\quad \times \mathbb{E} \left[\mathbb{1}_{\mathbb{A}_{[d^{2\alpha_\tau],1}^d}} \left| \mathcal{F}_{[d^{2\alpha_\tau],1}^d} \right. \right] . \end{aligned}$$

4.6. Proof of the Result for the Laplace distribution

Let us start with $D_{1,\tau}^{(1)}$:

$$D_{1,\tau}^{(1)} = \ell d^\alpha \mathbb{E} \left[Z_{[d^{2\alpha\tau},1]}^d \left(1 \wedge \exp \left\{ \sum_{i=1}^d \phi_d \left(X_{[d^{2\alpha\tau},i]}^d, Z_{[d^{2\alpha\tau},i]}^d \right) \right\} \right) \middle| \mathcal{F}_{[d^{2\alpha\tau},1]}^d \right],$$

where ϕ_d is given in (4.21). Then, by independence of the components of $Z_{[d^{2\alpha\tau},1]}^d$, we have

$$\begin{aligned} & \mathbb{E} \left[Z_{[d^{2\alpha\tau},1]}^d \left(1 \wedge \exp \left\{ \sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau},i]}^d, Z_{[d^{2\alpha\tau},i]}^d \right) \right\} \right) \middle| \mathcal{F}_{[d^{2\alpha\tau},1]}^d \right] \\ &= \mathbb{E} \left[Z_{[d^{2\alpha\tau},1]}^d \right] \mathbb{E} \left[1 \wedge \exp \left\{ \sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau},i]}^d, Z_{[d^{2\alpha\tau},i]}^d \right) \right\} \middle| \mathcal{F}_{[d^{2\alpha\tau},1]}^d \right] = 0. \end{aligned}$$

This allows us to write

$$\begin{aligned} \mathbb{E} \left[|D_{1,\tau}^{(1)}| \right] &\leq \ell d^\alpha \mathbb{E} \left[|Z_{[d^{2\alpha\tau},1]}^d| \right. \\ &\quad \left. \left| 1 \wedge \exp \left\{ \sum_{i=1}^d \phi_d \left(X_{[d^{2\alpha\tau},i]}^d, Z_{[d^{2\alpha\tau},i]}^d \right) \right\} - 1 \wedge \exp \left\{ \sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau},i]}^d, Z_{[d^{2\alpha\tau},i]}^d \right) \right\} \right] \right]. \end{aligned}$$

However, $x \mapsto 1 \wedge \exp(x)$ is a 1-Lipschitz function, thus

$$\mathbb{E} \left[|D_{1,\tau}^{(1)}| \right] \leq \ell d^\alpha \mathbb{E} \left[|Z_{[d^{2\alpha\tau},1]}^d| \left| \phi_d \left(X_{[d^{2\alpha\tau},1]}^d, Z_{[d^{2\alpha\tau},1]}^d \right) \right| \right],$$

and $D_{1,\tau}^{(1)} \rightarrow 0$ as $d \rightarrow \infty$ by Lemma 5 in Appendix 4.10.3.

For $D_{2,\tau}^{(1)}$, we observe that

$$-\frac{\sigma_d}{2} \mathbb{1}_{|X_{[d^{2\alpha\tau},1]}^d| < \sigma_d^{2m} r/2} \leq \frac{1}{\sigma^{2m-1} r} X_{[d^{2\alpha\tau},1]}^d \mathbb{1}_{|X_{[d^{2\alpha\tau},1]}^d| < \sigma_d^{2m} r/2} \leq \frac{\sigma_d}{2} \mathbb{1}_{|X_{[d^{2\alpha\tau},1]}^d| < \sigma_d^{2m} r/2}. \quad (4.48)$$

Distinguishing between $X_{[d^{2\alpha\tau},1]}^d < 0$ and $X_{[d^{2\alpha\tau},1]}^d \geq 0$, it follows that

$$\begin{aligned} |D_{2,\tau}^{(1)}| &\leq \left| \operatorname{sgn} \left(X_{[d^{2\alpha\tau},1]}^d \right) \right| \\ &\quad \times \left| \frac{h^L(\ell)}{2} - \ell d^\alpha \left(\frac{\sigma_d}{2} \mathbb{1}_{|X_{[d^{2\alpha\tau},1]}^d| \geq \sigma_d^{2m} r/2} + \frac{\sigma_d}{2} \mathbb{1}_{|X_{[d^{2\alpha\tau},1]}^d| < \sigma_d^{2m} r/2} \right) \mathbb{E} \left[\mathbb{1}_{A_{[d^{2\alpha\tau},1]}^d} \middle| \mathcal{F}_{[d^{2\alpha\tau},1]}^d \right] \right| \\ &\leq \frac{1}{2} \left| h^L(\ell) - \ell^2 \mathbb{E} \left[\mathbb{1}_{A_{[d^{2\alpha\tau},1]}^d} \middle| \mathcal{F}_{[d^{2\alpha\tau},1]}^d \right] \right|, \end{aligned}$$

where we recall that $\sigma_d = \ell d^{-\alpha}$ with $\alpha = 1/3$. Using the triangle inequality we obtain

$$\begin{aligned} 2\mathbb{E} \left[|D_{2,\tau}^{(1)}| \right] &\leq \mathbb{E} \left[\left| h^L(\ell) - \ell^2 \mathbb{E} \left[1 \wedge \exp \left(\sum_{i=1}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \middle| \mathcal{F}_{[d^{2\alpha\tau],1}^d} \right] \right| \right] \\ &\leq \mathbb{E} \left[\left| h^L(\ell) - \ell^2 \mathbb{E} \left[1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \middle| \mathcal{F}_{[d^{2\alpha\tau],1}^d} \right] \right| \right] \\ &\quad + \ell^2 \mathbb{E} \left[\left| 1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) - 1 \wedge \exp \left(\sum_{i=1}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \right| \right], \end{aligned}$$

where we used Jensen's inequality to remove the conditional expectation in the last term. Recalling that $x \mapsto 1 \wedge \exp(x)$ is 1-Lipschitz, we can then bound the second term

$$\begin{aligned} &\ell^2 \mathbb{E} \left[\left| 1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) - 1 \wedge \exp \left(\sum_{i=1}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \right| \right] \\ &\leq \ell^2 \mathbb{E} \left[\left| \phi_d \left(X_{[d^{2\alpha\tau],1}^d, Z_{[d^{2\alpha\tau],1}^d} \right) \right| \right], \tag{4.49} \\ &\leq \ell^2 \mathbb{E} \left[\left| \phi_d \left(X_{[d^{2\alpha\tau],1}^d, Z_{[d^{2\alpha\tau],1}^d} \right) \right|^2 \right]^{1/2}, \end{aligned}$$

where the final expectation converges to zero as $d \rightarrow \infty$ by Proposition 4.23. For the remaining term in $D_{2,\tau}^{(1)}$, since $(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d}]_{2 \leq i \leq n})$ is independent of $\mathcal{F}_{[d^{2\alpha\tau],1}^d}$, we have

$$\begin{aligned} &\ell^2 \mathbb{E} \left[\left| 1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \middle| \mathcal{F}_{[d^{2\alpha\tau],1}^d} \right| \right] \\ &= \ell^2 \mathbb{E} \left[\left| 1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \right| \right], \end{aligned}$$

and, using again the fact that $x \mapsto 1 \wedge \exp(x)$ is 1-Lipschitz, we have

$$\begin{aligned} &\left| h^L(\ell) - \ell^2 \mathbb{E} \left[1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \right] \right| \\ &\leq \left| h^L(\ell) - \ell^2 \mathbb{E} \left[1 \wedge \exp \left(\sum_{i=1}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \right] \right| \\ &\quad + \ell^2 \mathbb{E} \left[\left| \phi_d \left(X_{[d^{2\alpha\tau],1}^d, Z_{[d^{2\alpha\tau],1}^d} \right) \right| \right]. \end{aligned}$$

The last term goes to 0 as shown in (4.49), and, as $h^L(\ell) = \ell^2 a^L(\ell)$, with

$$a^L(\ell) = \lim_{d \rightarrow \infty} \mathbb{E} \left[1 \wedge \exp \left(\sum_{i=1}^d \phi_{d,i} \right) \right],$$

by Theorem 4.4, we obtain

$$\lim_{d \rightarrow \infty} \left| h^L(\ell) - \ell^2 \mathbb{E} \left[1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \right] \right| = 0 ,$$

showing that $D_{2,\tau}^{(1)} \rightarrow 0$ as $d \rightarrow \infty$. To obtain convergence of T_1^d , we observe that for any $\tau \in [s, t]$, $D_{1,\tau}^{(1)}$ and $D_{2,\tau}^{(1)}$ follow the same distributions as $D_{1,s}^{(1)}$ and $D_{2,s}^{(1)}$, since for any $k \in \mathbb{N}$, X_k^d has distribution π_d^L . Therefore, the convergence towards zero of $\mathbb{E}[|D_{1,\tau}^{(1)}|]$ and $\mathbb{E}[|D_{2,\tau}^{(1)}|]$ is uniform for $\tau \in [s, t]$, which gives us $T_1^d \rightarrow 0$ as $d \rightarrow \infty$.

Finally, consider T_2^d . Using analogous arguments to those used for T_1^d , we obtain

$$\begin{aligned} \mathbb{E} [|T_2^d|] &\leq C \int_s^t \frac{\ell^2}{2} \mathbb{E} \left[\left| \mathbb{E} \left[b_d \left(X_{[d^{2\alpha\tau],1}^d, Z_{[d^{2\alpha\tau],1}^d} \right)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha\tau],1}^d}^d} \middle| \mathcal{F}_{[d^{2\alpha\tau],1}^d}^d \right] - a^L(\ell, r) \right| \right] d\tau \\ &\leq C \int_s^t \frac{\ell^2}{2} \left\{ \mathbb{E} [|D_{1,\tau}^{(2)}|] + \mathbb{E} [|D_{2,\tau}^{(2)}|] \mathbb{E} [|D_{3,\tau}^{(2)}|] \right\} d\tau , \end{aligned}$$

where we define

$$\begin{aligned} D_{1,\tau}^{(2)} &= \mathbb{E} \left[\left(Z_{[d^{2\alpha\tau],1}^d} \right)^2 \mathbb{1}_{\mathbb{A}_{[d^{2\alpha\tau],1}^d}^d} \middle| \mathcal{F}_{[d^{2\alpha\tau],1}^d}^d \right] - a^L(\ell, r) , \\ D_{2,\tau}^{(2)} &= \left(\frac{\sigma_d}{2} \operatorname{sgn}(X_{[d^{2\alpha\tau],1}^d}) \mathbb{1}_{\{|X_{[d^{2\alpha\tau],1}^d|} \geq \sigma_d^{2m} r/2\}} \right. \\ &\quad \left. + \frac{1}{\sigma_d^{2m-1} r} X_{[d^{2\alpha\tau],1}^d} \mathbb{1}_{\{|X_{[d^{2\alpha\tau],1}^d|} < \sigma_d^{2m} r/2\}} \right)^2 \times \mathbb{E} \left[\mathbb{1}_{\mathbb{A}_{[d^{2\alpha\tau],1}^d}^d} \middle| \mathcal{F}_{[d^{2\alpha\tau],1}^d}^d \right] , \\ D_{3,\tau}^{(2)} &= 2 \left(\frac{\sigma_d}{2} \operatorname{sgn}(X_{[d^{2\alpha\tau],1}^d}) \mathbb{1}_{\{|X_{[d^{2\alpha\tau],1}^d|} \geq \sigma_d^{2m} r/2\}} \right. \\ &\quad \left. + \frac{1}{\sigma_d^{2m-1} r} X_{[d^{2\alpha\tau],1}^d} \mathbb{1}_{\{|X_{[d^{2\alpha\tau],1}^d|} < \sigma_d^{2m} r/2\}} \right) \mathbb{E} \left[Z_{[d^{2\alpha\tau],1}^d} \mathbb{1}_{\mathbb{A}_{[d^{2\alpha\tau],1}^d}^d} \middle| \mathcal{F}_{[d^{2\alpha\tau],1}^d}^d \right] . \end{aligned}$$

Using (4.48), Cauchy-Schwarz's inequality and the fact that the moments of $Z_{[d^{2\alpha\tau],1}^d}$ are bounded we have

$$\mathbb{E} [|D_{2,\tau}^{(2)}|] \leq \frac{\sigma_d^2}{4} \xrightarrow{d \rightarrow \infty} 0 , \quad \mathbb{E} [|D_{3,\tau}^{(2)}|] \leq C \sigma_d \xrightarrow{d \rightarrow \infty} 0 ,$$

since $\sigma_d = \ell d^{-\alpha}$ with $\alpha = 1/3$. The remaining term is bounded similarly to $D_{2,\tau}^{(1)}$, using the fact that $x \mapsto 1 \wedge \exp(x)$ is 1-Lipschitz, we have

$$\begin{aligned} &\mathbb{E} [|D_{3,\tau}^{(2)}|] \\ &\leq \mathbb{E} \left[\left| \mathbb{E} \left[\left(Z_{[d^{2\alpha\tau],1}^d} \right)^2 \left(1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{[d^{2\alpha\tau],i}^d, Z_{[d^{2\alpha\tau],i}^d} \right) \right) \right) \right] \middle| \mathcal{F}_{[d^{2\alpha\tau],1}^d}^d \right] - a^L(\ell, r) \right| \right] \\ &\quad + \mathbb{E} \left[\left(Z_{[d^{2\alpha\tau],1}^d} \right)^2 \left| \phi_d \left(X_{[d^{2\alpha\tau],1}^d, Z_{[d^{2\alpha\tau],1}^d} \right) \right| \right] . \end{aligned}$$

The second expectation is bounded as (4.49) using Cauchy-Schwarz's inequality and Proposition 4.23. For the first expectation, we use the conditional independence of the components of $Z_{\lfloor d^{2\alpha\tau} \rfloor}^d$ and write

$$\begin{aligned} & \mathbb{E} \left[\left(Z_{\lfloor d^{2\alpha\tau} \rfloor, 1}^d \right)^2 \left(1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{\lfloor d^{2\alpha\tau} \rfloor, i}^d, Z_{\lfloor d^{2\alpha\tau} \rfloor, i}^d \right) \right) \right) \middle| \mathcal{F}_{\lfloor d^{2\alpha\tau} \rfloor, 1}^d \right] \\ &= \mathbb{E} \left[\left(Z_{\lfloor d^{2\alpha\tau} \rfloor, 1}^d \right)^2 \right] \mathbb{E} \left[\left(1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{\lfloor d^{2\alpha\tau} \rfloor, i}^d, Z_{\lfloor d^{2\alpha\tau} \rfloor, i}^d \right) \right) \right) \right] \\ &= \mathbb{E} \left[\left(1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{\lfloor d^{2\alpha\tau} \rfloor, i}^d, Z_{\lfloor d^{2\alpha\tau} \rfloor, i}^d \right) \right) \right) \right]. \end{aligned}$$

It follows that $\mathbb{E}[D_{3,\tau}^{(2)}] \rightarrow 0$ as $d \rightarrow \infty$ since, by Theorem 4.4,

$$\left| \mathbb{E} \left[\left(1 \wedge \exp \left(\sum_{i=2}^d \phi_d \left(X_{\lfloor d^{2\alpha\tau} \rfloor, i}^d, Z_{\lfloor d^{2\alpha\tau} \rfloor, i}^d \right) \right) \right) \right] - a^L(\ell, r) \right| \rightarrow 0.$$

Combining the results for T_i^d , $i = 1, \dots, 5$ we obtain the result. \square

Acknowledgments

F.R.C. and G.O.R. acknowledge support from the EPSRC (grant # EP/R034710/1). G.O.R. acknowledges further support from the EPSRC (grant # EP/R018561/1) and the Alan Turing Institute. A.D. acknowledges support from the Lagrange Mathematics and Computing Research Center.

The authors would like to thank Éric Moulines for helpful discussions. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

4.7 Proof of Theorem 4.1

The proof of Theorem 4.1 follows that of [RR98b, Theorem 1, Theorem 2] and consists of four propositions showing convergence of the log-acceptance probability to a normal random variable and (weak) convergence of the process (4.11) to a Langevin diffusion.

We start by recalling and defining a number of quantities that we will use in the following proofs. Recall that $\sigma_d = \ell/d^\alpha$, that $\lambda_d = \sigma_d^{2m} r/2$ where $m \geq 1/2$ and $r > 0$ are to be chosen according to the different cases in Theorem 4.1. Recalling the expression of the proposal given in (4.9) and using the simplification given in (4.10), we define the proposal with starting point $\mathbf{x}^d \in \mathbb{R}^d$,

$$\mathbf{y}^d(\mathbf{x}^d, \mathbf{z}^d) = \mathbf{x}^d - \frac{\sigma_d^2}{2} \nabla G \left(\text{prox}_G^{\sigma_d^{2m} r/2}(\mathbf{x}^d) \right) + \sigma_d \mathbf{z}^d,$$

4.7. Proof of Theorem 4.1

where $\mathbf{z}^d \in \mathbb{R}^d$. Since $G(\mathbf{x}^d) = \sum_{i=1}^d g(x_i^d)$, the i -th component of the proposal only depends on the i -th components of \mathbf{x}^d and \mathbf{z}^d . Thus, for any $x, z \in \mathbb{R}$ we denote

$$y_d(x, z) = x - \frac{\sigma_d^2}{2} g' \left(\text{prox}_g^{\sigma_d^{2m_r/2}}(x) \right) + \sigma_d z .$$

The proposal for the chain $(X_k^d)_{k \geq 0}$ is then given by $Y_k^d = \mathbf{y}^d(X_k^d, Z_{k+1}^d) = (y_d(X_{k,i}^d, Z_{k+1,i}^d))_{i \in \{1, \dots, d\}}$. Let us define the generator of the discrete process $(X_k^d)_{k \geq 0}$ for all $V \in C_c^\infty(\mathbb{R}^d, \mathbb{R})$, i.e. infinitely differentiable \mathbb{R} -valued multivariate functions with compact support, and any $\mathbf{x}^d \in \mathbb{R}^d$,

$$\begin{aligned} L_d V(\mathbf{x}^d) &= d^{2\alpha} \mathbb{E} \left[\left[V(\mathbf{y}^d(\mathbf{x}^d, Z_1^d)) - V(\mathbf{x}^d) \right] \frac{\pi_d(\mathbf{y}^d(\mathbf{x}^d, Z_1^d)) q_d(\mathbf{y}^d(\mathbf{x}^d, Z_1^d), \mathbf{x}^d)}{\pi_d(\mathbf{x}^d) q_d(\mathbf{x}^d, \mathbf{y}^d(\mathbf{x}^d, Z_1^d))} \wedge 1 \right] \\ &= d^{2\alpha} \mathbb{E} \left[\left[V(\mathbf{y}^d(\mathbf{x}^d, Z_1^d)) - V(\mathbf{x}^d) \right] \prod_{i=1}^d \exp \left(\phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right] , \end{aligned}$$

where the expectation is w.r.t. $Z_1^d = (Z_{1,i}^d)_{i \in \{1, \dots, d\}}$, a d -dimensional standard normal random variable, and where we defined

$$\begin{aligned} \phi_d(x, z) &= \log \frac{\pi(y_d(x, z)) q(y_d(x, z), x)}{\pi(x) q(x, y_d(x, z))} \\ &= g(x) - g(y_d(x, z)) + \log q(y_d(x, z), x) - \log q(x, y_d(x, z)) . \end{aligned} \quad (4.50)$$

In the remainder we will work with one-dimensional functions $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$ applied to the first component of \mathbf{x}^d so that

$$\begin{aligned} L_d V(\mathbf{x}^d) &= d^{2\alpha} \mathbb{E} \left[\left[V(y_d(x_1^d, Z_{1,1}^d)) - V(x_1^d) \right] \frac{\pi_d(\mathbf{y}^d(\mathbf{x}^d, Z_1^d)) q_d(\mathbf{y}^d(\mathbf{x}^d, Z_1^d), \mathbf{x}^d)}{\pi_d(\mathbf{x}^d) q_d(\mathbf{x}^d, \mathbf{y}^d(\mathbf{x}^d, Z_1^d))} \wedge 1 \right] \\ &= d^{2\alpha} \mathbb{E} \left[\left[V(y_d(x_1^d, Z_{1,1}^d)) - V(x_1^d) \right] \prod_{i=1}^d \exp \left(\phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right] . \end{aligned} \quad (4.51)$$

We also define \tilde{L}_d to be a variant of L_d in which the first component of the acceptance ratio is omitted:

$$\tilde{L}_d V(\mathbf{x}^d) = d^{2\alpha} \mathbb{E} \left[\left[V(y_d(x_1^d, Z_{1,1}^d)) - V(x_1^d) \right] \prod_{i=2}^d \exp \left(\phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right] . \quad (4.52)$$

We further define the generator of the Langevin diffusion (4.13)

$$L V(x) = \frac{h(\ell, r)}{2} [V''(x) - g'(x)V'(x)] , \quad (4.53)$$

where $h(\ell, r) = \ell^2 a(\ell, r)$ is the speed of the diffusion and $a(\ell, r) = \lim_{d \rightarrow \infty} a_d(\ell, r)$ is given in Theorem 4.1.

We will make use of the derivatives of g in (4.8) up to order 8, which we denote by g', g'', g''' and $g^{(k)}$ for all $k = 4, \dots, 8$. We recall that $(g^\lambda)'$ is Lipschitz continuous with Lipschitz constant λ^{-1} [RW98, Proposition 12.19] and that $(g^\lambda)'(x) = \lambda^{-1}(\text{prox}_g^\lambda(x) - x)$, hence prox_g^λ is Lipschitz continuous with Lipschitz constant 1.

4.7.1 Auxiliary Results for the Proof of Case (a)

First, we characterize the limit behaviour of the acceptance ratio (4.12).

Proposition 4.10. *Under A18, A19 and A20, if $\alpha = 1/4$, $\beta = 1/8$ and $r > 0$, then*

(i) *the log-acceptance ratio (4.50), when $d \rightarrow \infty$, satisfies the following Taylor expansion*

$$\phi_d(x, z) = d^{-1/2}C_2(x, z) + d^{-3/4}C_3(x, z) + d^{-1}C_4(x, z) + C_5(x, z, \sigma_d),$$

where $C_2(x, z)$ is given in (4.57), C_3 and C_4 are polynomials in z and the derivatives of g , such that $\mathbb{E}[C_3(X_{0,1}^d, Z_{1,1}^d)] = 0$ and $\mathbb{E}[C_2(X_{0,1}^d, Z_{1,1}^d)^2] = -2\mathbb{E}[C_4(X_{0,1}^d, Z_{1,1}^d)]$;

(ii) *there exists sets $F_d \subseteq \mathbb{R}^d$ with $d^{2\alpha}\pi_d(F_d^c) \rightarrow 0$ such that*

$$\lim_{d \rightarrow \infty} \sup_{x^d \in F_d} \mathbb{E} \left[\left| \sum_{i=2}^d \phi_d(x_i^d, Z_{1,i}^d) - d^{-1/2} \sum_{i=2}^d C_2(x_i^d, Z_{1,i}^d) + \frac{\ell^4 K_1(r)^2}{2} \right| \right] = 0, \quad (4.54)$$

where $K_1(r)$ is given in Theorem 4.1-(a).

Proof. Take one component of the log-acceptance ratio

$$\phi_d(x, z) = g(x) - g(y_d(x, z)) + \log q(y_d(x, z), x) - \log q(x, y_d(x, z)),$$

with $y_d(x, z) = x - \sigma_d^2 g'(\text{prox}_g^{\sigma_d^{2m} r/2}(x))/2 + \sigma_d z$. We have that $\phi_d(x, z) = R_1(x, z, \sigma_d) + R_2(x, z, \sigma_d)$, where

$$\begin{aligned} R_1(x, z, \sigma) &= -g \left[x - \frac{\sigma^2}{2} g' \left(\text{prox}_g^{\sigma^{2m} r/2}(x) \right) + \sigma z \right] + g(x), \\ R_2(x, z, \sigma) &= \frac{1}{2} z^2 - \frac{1}{2} \left(z - \frac{\sigma}{2} g' \left(\text{prox}_g^{\sigma^{2m} r/2} \left[x + \sigma z - \frac{\sigma^2}{2} g' \left[\text{prox}_g^{\sigma^{2m} r/2}(x) \right] \right] \right) \right. \\ &\quad \left. - \frac{\sigma}{2} g' \left[\text{prox}_g^{\sigma^{2m} r/2}(x) \right] \right)^2. \end{aligned} \quad (4.55)$$

Following the approach of [RR98b] we approximate $\phi_d(x, z)$ with a Taylor expansion about $\sigma_d \rightarrow 0$.

(i) Using a Taylor expansion of order 5, we obtain

$$\phi_d(x, z) = d^{-1/2}C_2(x, z) + d^{-3/4}C_3(x, z) + d^{-1}C_4(x, z) + C_5(x, z, \sigma_d), \quad (4.56)$$

where

$$C_2(x, z) = \frac{\ell^2}{2} (-rzg''(x)g'(x)), \quad (4.57)$$

$C_3(x, z)$ and $C_4(x, z)$ are given in Section 4.9.1.1 and we use the integral form for the remainder

$$C_5(x, z, \sigma_d) = \int_0^{\sigma_d} \frac{\partial^5}{\partial \sigma^5} R(x, z, \sigma) \Big|_{\sigma=u} \frac{(\sigma_d - u)^4}{4!} du ,$$

with u between 0 and σ_d and the derivatives of R_1 and R_2 given in Appendix 4.9.2. In addition, integrating by parts and using the moments of $Z_{1,1}^d$ we find that $\mathbb{E}[C_2(X_{0,1}^d, Z_{1,1}^d)] = \mathbb{E}[C_3(X_{0,1}^d, Z_{1,1}^d)] = 0$ and

$$2\mathbb{E} [C_4(X_{0,1}^d, Z_{1,1}^d)] + \mathbb{E} [C_2(X_{0,1}^d, Z_{1,1}^d)^2] = 0 .$$

(ii) To construct the sets F_d , consider, for $j = 3, 4$, $F_{d,j} = F_{d,j}^1 \cap F_{d,j}^2$ where we define

$$F_{d,j}^1 = \left\{ \mathbf{x}^d \in \mathbb{R}^d : \left| \sum_{i=2}^d \mathbb{E} [C_j(x_i^d, Z_{1,i}^d) - C_j(X_{0,i}^d, Z_{1,i}^d)] \right| \leq d^{5/8} \right\} ,$$

and

$$F_{d,j}^2 = \left\{ \mathbf{x}^d \in \mathbb{R}^d : \left| \sum_{i=2}^d V_j(x_i^d) - \mathbb{E} [V_j(X_{0,i}^d)] \right| \leq d^{6/5} \right\} ,$$

where $V_j(x) := \text{Var}(C_j(x, Z_{1,1}^d))$. Using Markov's inequality and the fact that C_j, V_j are bounded by polynomials since g and its derivatives are bounded by polynomials, it is easy to show that $d^{1/2}\pi_d((F_{d,j}^1)^c) \rightarrow 0$ and $d^{1/2}\pi_d((F_{d,j}^2)^c) \rightarrow 0$, from which follows $d^{1/2}\pi_d(F_{d,j}^c) \rightarrow 0$ as $d \rightarrow \infty$. To prove L^1 convergence of C_j for $j = 3, 4$, observe that

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{i=2}^d C_j(x_i^d, Z_{1,i}^d) - \mathbb{E} [C_j(X_{0,1}^d, Z_{1,1}^d)] \right)^2 \right] \\ &= \sum_{i=2}^d V_j(x_i^d) + \left(\sum_{i=2}^d \mathbb{E} [C_j(x_i^d, Z_{1,i}^d) - C_j(X_{0,1}^d, Z_{1,1}^d)] \right)^2 , \end{aligned}$$

and that, for $\mathbf{x}^d \in F_{d,j}$, we have

$$\mathbb{E} \left[\left(\sum_{i=2}^d C_j(x_i^d, Z_{1,i}^d) - \mathbb{E} [C_j(X_{0,1}^d, Z_{1,1}^d)] \right)^2 \right] \leq \mathbb{E} [V_j(x_1^d)] (d-1) + d^{6/5} + d^{5/4} .$$

Thus, the third and fourth term in the Taylor expansion (4.56) converge in L^1 to 0 and $-\ell^4 K_1^2(r)/2$ respectively. Now, consider $C_5(x_i^d, Z_{1,i}^d, \sigma_d)$. We can bound $\frac{\partial^5 R}{\partial \sigma^5}(x, z, \sigma)$ with the derivatives of g evaluated at

$$x + \frac{\sigma^2}{2} \text{prox}_g^{\sigma^{2m} r/2}(x) + \sigma z \quad \text{and} \quad \text{prox}_g^{\sigma^{2m} r/2}(x) .$$

Under our assumptions, the derivatives of g are bounded by polynomials M_0 , it follows that there exist polynomials p of the form

$$A \left(1 + \left[\text{prox}_g^{\sigma^{2m}r/2}(x) \right]^N \right) (1 + z^N) (1 + x^N) (1 + \sigma^N) ,$$

for sufficiently large A and sufficiently large even integer N , such that

$$\begin{aligned} \left| g^{(k)} \left[x + \frac{\sigma^2}{2} \text{prox}_g^{\sigma^{2m}r/2}(x) + \sigma_d z \right] \right| &\vee \left| g^{(k)} \left[\text{prox}_g^{\sigma^{2m}r/2}(x) \right] \right| \\ &\leq p(\text{prox}_g^{\sigma^{2m}r/2}(x), x, z, \sigma_d) . \end{aligned}$$

In addition, $|\text{prox}_g^{\sigma^{2m}r/2}(x)| \leq C(1 + |x|)$ for some $C \geq 1$, and we can bound

$$p(\text{prox}_g^{\sigma^{2m}r/2}(x), x, z, \sigma) \leq A (1 + z^N) (1 + x^{2N}) (1 + \sigma^N) .$$

Therefore, we have

$$\begin{aligned} \mathbb{E} \left[\left| C_5(x_i^d, Z_{1,i}^d, \sigma_d) \right| \right] &\leq A \mathbb{E} \left[1 + (Z_{1,i}^d)^N \right] (1 + (x_i^d)^{2N}) \int_0^{\sigma_d} (1 + u^N) \frac{(\sigma_d - u)^4}{4!} du \\ &\leq A \mathbb{E} \left[1 + (Z_{1,i}^d)^N \right] (1 + (x_i^d)^{2N}) d^{-5/2} \\ &\leq A (1 + (x_i^d)^{2N}) d^{-5/2} , \end{aligned}$$

where the last inequality follows since all the moments of Z_1^d are bounded. Let us denote $p(x) = A (1 + x^{2N})$ and

$$F_{d,5} = \left\{ \mathbf{x}^d \in \mathbb{R}^d : \left| d^{-1} \sum_{i=1}^d p(x_i^d) - \mathbb{E} \left[p(X_{0,i}^d) \right] \right| < 1 \right\} .$$

By Chebychev's inequality we have $\pi_d(F_{d,5}^c) \leq \text{Var}(p(X_{0,1}^d))d^{-1}$. Additionally, for all $\mathbf{x}^d \in F_{d,5}$,

$$\begin{aligned} \sum_{i=2}^d \mathbb{E} \left[\left| C_5(x_i^d, Z_{1,i}^d, \sigma_d) \right| \right] &\leq \sum_{i=2}^d d^{-5/2} \left(\mathbb{E} \left[p(X_{0,1}^d) \right] + d^{-1} \right) \\ &\leq d^{-3/2} \left(\mathbb{E} \left[p(X_{0,1}^d) \right] + 1 \right) . \end{aligned}$$

Finally, set $F_d = \cap_{j=3}^5 F_{d,j}$. On F_d the last three terms of (4.56) converge uniformly in L^1 , and (4.54) follows using the triangle inequality. □

Next, we compare the generator L_d and \tilde{L}_d in (4.51) and (4.52) respectively.

Proposition 4.11. *Under A18, A19 and A20, if $\alpha = 1/4$, $\beta = 1/8$ and $r > 0$, there exists sets $S_d \subseteq \mathbb{R}^d$ with $d^{2\alpha}\pi_d(S_d^c) \rightarrow 0$ such that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$*

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in S_d} \left| L_d V(\mathbf{x}^d) - \tilde{L}_d V(\mathbf{x}^d) \right| = 0 ,$$

and

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in S_d} \mathbb{E} \left[\left| \left(\exp \left(\sum_{i=1}^d \phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right) - \left(\exp \left(\sum_{i=2}^d \phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right) \right| \right] = 0 .$$

Proof. The function $x \mapsto \exp(x) \wedge 1$ is Lipschitz continuous with Lipschitz constant 1, hence

$$\left| L_d V(\mathbf{x}^d) - \tilde{L}_d V(\mathbf{x}^d) \right| \leq d^{2\alpha} \mathbb{E} \left[\left| V \left(y_d(x_1^d, Z_{1,1}^d) \right) - V(x_1^d) \right| |R(x_1^d, Z_{1,1}^d, \sigma_d)| \right] ,$$

where $R(x, z, \sigma) = R_1(x, z, \sigma) + R_2(x, z, \sigma)$ as in (4.55). Using a Taylor expansion of order 1 about $\sigma = 0$ with integral remainder:

$$R(x, z, \sigma) = R(x, z, 0) + \frac{\partial R}{\partial \sigma}(x, z, \sigma) \Big|_{\sigma=0} \sigma + \int_0^\sigma \frac{\partial^2 R}{\partial \sigma^2}(x, z, \sigma) \Big|_{\sigma=u} (\sigma - u) du ,$$

we obtain

$$R(x, z, \sigma) = \int_0^\sigma \frac{\partial^2 R}{\partial \sigma^2}(x, z, \sigma) \Big|_{\sigma=u} (\sigma - u) du ,$$

where $\frac{\partial^2 R}{\partial \sigma^2}(x, z, \sigma)$ is bounded by the derivatives of g evaluated at

$$x + \frac{\sigma^2}{2} \text{prox}_g^{\sigma^{2m}r/2}(x) + \sigma z \quad \text{and} \quad \text{prox}_g^{\sigma^{2m}r/2}(x) .$$

Under our assumptions, the derivatives of g are bounded by polynomials M_0 , it follows that there exist polynomials p of the form

$$A \left(1 + \left[\text{prox}_g^{\sigma^{2m}r/2}(x) \right]^N \right) (1 + z^N) (1 + x^N) (1 + \sigma^N) ,$$

for sufficiently large A and sufficiently large even integer N , such that

$$\left| g^{(k)} \left[x + \frac{\sigma^2}{2} \text{prox}_g^{\sigma^{2m}r/2}(x) + \sigma z \right] \right| \vee \left| g^{(k)} \left[\text{prox}_g^{\sigma^{2m}r/2}(x) \right] \right| \leq p(\text{prox}_g^{\sigma^{2m}r/2}(x), x, z, \sigma) .$$

Proceeding as in Proposition 4.10, we can bound

$$p(\text{prox}_g^{\sigma^{2m}r/2}(x), x, z, \sigma) \leq A (1 + z^N) (1 + x^{2N}) (1 + \sigma^N) .$$

Therefore, we have

$$\begin{aligned} \left| R(x_1^d, Z_{1,1}^d, \sigma_d) \right| &\leq A \left(1 + (Z_{1,1}^d)^N \right) \left(1 + (x_1^d)^{2N} \right) \\ &\quad \times \int_0^{\sigma_d} (1 + u^N)(\sigma_d - u) du \leq A \left(1 + (Z_{1,1}^d)^N \right) \left(1 + (x_1^d)^{2N} \right) \frac{\sigma_d^2}{2}. \end{aligned} \quad (4.58)$$

Since $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$, there exists a constant C such that

$$\begin{aligned} \left| V(y_d(x_1^d, Z_{1,1}^d)) - V(x_1^d) \right| &\leq C \left| y_d(x_1^d, Z_{1,1}^d) - x_1^d \right| \\ &\leq C \left(\sigma_d |Z_{1,1}^d| + \frac{\sigma_d^2}{2} \left| g' \left(\text{prox}_g^{\sigma_d^{2m} r/2}(x_1^d) \right) \right| \right). \end{aligned}$$

Recalling that $g'(\text{prox}_g^\lambda(x)) = (g^\lambda)'(x)$ with $(g^\lambda)'$ Lipschitz continuous with Lipschitz constant λ^{-1} , we have

$$\left| g' \left(\text{prox}_g^{\sigma_d^{2m} r/2}(x_1^d) \right) \right| \leq \frac{2}{\sigma_d^{2m} r} (1 + |x_1^d|),$$

and

$$\begin{aligned} \left| V(y_d(x_1^d, Z_{1,1}^d)) - V(x_1^d) \right| &\leq C \left(\sigma_d |Z_{1,1}^d| + \frac{\sigma_d^{2-2m}}{r} (1 + |x_1^d|) \right) \\ &\leq C \sigma_d \left(|Z_{1,1}^d| + \frac{1}{r} (1 + |x_1^d|) \right), \end{aligned} \quad (4.59)$$

since $m = 1/2$. Combining (4.58) and (4.59) we obtain

$$\begin{aligned} d^{2\alpha} \left| V(y_d(x_1^d, Z_{1,1}^d)) - V(x_1^d) \right| \left| R(x_1^d, Z_{1,1}^d, \sigma_d) \right| \\ \leq C \sigma_d \left(1 + (Z_{1,1}^d)^N \right) \left(1 + (x_1^d)^{2N} \right) \left(|Z_{1,1}^d| + \frac{1}{r} (1 + |x_1^d|) \right), \end{aligned}$$

for some $C > 0$.

Set S_d to be the set in which $1 + (x_1^d)^{2N+1} \leq d^{\alpha/2}$, applying Markov's inequality we obtain

$$d^{2\alpha} \pi_d(S_d^c) = d^{2\alpha} \pi_d \left(\left(1 + (x_1^d)^{2N+1} \right)^5 \geq d^{5\alpha/2} \right) \leq d^{-\alpha/2} \mathbb{E} \left[\left(1 + (x_1^d)^{2N+1} \right)^5 \right] \xrightarrow{d \rightarrow \infty} 0.$$

Recalling that $|Z_{1,1}^d|$ and $1 + (Z_{1,1}^d)^N$ are bounded, we have that

$$\sup_{x^d \in S_d} \left| L_d V(x^d) - \tilde{L}_d V(x^d) \right| \leq C d^{\alpha/2} \frac{\ell}{d^\alpha} \xrightarrow{d \rightarrow \infty} 0.$$

The second results follows from (4.58) using the same argument. \square

The following result considers the convergence to the generator of the Langevin diffusion (4.53).

Proposition 4.12. *Under A18, A19 and A20, if $\alpha = 1/4$, $\beta = 1/8$ and $r > 0$, there exists sets $T_d \subseteq \mathbb{R}^d$ with $d^{2\alpha}\pi_d(T_d^c) \rightarrow 0$ as $d \rightarrow \infty$, such that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$*

$$\lim_{d \rightarrow \infty} \sup_{x^d \in T_d} \left| d^{2\alpha} \mathbb{E} \left[V \left(y_d(x_1^d, Z_{1,1}^d) \right) - V(x_1^d) \right] - \frac{\ell^2}{2} (V''(x_1^d) + g'(x_1^d)V'(x_1^d)) \right| = 0.$$

Proof. Take

$$y_d(x_1^d, Z_{1,1}^d) = x_1^d + \frac{\sigma_d^2}{2} g' \left(\text{prox}_{g^{\sigma_d^{2m} r/2}}(x_1^d) \right) + \sigma_d Z_{1,1}^d,$$

and use a Taylor expansion of order 2 of

$$W(x, z, \sigma) = V \left[x + \frac{\sigma^2}{2} g' \left(\text{prox}_{g^{\sigma^{2m} r/2}}(x) \right) + \sigma z \right],$$

about $\sigma = 0$ with integral remainder:

$$\begin{aligned} W(x, z, \sigma) &= W(x, z, 0) + \frac{\partial W}{\partial \sigma}(x, z, \sigma) \Big|_{\sigma=0} \sigma + \frac{1}{2} \frac{\partial^2 W}{\partial \sigma^2}(x, z, \sigma) \Big|_{\sigma=0} \sigma^2 \\ &\quad + \int_0^\sigma \frac{\partial^3 W}{\partial \sigma^3}(x, z, \sigma) \Big|_{\sigma=u} \frac{(\sigma - u)^2}{2} du. \end{aligned}$$

Using the derivatives

$$\begin{aligned} W(x, z, 0) &= V(x), \quad \frac{\partial W}{\partial \sigma}(x, z, \sigma) \Big|_{\sigma=0} = V'(x)z, \\ \frac{\partial^2 W}{\partial \sigma^2}(x, z, \sigma) &= V''(x)z^2 + V'(x)g'(x), \end{aligned}$$

and recalling that $\mathbb{E} [Z_{1,1}^d] = 0$, $\mathbb{E} [(Z_{1,1}^d)^2] = 1$, we have

$$\begin{aligned} \mathbb{E} \left[V \left(y_d(x_1^d, Z_{1,1}^d) \right) - V(x_1^d) \right] &= \frac{\sigma_d^2}{2} \left[V''(x_1^d) + V'(x_1^d)g'(x_1^d) \right] \\ &\quad + \mathbb{E} \left[\int_0^{\sigma_d} \frac{\partial^3 W}{\partial \sigma^3}(x_1^d, Z_{1,1}^d, \sigma) \Big|_{\sigma=u} \frac{(\sigma_d - u)^2}{2} du \right]. \end{aligned}$$

Proceeding as in the previous proposition, we can bound

$$\left| \int_0^{\sigma_d} \frac{\partial^3 W}{\partial \sigma^3}(x_1^d, Z_{1,1}^d, \sigma) \Big|_{\sigma=u} \frac{(\sigma_d - u)^2}{2} du \right| \leq A \left(1 + (Z_{1,1}^d)^N \right) \left(1 + (x_1^d)^{2N} \right) d^{-3\alpha}.$$

Setting T_d to be the set in which $(1 + (x_1^d)^{2N}) \leq d^{\alpha/2}$, the result follows by applying Markov's inequality as in Proposition 4.11. \square

Before proceeding to stating and proving the last auxiliary result, let us denote by $\psi_1 : \mathbb{R} \rightarrow [0, +\infty)$ the characteristic function of the distribution $N(0, \ell^4 K_1^2(r))$, where $K_1^2(r)$ is given in Theorem 4.1–(a),

$$\psi_1(t) = \exp(-t^2 \ell^4 K_1^2(r)/2) ,$$

and by $\psi_1^d(\mathbf{x}^d; t) = \int \exp(itw) \mathcal{Q}_1^d(\mathbf{x}^d; dw)$ the characteristic functions associated with the law

$$\mathcal{Q}_1^d(\mathbf{x}^d; \cdot) = \mathcal{L} \left\{ d^{-1/2} \sum_{i=2}^d C_2(x_i^d, Z_{1,i}^d) \right\} .$$

Proposition 4.13. *Under A18, A19 and A20, if $\alpha = 1/4$, $\beta = 1/8$ and $r > 0$, there exists a sequence of sets $H_d \subseteq \mathbb{R}^d$ such that*

(i) $\lim_{d \rightarrow \infty} d^{2\alpha} \pi_d(H_d^c) = 0$,

(ii) for all $t \in \mathbb{R}$, $\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in H_d} |\psi_1^d(\mathbf{x}^d; t) - \psi_1(t)| = 0$,

(iii) for all bounded continuous function $\chi : \mathbb{R} \rightarrow \mathbb{R}$,

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in H_d} \left| \int_{\mathbb{R}} \mathcal{Q}_1^d(\mathbf{x}^d; du) \chi(u) - \left(2\pi \ell^4 K_1^2(r) \right)^{-1/2} \int_{\mathbb{R}} \chi(u) e^{-u^2/(2\ell^4 K_1^2(r))} du \right| = 0 ,$$

(iv) in particular,

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in H_d} \left| \mathbb{E} \left[1 \wedge \exp \left(d^{-1/2} \sum_{i=2}^d C_2(x_i^d, Z_{1,i}^d) - \frac{\ell^4 K_1^2(r)}{2} \right) \right] - 2\Phi \left(-\frac{\ell^2 K_1(r)}{2} \right) \right| = 0 ,$$

where Φ is the distribution function of the standard normal random variable.

Proof. (i) Define the functions $h_j(x) = [-g''(x)g'(x)]^j$ with $j = 1, \dots, 4$ and let $H_d = H_{d,1} \cap H_{d,2}$ where

$$H_{d,1} = \left\{ \mathbf{x}^d \in \mathbb{R}^d : \left| \frac{1}{d} \sum_{i=2}^d h_j(x_i^d) - \int_{\mathbb{R}} h_j(u) \pi(u) du \right| \leq d^{1/3} \text{ for } j = 1, \dots, 4 \right\} ,$$

$$H_{d,2} = \left\{ \mathbf{x}^d \in \mathbb{R}^d : |h_j(x_i^d)| \leq d^{2/3} \text{ for } i = 1, \dots, d \text{ and } j = 1, \dots, 4 \right\} .$$

Using Chebychev's inequality, the fact that the derivatives of g are bounded by polynomials and that π has finite moments, we have $d^{1/2} \pi_d((H_{d,1})^c) \rightarrow 0$ as $d \rightarrow \infty$. Similarly, by Markov's inequality we have $d^{1/2} \pi_d((H_{d,2})^c) \rightarrow 0$ as $d \rightarrow \infty$.

(ii) We follow [RR98b, Lemma 3(b)] and decompose

$$\begin{aligned}
 |\psi_1^d(\mathbf{x}^d; t) - \psi_1(t)| &\leq \left| \psi_1^d(\mathbf{x}^d; t) - \prod_{i=2}^d \left(1 - \frac{t^2}{2d} v(x_i^d) \right) \right| \\
 &\quad + \left| \prod_{i=2}^d \left(1 - \frac{t^2}{2d} v(x_i^d) \right) - \prod_{i=2}^d \exp \left(-t^2 \frac{v(x_i^d)}{2d} \right) \right| \\
 &\quad + \left| \prod_{i=2}^d \exp \left(-t^2 \frac{v(x_i^d)}{2d} \right) - \exp \left(-t^2 \frac{\ell^4 K_1(r)^2}{2} \right) \right|,
 \end{aligned} \tag{4.60}$$

where $v(x_i^d) = \text{Var}(C_2(x_i^d, Z_{1,i}^d)) = \mathbb{E}[C_2(x_i^d, Z_{1,i}^d)^2]$, where the expectation is taken w.r.t. $Z_{1,i}^d$. For the first term, decompose the characteristic function $\psi_1^d(\mathbf{x}^d; t) = \prod_{i=2}^d \theta_i^d(x_i^d; t)$ as the product of the characteristic functions of $d^{-1/2}W_i$ where we define $W_i = C_2(x_i^d, Z_{1,i}^d)$, using [Dur19, equation (3.3.3)] as in the proof of [Dur19, Theorem 3.4.10] we obtain

$$\begin{aligned}
 \left| \theta_i^d(x_i^d; t) - \left(1 - \frac{t^2}{2d} v(x_i^d) \right) \right| &\leq \mathbb{E} \left[\frac{|t|^3}{d^{3/2}} \frac{|W_i|^3}{3!} \wedge \frac{2|t|^2}{d} \frac{|W_i|^2}{2!} \right] \\
 &\leq \mathbb{E} \left[\frac{|t|^3}{d^{3/2}} \frac{|W_i|^3}{3!}; |W_i| \leq d^{1/2} \varepsilon \right] + \frac{t^2}{d} \mathbb{E} \left[|W_i|^2; |W_i| > d^{1/2} \varepsilon \right] \\
 &\leq \frac{\varepsilon |t|^3}{6d} \mathbb{E} \left[|W_i|^2 \right] + \frac{t^2}{\varepsilon^2 d^2} \mathbb{E} \left[|W_i|^4 \right],
 \end{aligned}$$

for any $\varepsilon > 0$. For sufficiently large d , we have that $t^2 v(x_i^d)/(2d) \leq 1$ for $\mathbf{x} \in H_{d,2}$, and we can use [Dur19, Lemma 3.4.3]

$$\begin{aligned}
 \left| \psi_j^d(\mathbf{x}^d; t) - \prod_{i=2}^d \left(1 - \frac{t^2}{2d} v(x_i^d) \right) \right| &\leq \sum_{i=2}^d \left(\frac{\varepsilon |t|^3}{6d} \mathbb{E} \left[|W_i|^2 \right] + \frac{t^2}{\varepsilon^2 d^2} \mathbb{E} \left[|W_i|^4 \right] \right) \\
 &\leq \frac{\varepsilon \ell^4 |t|^3}{6} (K_1^2(r) + D_1 d^{-1/3}) + \frac{t^2}{\varepsilon^2 d} (\mathbb{E} \left[|W_i|^4 \right] + D_2 \ell^8 d^{-1/4}),
 \end{aligned}$$

where the last inequality follows from the fact that $\mathbf{x}^d \in H_{d,2}$ and D_1, D_2 are positive constants. For any $\delta > 0$ we can chose ε small enough so that the first term in the above is less than $\delta/2$ and we can chose d sufficiently large to make the second term less than $\delta/2$. Thus, for any $\delta > 0$ we can find $\varepsilon > 0$ and $d \in \mathbb{N}$ such that

$$\left| \psi_1^d(\mathbf{x}^d; t) - \prod_{i=2}^d \left(1 - \frac{t^2}{2d} v(x_i^d) \right) \right| < \delta,$$

the uniform convergence then follows. The second term in (4.60) converges to 0 uniformly for all $\mathbf{x}^d \in H_{d,1}$; while for the third term in (4.60) we use again [Dur19, Lemma 3.4.3]

$$\left| \prod_{i=2}^d \exp \left(-t^2 \frac{v(x_i^d)}{2d} \right) - \exp \left(-t^2 \frac{\ell^4 K_1(r)^2}{2} \right) \right| \leq \sum_{i=2}^d \frac{t^4 v(x_i^d)^2}{4d^2},$$

which goes to zero when $d \rightarrow \infty$, for all $\mathbf{x}^d \in H_{d,2}$. The result then follows.

- (iii) Let $\chi : \mathbb{R} \rightarrow \mathbb{R}$ be a bounded and continuous function. Define the sequence $\{\mathbf{x}^d : d \in \mathbb{N}^*\}$, where, for any $d \in \mathbb{N}^*$, $\mathbf{x}^d \in H_d$ satisfies,

$$\begin{aligned} & \sup_{y^d \in H_d} \left| \int_{\mathbb{R}} \mathcal{Q}_1^d(y_i^d; du) \chi(u) - \left(2\pi\ell^4 K_1^2(r)\right)^{-1/2} \int_{\mathbb{R}} \chi(u) e^{-u^2/(2\ell^4 K_1^2(r))} du \right| \\ & \leq \left| \int_{\mathbb{R}} \mathcal{Q}_1^d(x_i^d; du) \chi(u) - \left(2\pi\ell^4 K_1^2(r)\right)^{-1/2} \int_{\mathbb{R}} \chi(u) e^{-u^2/(2\ell^4 K_1^2(r))} du \right| + d^{-1}. \end{aligned}$$

Then, using (ii) and Lévy's continuity theorem (e.g. [Shi96, Theorem 1, page 322]), we obtain

$$\lim_{d \rightarrow \infty} \left| \int_{\mathbb{R}} \mathcal{Q}_1^d(x_i^d; du) \chi(u) - \left(2\pi\ell^4 K_1^2(r)\right)^{-1/2} \int_{\mathbb{R}} \chi(u) e^{-u^2/(2\ell^4 K_1^2(r))} du \right| = 0.$$

Combining this limit with the definition of the sequence $\{\mathbf{x}^d : d \in \mathbb{N}^*\}$, we conclude the proof.

- (iv) This statement follows directly from (iii) and [RGG97b, Proposition 2.4]. □

4.7.2 Auxiliary Results for the Proof of Case (b)

First, we characterize the limit behaviour of the acceptance ratio (4.12). The following result is an extension of [RR98b, Lemma 1].

Proposition 4.14. *Under A18, A19, A20 and A21, if $\alpha = 1/6$, $\beta = 1/6$ and $r > 0$, then*

- (i) *the log-acceptance ratio (4.50), when $d \rightarrow \infty$, satisfies the following Taylor expansion*

$$\begin{aligned} \phi_d(x, z) &= d^{-1/2} C_3(x, z) + d^{-2/3} C_4(x, z) \\ &\quad + d^{-5/6} C_5(x, z) + d^{-1} C_6(x, z) + C_7(x, z, \sigma_d), \end{aligned}$$

where C_3 is given in (4.61), C_4, C_5, C_6 are polynomials in z and the derivatives of g , such that $\mathbb{E}[C_j(X_{0,1}^d, Z_{1,1}^d)] = 0$ for $j = 3, 4, 5$ and $\mathbb{E}[C_3(X_{0,1}^d, Z_{1,1}^d)^2] = -2\mathbb{E}[C_6(X_{0,1}^d, Z_{1,1}^d)]$,

- (ii) *there exists sets $F_d \subseteq \mathbb{R}^d$ with $d^{2\alpha} \pi_d(F_d^c) \rightarrow 0$ such that*

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in F_d} \mathbb{E} \left[\left| \sum_{i=2}^d \phi_d(x_i^d, Z_{1,i}^d) - d^{-1/2} \sum_{i=2}^d C_3(x_i^d, Z_{1,i}^d) + \frac{\ell^6 K_2(r)^2}{2} \right| \right] = 0,$$

where $K_2(r)$ is given in Theorem 4.1-(b).

4.7. Proof of Theorem 4.1

Proof. Take one component of the log-acceptance ratio

$$\phi_d(x, z) = g(x) - g(y_d(x, z)) + \log q(y_d(x, z), x) - \log q(x, y_d(x, z)) ,$$

with $y_d(x, z) = x - \sigma_d^2 g'(\text{prox}_g^{\sigma_d^2 m r / 2}(x)) + \sigma_d z$. Proceeding as in the proof for case (a), we have that $\phi_d(x, z) = R_1(x, z, \sigma_d) + R_2(x, z, \sigma_d)$ where R_1, R_2 are given in (4.55). Following the approach of [RR98b] we approximate $\phi_d(x, z)$ with a Taylor expansion about $\sigma_d = 0$.

(i) Using a Taylor expansion of order 7, we find that

$$\begin{aligned} \phi_d(x, z) &= d^{-1/2} C_3(x, z) + d^{-2/3} C_4(x, z) + d^{-5/6} C_5(x, z) \\ &\quad + d^{-1} C_6(x, z) + C_7(x, z, \sigma_d) , \end{aligned}$$

where

$$C_3(x, z) = \frac{\ell^3}{6} \left(\frac{1}{2} g'''(x) z^3 - \frac{3}{2} g''(x) g'(x) z (1 + 2r) \right) , \quad (4.61)$$

$C_4(x, z), C_5(x, z)$ and $C_6(x, z)$ are given in Section 4.9.1.2 and integral form of the remainder

$$C_7(x, z, \sigma_d) = \int_0^{\sigma_d} \frac{\partial^7}{\partial \sigma^7} R(x, z, \sigma) \Big|_{\sigma=u} \frac{(\sigma_d - u)^6}{6!} du ,$$

with u between 0 and σ_d and the derivatives of R_1 and R_2 are given in Appendix 4.9.2. In addition, integrating by parts and using the moments of the standard normal $Z_{1,1}^d$, we find that $\mathbb{E}[C_3(X_{0,1}^d, Z_{1,1}^d)] = \mathbb{E}[C_4(X_{0,1}^d, Z_{1,1}^d)] = \mathbb{E}[C_5(X_{0,1}^d, Z_{1,1}^d)] = 0$ and

$$\begin{aligned} \mathbb{E} [C_6(X_{0,1}^d, Z_{1,1}^d)] &= -\frac{\ell^6}{16} (r + 2r^2) \mathbb{E} \left[\left(g''(X_{0,1}^d) g'(X_{0,1}^d) \right)^2 \right] \\ &\quad - \frac{\ell^6}{16} \left(\frac{1}{2} + r \right) \mathbb{E} [g''(X_{0,1}^d)^3] - \frac{5\ell^6}{96} \mathbb{E} [g'''(X_{0,1}^d)^2] , \end{aligned}$$

which shows that $\mathbb{E}[C_3(X_{0,1}^d, Z_{1,1}^d)^2 + 2C_6(X_{0,1}^d, Z_{1,1}^d)] = 0$.

(ii) The proof of this result follows using the same steps as case (a) and is analogous to that of [RR98b, Lemma 1]. □

Next, we compare the generator L_d and \tilde{L}_d in (4.51) and (4.52) respectively, extending [RR98b, Theorem 3].

Proposition 4.15. *Under A18, A19, A20 and A21, if $\alpha = 1/6$, $\beta = 1/6$ and $r > 0$, there exists sets $S_d \subseteq \mathbb{R}^d$ with $d^{2\alpha} \pi_d(S_d^c) \rightarrow 0$ such that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$*

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in S_d} \left| L_d V(\mathbf{x}^d) - \tilde{L}_d V(\mathbf{x}^d) \right| = 0 ,$$

and

$$\lim_{d \rightarrow \infty} \sup_{x^d \in S_d} \mathbb{E} \left[\left| \left(\exp \left(\sum_{i=1}^d \phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right) - \left(\exp \left(\sum_{i=2}^d \phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right) \right| \right] = 0.$$

Proof. Proceeding as in Proposition 4.11, for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$, there exists a constant C such that

$$\begin{aligned} \left| V \left(y_d(x_1^d, Z_{1,1}^d) \right) - V(x_1^d) \right| &\leq C \left| y_d(x_1^d, Z_{1,1}^d) - x_1^d \right| \\ &\leq C \left(\sigma_d |Z_{1,1}^d| + \frac{\sigma_d^2}{2} \left| g' \left(\text{prox}_{g^d}^{\sigma_d^{2m} r/2} (x_1^d) \right) \right| \right). \end{aligned}$$

Under A21, g' is Lipschitz continuous and we have, for some $C \geq 1$,

$$\left| g' \left(\text{prox}_{g^d}^{\sigma_d^{2m} r/2} (x_1^d) \right) \right| \leq C \left(1 + \left| \text{prox}_{g^d}^{\sigma_d^{2m} r/2} (x_1^d) \right| \right) \leq C(1 + |x_1^d|),$$

where we used the fact that prox_g^λ is 1-Lipschitz continuous for all $\lambda > 0$. The result then follows similarly to Proposition 4.11 and [RR98b, Theorem 3]. \square

The following result considers the convergence to the generator of the Langevin diffusion (4.53) and is a generalization of [RR98b, Lemma 2].

Proposition 4.16. *Under A18, A19, A20 and A21, if $\alpha = 1/6$, $\beta = 1/6$ and $r > 0$, there exists sets $T_d \subseteq \mathbb{R}^d$ with $d^{2\alpha} \pi_d(T_d^c) \rightarrow 0$ as $d \rightarrow \infty$ such that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$*

$$\lim_{d \rightarrow \infty} \sup_{x^d \in T_d} \left| d^{2\alpha} \mathbb{E} \left[V \left(y_d(x_1^d, Z_{1,1}^d) \right) - V(x_1^d) \right] - \frac{\ell^2}{2} (V''(x_1^d) + g'(x_1^d) V'(x_1^d)) \right| = 0.$$

Proof. The proof is identical to that of Proposition 4.12. \square

Before proceeding to stating and proving the last auxiliary result, let us denote by $\psi_2 : \mathbb{R} \rightarrow [0, +\infty)$ the characteristic function of the distribution $N(0, \ell^6 K_2^2(r))$, where $K_2^2(r)$ is given in Theorem 4.1-(b),

$$\psi_2(t) = \exp(-t^2 \ell^6 K_2^2(r)/2),$$

and by $\psi_2^d(\mathbf{x}^d; t) = \int \exp(itw) \mathcal{Q}_2^d(\mathbf{x}^d; dw)$ the characteristic functions associated with the law

$$\mathcal{Q}_2^d(\mathbf{x}^d; \cdot) = \mathcal{L} \left\{ d^{-1/2} \sum_{i=2}^d C_3(x_i^d, Z_{1,i}^d) \right\}.$$

The following result extends [RR98b, Lemma 3].

Proposition 4.17. *Under A18, A19, A20 and A21, if $\alpha = 1/6$, $\beta = 1/6$ and $r > 0$, there exists a sequence of sets $H_d \subseteq \mathbb{R}^d$ such that*

4.7. Proof of Theorem 4.1

(i) $\lim_{d \rightarrow \infty} d^{2\alpha} \pi_d(H_d^c) = 0$,

(ii) for all $t \in \mathbb{R}$, $\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in H_d} |\psi_2^d(\mathbf{x}^d; t) - \psi_2(t)| = 0$,

(iii) for all bounded continuous function $\chi : \mathbb{R} \rightarrow \mathbb{R}$,

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in H_d} \left| \int_{\mathbb{R}} \mathcal{Q}_2^d(\mathbf{x}^d; du) \chi(u) - \left(2\pi\ell^6 K_2^2(r)\right)^{-1/2} \int_{\mathbb{R}} \chi(u) e^{-u^2/(2\ell^6 K_2^2(r))} du \right| = 0,$$

(iv) in particular,

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in H_d} \left| \mathbb{E} \left[1 \wedge \exp \left(d^{-1/2} \sum_{i=2}^d C_3(x_i^d, Z_{1,i}^d) - \frac{\ell^6 K_2^2(r)}{2} \right) \right] - 2\Phi \left(-\frac{\ell^3 K_2(r)}{2} \right) \right| = 0,$$

where Φ is the distribution function of the standard normal random variable.

Proof. (i) The proof is analogous to that of Proposition 4.13 and follows the same steps of that of [RR98b, Lemma 3(a)].

(ii) The proof is analogous to that of Proposition 4.13 and follows the same steps of that of [RR98b, Lemma 3(b)].

(iii) Following the steps of (iii) in Proposition 4.13 and the Lévy's continuity Theorem (e.g. [Shi96, Theorem 1, page 322]) bring the result.

(iv) This statement follows directly from (iii) and [RGG97b, Proposition 2.4]. \square

4.7.3 Auxiliary Results for the Proof of Case (c)

First, we characterize the limit behaviour of the acceptance ratio (4.12).

Proposition 4.18. *Under A18, A19, A20 and A21 and, if $\alpha = 1/6$, $\beta = m/6$ for $m > 1$ and $r > 0$, then*

(i) *the log-acceptance ratio (4.50), when $d \rightarrow \infty$, satisfies the following Taylor expansion*

$$\begin{aligned} \phi_d(x, z) &= d^{-1/2} C_3(x, z) + d^{-2/3} C_4(x, z) \\ &\quad + d^{-5/6} C_5(x, z) + d^{-1} C_6(x, z) + C_7(x, z, \sigma_d), \end{aligned}$$

where C_3 is given in (4.61), C_4, C_5, C_6 are polynomials in z and the derivatives of g , such that $\mathbb{E}[C_j(X_{0,1}^d, Z_{1,1}^d)] = 0$ for $j = 3, 4, 5$ and $\mathbb{E}[C_3(X_{0,1}^d, Z_{1,1}^d)^2] = -2\mathbb{E}[C_6(X_{0,1}^d, Z_{1,1}^d)]$,

(ii) *there exists sets $F_d \subseteq \mathbb{R}^d$ with $d^{2\alpha} \pi_d(F_d^c) \rightarrow 0$ such that*

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in F_d} \mathbb{E} \left[\left| \sum_{i=2}^d \phi_d(x_i^d, Z_{1,i}^d) - d^{-1/2} \sum_{i=2}^d C_3(x_i^d, Z_{1,i}^d) + \frac{\ell^6 K_2^2(0)}{2} \right| \right] = 0,$$

where K_2 is given in Theorem 4.1-(b).

Proof. Take one component of the log-acceptance ratio

$$\phi_d(x, z) = g(x) - g(y_d(x, z)) + \log q(y_d(x, z), x) - \log q(x, y_d(x, z)) ,$$

with $y_d(x, z) = x - \sigma_d^2 g'(\text{prox}_g^{\sigma_d^{2mr/2}}(x)) + \sigma_d z$. Proceeding as in the proof for case (a), we have that $\phi_d(x, z) = R_1(x, z, \sigma_d) + R_2(x, z, \sigma_d)$ where R_1, R_2 are given in (4.55). Following the approach of [RR98b] we approximate $\phi_d(x, z)$ with a Taylor expansion about $\sigma_d = 0$.

(i) Using a Taylor expansion of order 7, we find that

$$\begin{aligned} \phi_d(x, z) &= d^{-1/2} C_3(x, z) + d^{-2/3} C_4(x, z) + d^{-5/6} C_5(x, z) \\ &\quad + d^{-1} C_6(x, z) + C_7(x, z, \sigma_d) , \end{aligned}$$

where

$$C_3(x, z) = \frac{\ell^3}{6} \left(\frac{1}{2} g'''(x) z^3 - \frac{3}{2} g''(x) g'(x) z \right) ,$$

$C_4(x, z), C_5(x, z)$ and $C_6(x, z)$ are given in Section 4.9.1.3 and integral form of the remainder

$$C_7(x, z, \sigma_d) = \int_0^{\sigma_d} \frac{\partial^7}{\partial \sigma^7} R(x, z, \sigma) \Big|_{\sigma=u} \frac{(\sigma_d - u)^6}{6!} du ,$$

with u between 0 and σ_d and the derivatives of R_1 and R_2 are given in Appendix 4.9.2.

In addition, integrating by parts and using the moments of $Z_{1,1}^d$ we find that $\mathbb{E}[C_3(X_{0,1}^d, Z_{1,1}^d)] = \mathbb{E}[C_4(X_{0,1}^d, Z_{1,1}^d)] = \mathbb{E}[C_5(X_{0,1}^d, Z_{1,1}^d)] = 0$ and

$$\mathbb{E}[C_6(X_{0,1}^d, Z_{1,1}^d)] = \ell^6 \left(-\frac{1}{32} \mathbb{E} [g''(X_{0,1}^d)^3] - \frac{5}{96} \mathbb{E} [g'''(X_{0,1}^d)^2] \right) ,$$

which shows that $\mathbb{E}[C_3(X_{0,1}^d, Z_{1,1}^d)^2 + 2C_6(X_{0,1}^d, Z_{1,1}^d)] = 0$.

(ii) The proof of this result follows using the same steps as case (a) and is analogous to that of [RR98b, Lemma 1]. □

Next, we compare the generator L_d and \tilde{L}_d in (4.51) and (4.52) respectively.

Proposition 4.19. *Under A18, A19, A20 and A21, if $\alpha = 1/6$, $\beta = m/6$ for $m > 1$ and $r > 0$, there exists sets $S_d \subseteq \mathbb{R}^d$ with $d^{2\alpha} \pi_d(S_d^c) \rightarrow 0$ as $d \rightarrow \infty$ such that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$*

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in S_d} \left| L_d V(\mathbf{x}^d) - \tilde{L}_d V(\mathbf{x}^d) \right| = 0 ,$$

and

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in S_d} \mathbb{E} \left[\left| \left(\exp \left(\sum_{i=1}^d \phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right) - \left(\exp \left(\sum_{i=2}^d \phi_d(x_i^d, Z_{1,i}^d) \right) \wedge 1 \right) \right| \right] = 0 .$$

4.7. Proof of Theorem 4.1

Proof. The proof is identical to that of Proposition 4.15. \square

The following result considers the convergence to the generator of the Langevin diffusion (4.53).

Proposition 4.20. *Under A18, A19, A20 and A21, if $\alpha = 1/6$, $\beta = m/6$ for $m > 1$ and $r > 0$, there exists sets $T_d \subseteq \mathbb{R}^d$ with $d^{2\alpha}\pi_d(T_d^c) \rightarrow 0$ as $d \rightarrow \infty$ such that for any $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$*

$$\lim_{d \rightarrow \infty} \sup_{x^d \in T_d} \left| d^{2\alpha} \mathbb{E} \left[V \left(y_d(x_1^d, Z_{1,1}^d) \right) - V(x_1^d) \right] - \frac{\ell^2}{2} (V''(x_1^d) + g'(x_1^d)V'(x_1^d)) \right| = 0 .$$

Proof. The proof is identical to that of Proposition 4.12. \square

Before proceeding to stating and proving the last auxiliary result, let us denote by $\psi_3 : \mathbb{R} \rightarrow [0, +\infty)$ the characteristic function of the distribution $N(0, \ell^6 K_2^2(0))$, where K_2^2 is given in Theorem 4.1-(b),

$$\psi_3(t) = \exp(-t^2 \ell^6 K_2^2(0)/2) ,$$

and by $\psi_3^d(\mathbf{x}^d; t) = \int \exp(itw) \mathcal{Q}_3^d(\mathbf{x}^d; dw)$ the characteristic functions associated with the law

$$\mathcal{Q}_3^d(\mathbf{x}^d; \cdot) = \mathcal{L} \left\{ d^{-1/2} \sum_{i=2}^d C_3(x_i^d, Z_{1,i}^d) \right\} .$$

Proposition 4.21. *Under A18, A19, A20 and A21, if $\alpha = 1/6$, $\beta = m/6$ for $m > 1$ and $r > 0$, there exists a sequence of sets $H_d \subseteq \mathbb{R}^d$ such that*

(i) $\lim_{d \rightarrow \infty} d^{2\alpha} \pi_d(H_d^c) = 0$,

(ii) for all $t \in \mathbb{R}$, $\lim_{d \rightarrow \infty} \sup_{x^d \in H_d} |\psi_3^d(\mathbf{x}^d; t) - \psi_3(t)| = 0$,

(iii) for all bounded continuous function $\chi : \mathbb{R} \rightarrow \mathbb{R}$,

$$\lim_{d \rightarrow \infty} \sup_{x^d \in H_d} \left| \int_{\mathbb{R}} \mathcal{Q}_3^d(\mathbf{x}^d; du) \chi(u) - \left(2\pi \ell^6 K_2^2(0) \right)^{-1/2} \int_{\mathbb{R}} \chi(u) e^{-u^2/(2\ell^6 K_2^2(0))} du \right| = 0 ,$$

(iv) in particular,

$$\lim_{d \rightarrow \infty} \sup_{x^d \in H_d} \left| \mathbb{E} \left[1 \wedge \exp \left(d^{-1/2} \sum_{i=2}^d C_3(x_i^d, Z_{1,i}^d) - \frac{\ell^6 K_2^2(0)}{2} \right) \right] - 2\Phi \left(-\frac{\ell^3 K_2(0)}{2} \right) \right| = 0 ,$$

where Φ is the distribution function of the standard normal random variable.

Proof. (i) The proof is analogous to that of Proposition 4.13 and follows the same steps of that of [RR98b, Lemma 3(a)].

- (ii) The proof is analogous to that of Proposition 4.13 and follows the same steps of that of [RR98b, Lemma 3(b)].
- (iii) Following the steps of (iii) in Proposition 4.13 and the Lévy's continuity Theorem (e.g. [Shi96, Theorem 1, page 322]) bring the result.
- (iv) This statement follows directly from (iii) and [RGG97b, Proposition 2.4]. □

4.7.4 Proof of Theorem 4.1

Proof of Theorem 4.1. (a) The asymptotic acceptance rate follows by combining Propositions 4.10–4.12 with part (iv) of Proposition 4.13 as in the proof of [RR98b, Theorem 1]. To prove the weak convergence of the process it suffices to show that there exists events $F_d^* \in \mathbb{R}^d$ such that for all $t > 0$

$$\lim_{d \rightarrow \infty} \mathbb{P} \left(L_t^d \in F_d^* \text{ for all } 0 \leq s \leq t \right) = 1 ,$$

and

$$\lim_{d \rightarrow \infty} \sup_{\mathbf{x}^d \in F_d^*} \left| L_d V(\mathbf{x}^d) - LV(\mathbf{x}^d) \right| .$$

for all $V \in C_c^\infty(\mathbb{R}, \mathbb{R})$ [EK09, Chapter 4, Corollary 8.7]. We take $F_d^* = F_d \cap S_d \cap T_d \cap H_d$. Then, $d^{2\alpha} \pi_d((F_d^*)^c) \rightarrow 0$ and

$$\lim_{d \rightarrow \infty} \mathbb{P} \left(L_t^d \in F_d^* \text{ for all } 0 \leq s \leq t \right) = 1 ,$$

for all fixed t . Combining Propositions 4.10–4.13 we obtain convergence of the generators.

To obtain the value of $a(\ell, r)$ maximizing the speed, we observe that $K_1^2(r)$ is a function of the ratio $r = c^2/\ell^{2m} = c^2/\ell$ only, we can take $c \propto \ell^{1/2}$ so that $K_1^2(r)$ is constant for given c . Using the same substitution as in [RR98b, Theorem 2] we find that $h(\ell, r)$ is maximized at the unique value of ℓ such that $a(\ell, r) = 0.452$.

- (b) The proof is analogous to that of case (a) replacing Propositions 4.10, 4.11, 4.12 and 4.13 with Propositions 4.14, 4.15, 4.16 and 4.17. To obtain the value of $a(\ell, r)$ maximizing the speed, we observe that $K_2^2(r)$ is a function of the ratio $r = c^2/\ell^{2m} = c^2/\ell^2$ only, we can take $c \propto \ell$ so that $K_2^2(r)$ is constant for given c . Using the same substitution as in [RR98b, Theorem 2] we find that $h(\ell, r)$ is maximized at the unique value of ℓ such that $a(\ell, r) = 0.574$.
- (c) The proof is analogous to that of case (a) replacing Propositions 4.10, 4.11, 4.12 and 4.13 with Propositions 4.18, 4.19, 4.20 and 4.21. To obtain the value of $a(\ell, r)$ maximizing the speed, we observe that $K_2^2(0)$ is constant w.r.t. r , we can use the same substitution as in [RR98b, Theorem 2] we find that $h(\ell, r)$ is maximized at the unique value of ℓ such that $a(\ell, r) = 0.574$. □

Case	Figure	Algorithm	α	β	m	r
(a)	4.5	Proximal MALA	1/4	1/8	1/2	1
(b)	4.6	P-MALA	1/6	1/6	1	1
(c)	4.7	Proximal MALA	1/6	1/2	3	2
—	4.8	MALA	1/6	1/6	1	≈ 0

Table 4.1: Algorithm setting for the simulation study on the Gaussian target.

4.8 Numerical Experiments

4.8.1 Differentiable Targets

We collect here a number of numerical experiments confirming the results in Section 4.3.2. To do so, we consider the Gaussian distribution in Example 4.2 and four algorithmic settings summarized in Table 4.1 which correspond to the three cases identified in Theorem 4.1 and MALA.

The first plot in Figure 4.5–4.8 show that for values of α different from those identified in Theorem 4.1 the acceptance ratio $a_d(\ell, r)$ becomes degenerate as d increases. For the values of α identified in Theorem 4.1 we analyze the influence of ℓ on the acceptance $a_d(\ell, r)$ (second plot), obtaining for $d \rightarrow +\infty$ the expression given by Theorem 4.1–(a) for Figure 4.5, the expression in Theorem 4.1–(b) for Figures 4.6 and 4.8 and that in Theorem 4.1–(c) for Figure 4.7.

Finally, we consider the relationship between acceptance ratio $a_d(\ell, r)$ and the speed of the diffusion $h(\ell, r)$ approximated by the expected squared jumping distance (see, e.g. [GRG96])

$$\text{ESJD}_d := d^{2\alpha} \mathbb{E} \left[(X_0^d - X_1^d)^2 \right].$$

Looking at the last plot in Figure 4.5–4.8 we find that, even for relatively small values of d , the shape of the plot of ESJD_d as a function of the acceptance $a_d(\ell, r)$ is similar to that of the theoretical limit. This suggests that tuning the acceptance ratio to be approximately 0.452 when $\alpha = 1/4, \beta = 1/8$ and approximately 0.574 when $\alpha = 1/6, \beta = m/6$ with $m \geq 1$ should generally guarantee high efficiency.

4.8.2 Laplace Target

We collect here a number of numerical experiments confirming the results for the Laplace distribution in Section 4.3.2. Similarly to Section 4.8.1 we consider three algorithmic settings, summarized in Table 4.2.

The first plot in Figures 4.9–4.11 shows that for $\alpha \neq 1/3$ the acceptance ratio $a_d(\ell, r)$ becomes degenerate as d increases; while the second plot shows that $(a_d(\ell, r))_{d \in \mathbb{N}^*}$ and $(\text{ESJD}_d)_{d \in \mathbb{N}^*}$ converge to $a^L(\ell)$ and $h^L(\ell)$ given in Theorems 4.4 and 4.7, respectively. In the case $m = 3, r = 2$ in Figure 4.11, we find that the behaviour for low values of




Figure 4.5: **Case (a)**: Proximal MALA with Gaussian target and $m = 1/2, r = 1$. Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); $ESJD_d$ as a function of the acceptance rate $a_d(\ell, r)$ (third).

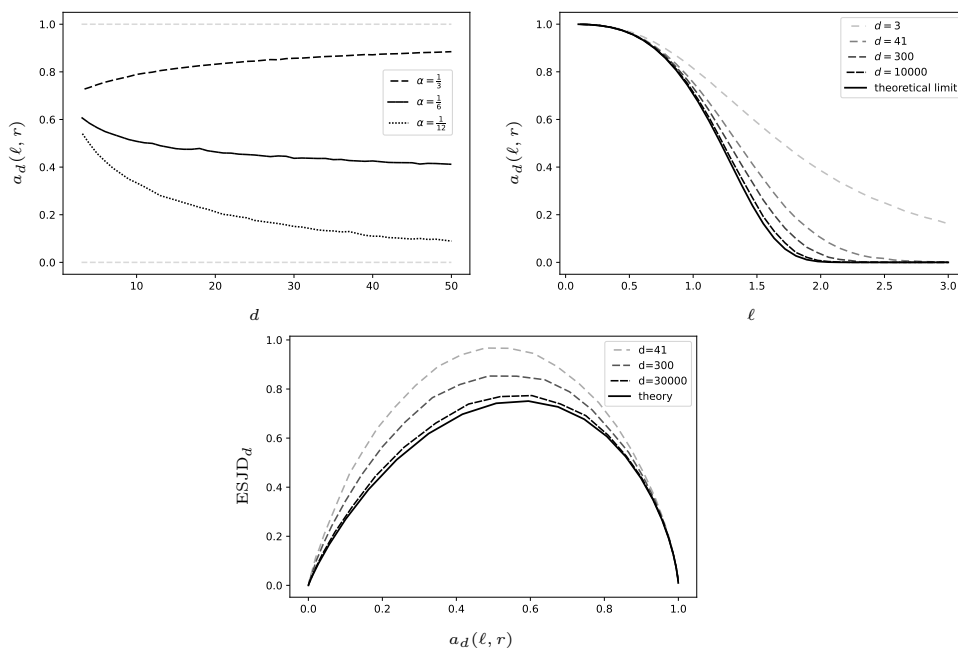


Figure 4.6: **Case (b)**: Proximal MALA with Gaussian target and $m = 1, r = 1$ (P-MALA). Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); ESJD_d as a function of the acceptance rate $\alpha_d(\ell, r)$ (third).




Figure 4.7: **Case (c)**: Proximal MALA with Gaussian target and $m = 3, r = 2$. Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); ESJD_d as a function of the acceptance rate $a_d(\ell, r)$ (third).

4.8. Numerical Experiments

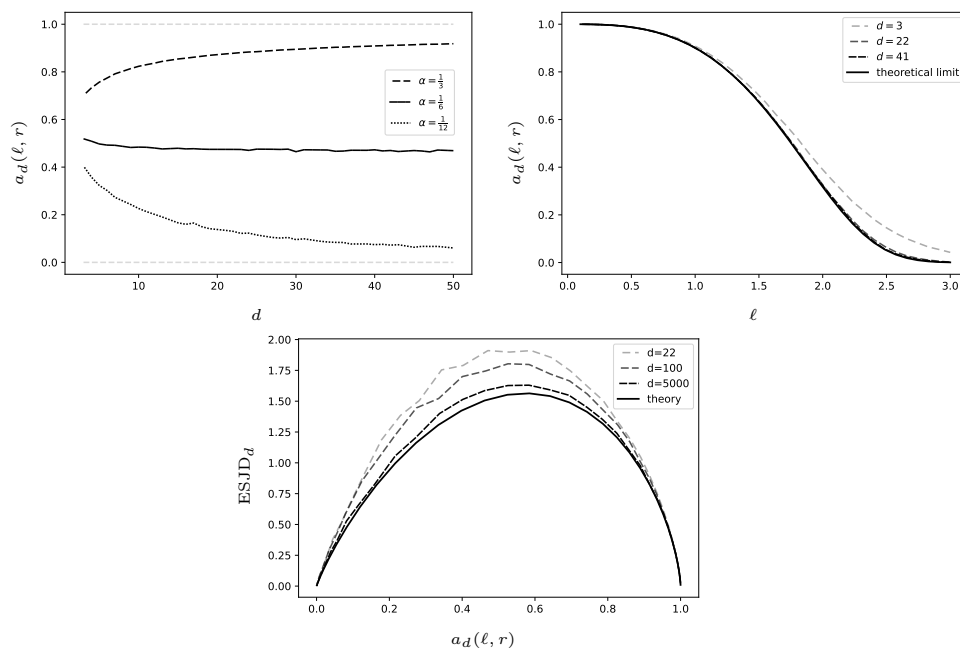


Figure 4.8: Proximal MALA with Gaussian target and $m = 1, r \rightarrow 0$ (MALA). Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); ESJD_d as a function of the acceptance rate $\alpha_d(\ell, r)$ (third).

Figure	Algorithm	α	β	m	r
4.9	sG-MALA	1/3	1/3	1	0
4.10	P-MALA	1/3	1/3	1	1
4.11	Proximal MALA	1/3	1	3	2

Table 4.2: Algorithm setting for the simulation study on the Laplace target.




Figure 4.9: Proximal MALA with Laplace target and $m = 1, r = 0$ (sG-MALA). Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); ESJD_d as a function of the acceptance rate $a_d(\ell, r)$ (third).

d significantly differs from the limiting one. For values of d lower than 130 the ESJD_d achieves its maximum at a value of the acceptance $a_d(\ell, r)$ different from that predicted by Theorem 4.7. In practice, this might mean that for low dimensional settings the recommended choice of $a(\ell, c) = 0.360$ is far from optimal. Similar behaviours for small d have also been observed in the case of RWM and MALA (e.g., [SR+09, Section 2.1]).

4.8.3 Mix of a Laplace and differentiable target

We collect here a the rest of the numerical experiments illustrating the scaling the results for the density defined in (4.19), discussed in Section 4.4.1. Similarly to Section 4.8.2 we consider three algorithmic settings, summarized in Table 4.3.

The first plot in Figures 4.13–4.12 shows that for $\alpha \neq 1/3$ the acceptance ratio $a_d(\ell, r)$ becomes degenerate as d increases; while the second plot shows that $(a_d(\ell, r))_{d \in \mathbb{N}^*}$ and




Figure 4.10: Proximal MALA with Laplace target and $m = 1, r = 1$ (P-MALA). Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); ESJD_d as a function of the acceptance rate $\alpha_d(\ell, r)$ (third).

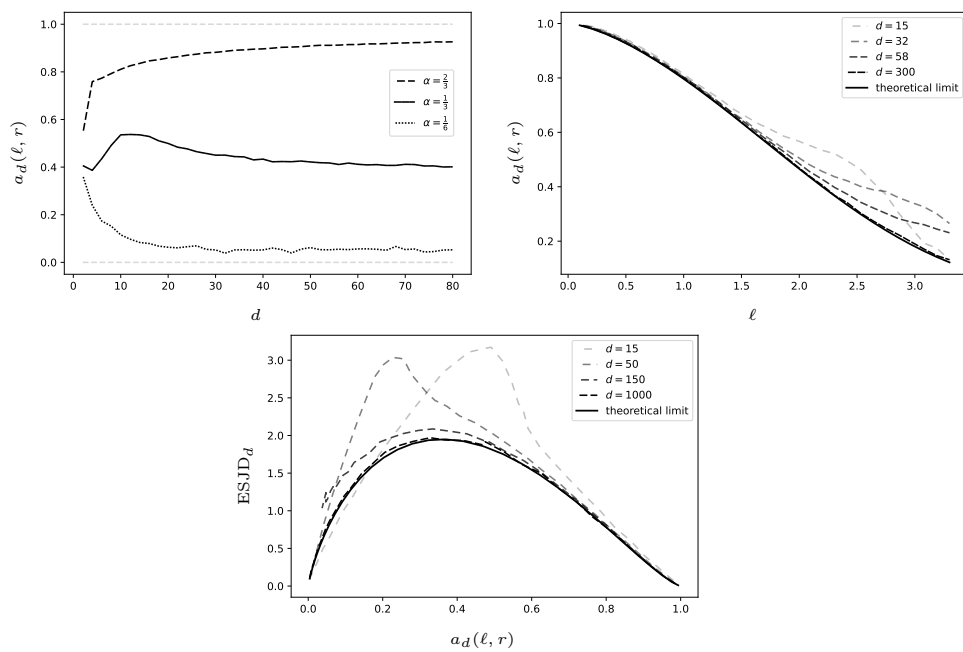


Figure 4.11: Proximal MALA with Laplace target and $m = 3, r = 2$. Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); ESJD_d as a function of the acceptance rate $\alpha_d(\ell, r)$ (third).

Figure	Algorithm	α	β	m	r
4.12	sG-MALA	1/3	1/3	1	0
4.13	P-MALA	1/3	1/3	1	1
4.4	Proximal MALA	1/3	1	3	2

Table 4.3: Algorithm setting for the simulation study on the mixed Laplace-normal target.

$(\text{ESJD}_d)_{d \in \mathbb{N}^*}$ converge as $d \rightarrow \infty$.

It is interesting to note that the convergence of these plots happens faster with respect to the dimension d , compared to the Laplace distribution. It is also surprising that out of the three settings tested for the parameters r and m , the one that seems to converge faster, and that looks most stable, is $m = 3$ and $r = 2$, which was never the case for the other densities. This may suggest that for such distributions, a general Proximal MALA algorithm has better properties of convergence than sG-MALA and P-MALA.

4.9 Taylor Expansions for the Results on Differentiable Targets

4.9.1 Coefficients of the Taylor Expansion

We collect here the coefficients of the Taylor expansions in Propositions 4.10, 4.14 and 4.18.

4.9.1.1 Case (a)

If $\alpha = 1/4$, $\beta = 1/8$ and $r > 0$, then the log-acceptance ratio (4.50) satisfies

$$\phi_d(x, z) = d^{-1/2}C_2(x, z) + d^{-3/4}C_3(x, z) + d^{-1}C_4(x, z) + C_5(x, z, \sigma_d) ,$$

where

$$C_2(x, z) = \frac{\ell^2}{2} (-rzg''(x)g'(x)) ,$$




Figure 4.12: Proximal MALA with mixed Laplace-normal target and $m = 1, r = 0$ (sG-MALA). Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); $ESJD_d$ as a function of the acceptance rate $a_d(\ell, r)$ (third).

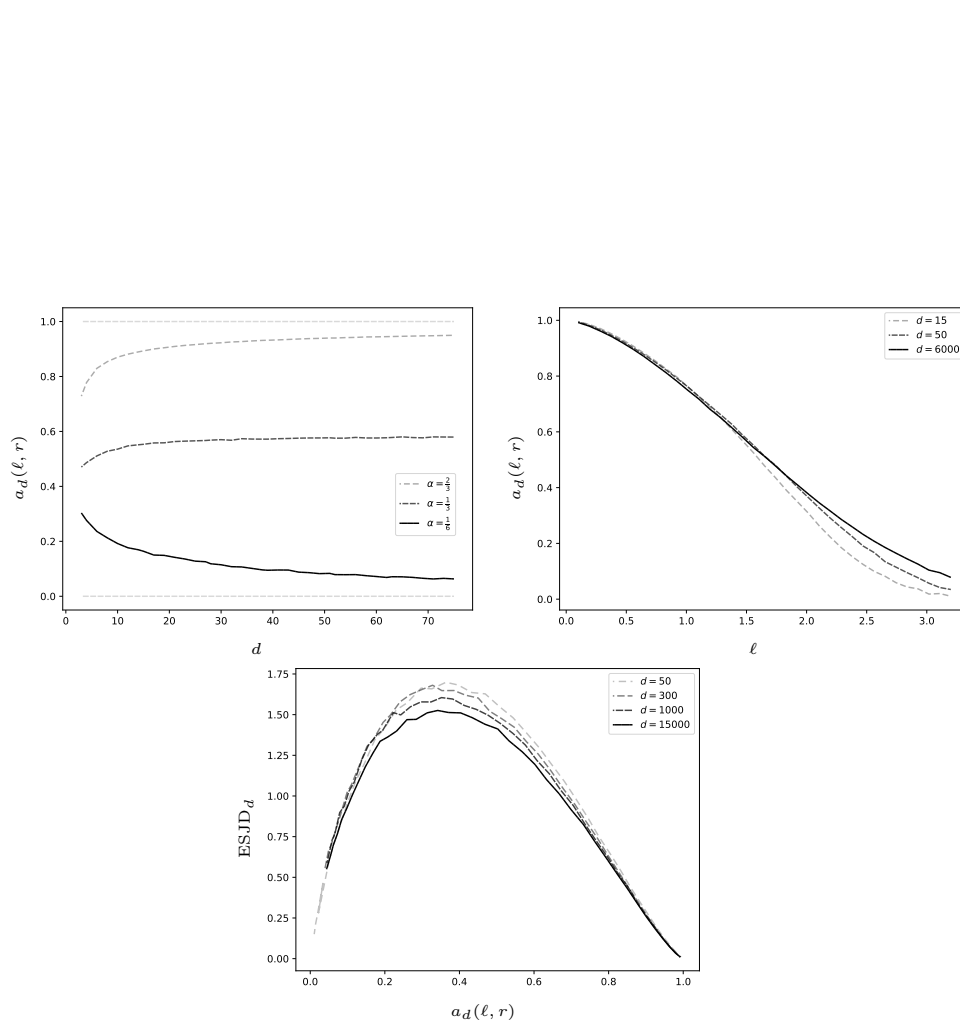


Figure 4.13: Proximal MALA with mixed Laplace-normal target and $m = 1, r = 1$ (P-MALA). Average acceptance rate for different choices of α (first); acceptance rate as a function of ℓ for increasing dimension d (second); ESJD_d as a function of the acceptance rate $a_d(\ell, r)$ (third).

and

$$\begin{aligned}
 C_3(x, z) &= \frac{\ell^3}{6} \left(\frac{z^3}{2} g'''(x) - \frac{3}{2} z^2 r g'(x) g'''(x) - \frac{3}{2} r z^2 [g''(x)]^2 + \frac{3}{4} z r^2 g'''(x) [g'(x)]^2 \right. \\
 &\quad \left. - \frac{3}{2} z g'(x) g''(x) + \frac{3}{2} r [g'(x)]^2 g''(x) + 3r^2 z [g'(x)]^2 g''(x) \right), \\
 C_4(x, z) &= \frac{\ell^4}{24} \left\{ g^{(4)}(x) \left(z^4 - \frac{zr^3}{2} [g'(x)]^3 - 3z^3 r g'(x) + \frac{3}{2} z^2 r^2 [g'(x)]^2 \right) \right. \\
 &\quad + g'''(x) \left(-6z^2 g'(x) - 9rz^3 g''(x) + 9z^2 r^2 g'(x) g''(x) + 9rz [g'(x)]^2 \right. \\
 &\quad \left. - \frac{9}{2} z r^3 [g'(x)]^2 g''(x) - \frac{3}{2} r^2 [g'(x)]^3 \right) \\
 &\quad \left. + 12rz g'(x) [g''(x)]^2 + 3g''(x) [g'(x)]^2 - 3z^2 [g''(x)]^2 + 3z^2 r^2 [g''(x)]^3 \right. \\
 &\quad \left. - 6r^2 [g'(x) g''(x)]^2 - 3zr^3 g'(x) [g''(x)]^3 \right\},
 \end{aligned}$$

and we use the integral form for the remainder

$$C_5(x, z, \sigma_d) = \int_0^{\sigma_d} \frac{\partial^5}{\partial \sigma^5} R(x, z, \sigma) \Big|_{\sigma=u} \frac{(\sigma_d - u)^4}{4!} du,$$

with u between 0 and σ_d and the derivatives of R_1 and R_2 given in Appendix 4.9.2.

4.9.1.2 Case (b)

If $\alpha = 1/6$, $\beta = 1/6$ and $r > 0$, the the log-acceptance ratio (4.50) satisfies

$$\begin{aligned}
 \phi_d(x, z) &= d^{-1/2} C_3(x, z) + d^{-2/3} C_4(x, z) \\
 &\quad + d^{-5/6} C_5(x, z) + d^{-1} C_6(x, z) + C_7(x, z, \sigma_d),
 \end{aligned}$$

where

$$C_3(x, z) = \frac{\ell^3}{6} \left(\frac{1}{2} g'''(x) z^3 - \frac{3}{2} g''(x) g'(x) z (1 + 2r) \right),$$

and

$$\begin{aligned}
 C_4(x, z) &= \frac{\ell^4}{24} \left(z^4 g^{(4)}(x) - 6z^2 g'''(x)g'(x)(1+r) \right. \\
 &\quad \left. - 3z^2 [g''(x)]^2 (1+2r) + 3g''(x) [g'(x)]^2 (1+2r) \right), \\
 C_5(x, z) &= \frac{\ell^5}{120} \left(\frac{3}{2} z^5 g^{(5)}(x) - 15z^3 g^{(4)}(x)g'(x)(1+r) + 15z [g'(x)]^2 g'''(x) \left(\frac{3}{2} + 3r + r^2 \right) \right. \\
 &\quad \left. + 15z(1+4r+2r^2)g'(x) [g''(x)]^2 - 15z^3 g''(x)g'''(x)(1+3r) \right), \\
 C_6(x, z) &= \frac{\ell^6}{720} \left(2z^6 g^{(6)}(x) - 30(1+r)z^4 g'(x)g^{(5)}(x) + 45(2+4r+r^2)z^2 [g'(x)]^2 g^{(4)}(x) \right. \\
 &\quad \left. + 90(r+r^2)z^2 [g''(x)]^3 - 15(2+6r+3r^2)g'''(x) [g'(x)]^3 \right. \\
 &\quad \left. - 30(1+4r)z^4 g''(x)g^{(4)}(x) + 45(3+16r+6r^2)z^2 g'(x)g''(x)g'''(x) \right. \\
 &\quad \left. - \frac{45}{2}(1+4r)z^4 [g'''(x)]^2 - \frac{45}{2}(1+8r+8r^2) [g'(x)g''(x)]^2 \right),
 \end{aligned}$$

and integral form of the remainder

$$C_7(x, z, \sigma_d) = \int_0^{\sigma_d} \frac{\partial^7}{\partial \sigma^7} R(x, z, \sigma) \Big|_{\sigma=u} \frac{(\sigma_d - u)^6}{6!} du,$$

with u between 0 and σ_d and the derivatives of R_1 and R_2 are given in Appendix 4.9.2.

4.9.1.3 Case (c)

If $\alpha = 1/6$, $\beta = m/6$ for $m > 1$ and $r > 0$, then the log-acceptance ratio (4.50) satisfies

$$\begin{aligned}
 \phi_d(x, z) &= d^{-1/2} C_3(x, z) + d^{-2/3} C_4(x, z) \\
 &\quad + d^{-5/6} C_5(x, z) + d^{-1} C_6(x, z) + C_7(x, z, \sigma_d),
 \end{aligned}$$

where

$$C_3(x, z) = \frac{\ell^3}{6} \left(\frac{1}{2} g'''(x) z^3 - \frac{3}{2} g''(x) g'(x) z \right)$$

and

$$\begin{aligned}
 C_4(x, z) &= \frac{\ell^4}{24} \begin{cases} z^4 g^{(4)}(x) - 6z^2 g'''(x)g'(x) - 3z^2 [g''(x)]^2 + 3g''(x) [g'(x)]^2 \\ \quad - 12zr g'(x)g''(x) & \text{if } m = 3/2 \\ z^4 g^{(4)}(x) - 6z^2 g'''(x)g'(x) - 3z^2 [g''(x)]^2 + 3g''(x) [g'(x)]^2 \\ \quad \text{otherwise} \end{cases}, \\
 C_5(x, z) &= \frac{\ell^5}{120} \begin{cases} \frac{3}{2}z^5 g^{(5)}(x) - 15z^3 g^{(4)}(x)g'(x) + \frac{45}{2}z [g'(x)]^2 g'''(x) + 15zg'(x) [g''(x)]^2 \\ \quad - 15z^3 g''(x)g'''(x) - 30z^2 r [g''(x)]^2 + 30r [g'(x)]^2 g''(x) \\ \quad - 30rz^2 g'(x)g'''(x) & \text{if } m = 3/2 \\ \frac{3}{2}z^5 g^{(5)}(x) - 15z^3 g^{(4)}(x)g'(x) + \frac{45}{2}z [g'(x)]^2 g'''(x) + 15zg'(x) [g''(x)]^2 \\ \quad - 15z^3 g''(x)g'''(x) - 60zr g'(x)g''(x) & \text{if } m = 2 \\ \frac{3}{2}z^5 g^{(5)}(x) - 15z^3 g^{(4)}(x)g'(x) + \frac{45}{2}z [g'(x)]^2 g'''(x) \\ \quad + 15zg'(x) [g''(x)]^2 - 15z^3 g''(x)g'''(x) & \text{otherwise} \end{cases}, \\
 C_6(x, z) &= \frac{\ell^6}{720} \begin{cases} 2z^6 g^{(6)}(x) - 30z^4 g'(x)g^{(5)}(x) + 90z^2 [g'(x)]^2 g^{(4)}(x) - 30g'''(x) [g'(x)]^3 \\ \quad - \frac{45}{2} [g'(x)g''(x)]^2 - 30z^4 g''(x)g^{(4)}(x) + 135z^2 g'(x)g''(x)g'''(x) \\ \quad - \frac{45}{2} z^4 [g'''(x)]^2 - 90z^3 r g'(x)g^{(4)}(x) + 540rz g'(x) [g''(x)]^2 \\ \quad + 180rz g'(x)g''(x) + 270rz g'''(x) [g'(x)]^2 - 45zr g''(x)g'''(x) \\ \quad + 90rz^3 g'''(x) & \text{if } m = 3/2 \\ 2z^6 g^{(6)}(x) - 30z^4 g'(x)g^{(5)}(x) + 90z^2 [g'(x)]^2 g^{(4)}(x) - 30g'''(x) [g'(x)]^3 \\ \quad - \frac{45}{2} [g'(x)g''(x)]^2 - 30z^4 g''(x)g^{(4)}(x) + 135z^2 g'(x)g''(x)g'''(x) \\ \quad - \frac{45}{2} z^4 [g'''(x)]^2 - 180z^2 r g'(x)g'''(x) - 180z^2 r [g''(x)]^2 \\ \quad + 180r g''(x) [g'(x)]^2 & \text{if } m = 2 \\ 2z^6 g^{(6)}(x) - 30z^4 g'(x)g^{(5)}(x) + 90z^2 [g'(x)]^2 g^{(4)}(x) - 30g'''(x) [g'(x)]^3 \\ \quad - \frac{45}{2} [g'(x)g''(x)]^2 - 30z^4 g''(x)g^{(4)}(x) + 135z^2 g'(x)g''(x)g'''(x) \\ \quad - \frac{45}{2} z^4 [g'''(x)]^2 - 360zr g'(x)g''(x) & \text{if } m = 5/2 \\ 2z^6 g^{(6)}(x) - 30z^4 g'(x)g^{(5)}(x) + 90z^2 [g'(x)]^2 g^{(4)}(x) - 30g'''(x) [g'(x)]^3 \\ \quad - \frac{45}{2} [g'(x)g''(x)]^2 - 30z^4 g''(x)g^{(4)}(x) + 135z^2 g'(x)g''(x)g'''(x) \\ \quad - \frac{45}{2} z^4 [g'''(x)]^2 & \text{otherwise} \end{cases},
 \end{aligned}$$

and integral form of the remainder

$$C_7(x, z, \sigma_d) = \int_0^{\sigma_d} \frac{\partial^7}{\partial \sigma^7} R(x, z, \sigma) \Big|_{\sigma=u} \frac{(\sigma_d - u)^6}{6!} du,$$

with u between 0 and σ_d and the derivatives of R_1 and R_2 are given in Appendix 4.9.2.

4.9.2 Taylor Expansions of the Log-acceptance Ratio

4.9.2.1 R_1

Recall that $R_1(x, z, \sigma) = -g \left[x - \frac{\sigma^2}{2} g' \left(\text{prox}_g^{\sigma^{2m_r/2}}(x) \right) + \sigma z \right] + g(x)$. We compute the derivatives of R_1 w.r.t. σ evaluated at 0:

$$\begin{aligned} R_1(x, z, 0) &= 0, \\ \frac{\partial R_1}{\partial \sigma}(x, z, \sigma)|_{\sigma=0} &= -g'(x)z, \\ \frac{\partial^2 R_1}{\partial \sigma^2}(x, z, \sigma)|_{\sigma=0} &= -z^2 g''(x) + [g'(x)]^2, \\ \frac{\partial^3 R_1}{\partial \sigma^3}(x, z, \sigma)|_{\sigma=0} &= -z^3 g'''(x) + 3g'(x)g''(x) \left[z + \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right], \\ \frac{\partial^4 R_1}{\partial \sigma^4}(x, z, \sigma)|_{\sigma=0} &= -z^4 g^{(4)}(x) + 6z^2 g'''(x)g'(x) - 3g''(x) [g'(x)]^2 \\ &\quad + 12z [g''(x)]^2 \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \\ &\quad + 6g'(x) \\ &\quad \times \left[g''(x) \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} + \left(\frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right)^2 g'''(x) \right]. \end{aligned}$$

In addition, for $m > 1/2$ we will also use

$$\begin{aligned}
 \frac{\partial^5 R_1}{\partial \sigma^5}(x, z, \sigma)|_{\sigma=0} &= -z^5 g^{(5)}(x) + 10z^3 g^{(4)}(x)g'(x) - 15zg'''(x) [g'(x)]^2 \\
 &\quad + 30z [g''(x)]^2 \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad + 10g'(x)g''(x) \frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0}, \\
 \frac{\partial^6 R_1}{\partial \sigma^6}(x, z, \sigma)|_{\sigma=0} &= -z^6 g^{(6)}(x) + 15z^4 g^{(5)}(x)g'(x) - 45z^2 g^{(4)}(x) [g'(x)]^2 + 15g'''(x) [g'(x)]^3 \\
 &\quad - 90g'(x) [g''(x)]^2 \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad - 60zg''(x) \frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad + 90g''(x)g'''(x) \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad + 15g'(x) \\
 &\quad \times \left(g''(x) \frac{\partial^4}{\partial \sigma^4} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + 3 \left[\frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right]^2 g'''(x) \right).
 \end{aligned}$$

4.9.2.2 R_2

Recall that

$$\begin{aligned}
 R_2(x, z, \sigma) &= \frac{1}{2}z^2 \\
 &\quad - \frac{1}{2} \left(z - \frac{\sigma}{2}g' \left(\text{prox}_g^{\sigma^{2m}r/2} \left[x + \sigma z - \frac{\sigma^2}{2}g' \left[\text{prox}_g^{\sigma^{2m}r/2}(x) \right] \right] \right) - \frac{\sigma}{2}g' \left[\text{prox}_g^{\sigma^{2m}r/2}(x) \right] \right)^2.
 \end{aligned}$$

We compute the derivatives of R_2 w.r.t. σ evaluated at 0:

$$\begin{aligned}
 R_2(x, z, 0) &= 0, \\
 \frac{\partial R_2}{\partial \sigma}(x, z, \sigma)|_{\sigma=0} &= zg'(x), \\
 \frac{\partial^2 R_2}{\partial \sigma^2}(x, z, \sigma)|_{\sigma=0} &= -[g'(x)]^2 + zg''(x) \left[z + 2 \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right], \\
 \frac{\partial^3 R_2}{\partial \sigma^3}(x, z, \sigma)|_{\sigma=0} &= -3g'(x)g''(x) \left[z + 2 \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right] \\
 &\quad + \frac{3}{2}z \left(\left[z + \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right]^2 g'''(x) \right. \\
 &\quad + \left. \left[-g'(x) + 2z \frac{\partial^2}{\partial \sigma \partial x} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right] g''(x) \right. \\
 &\quad + \left. \left[\frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right]^2 g'''(x) + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} g''(x) \right), \\
 \frac{\partial^4 R_2}{\partial \sigma^4}(x, z, \sigma)|_{\sigma=0} &= -3[g''(x)]^2 \left[z + 2 \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right]^2 \\
 &\quad - 6g'(x) \left(\left[z + \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right]^2 g'''(x) \right. \\
 &\quad + \left. \left[-g'(x) + 2z \frac{\partial^2}{\partial \sigma \partial x} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right] g''(x) \right. \\
 &\quad + \left. \left[\frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right]^2 g'''(x) + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} g''(x) \right) \\
 &\quad + 2zg^{(4)}(x) \left[z + \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right]^3 \\
 &\quad + 6zg'''(x) \left[z + \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right] \\
 &\quad \quad \times \left[-g'(x) + 2z \frac{\partial^2}{\partial \sigma \partial x} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right] \\
 &\quad + 2zg''(x) \left[-3g''(x) \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} - 3g'(x) \frac{\partial^2}{\partial \sigma \partial x} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right. \\
 &\quad \quad + 3z^2 \frac{\partial^3}{\partial \sigma \partial x^2} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} + 3z \frac{\partial^3}{\partial \sigma^2 \partial x} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \\
 &\quad \quad \left. + \frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right] \\
 &\quad + 2zg^{(4)}(x) \left[\frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \right]^3 + 2zg''(x) \frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \\
 &\quad + 6z \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m_r/2}}(x)|_{\sigma=0} g'''(x).
 \end{aligned}$$

We then proceed to get the derivatives needed for $m > 1/2$:

$$\begin{aligned}
 \frac{\partial^5 R_2}{\partial \sigma^5}(x, z, \sigma)|_{\sigma=0} &= -15zg''(x) \left(z^2g'''(x) + \left[-g'(x) + 2\frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] g''(x) \right) \\
 &\quad - 5g'(x) \left(2z^3g^{(4)}(x) + 6zg'''(x) \left[-g'(x) + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \right. \\
 &\quad \quad \left. + 2g''(x) \left[3z\frac{\partial^3}{\partial \sigma^2 \partial x} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + \frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \right. \\
 &\quad \quad \left. + 2g''(x)\frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right) \\
 &\quad + \frac{5}{2}g^{(5)}(x)z^5 + \frac{5}{2}zg''(x)\frac{\partial^4}{\partial \sigma^4} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad + 15g^{(4)}(x)z^3 \left[-g'(x) + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \\
 &\quad + \frac{15}{2}zg'''(x) \left[-g'(x) + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right]^2 \\
 &\quad + \frac{15}{2}zg'''(x) \left[\frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right]^2 \\
 &\quad + 10z^2g'''(x) \left(3z\frac{\partial^3}{\partial \sigma^2 \partial x} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + \frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right) \\
 &\quad + \frac{5}{2}zg''(x) \left(\frac{\partial^4}{\partial \sigma^4} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} - 6g''(x)\frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right. \\
 &\quad \quad \left. - 6g'(x)\frac{\partial^3}{\partial \sigma^2 \partial x} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + 6z^2\frac{\partial^4}{\partial \sigma^2 \partial x^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right. \\
 &\quad \quad \left. + 4z\frac{\partial^4}{\partial \sigma^3 \partial x} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right)
 \end{aligned}$$

and

$$\begin{aligned}
 \frac{\partial^6 R_2}{\partial \sigma^6}(x, z, \sigma)|_{\sigma=0} &= -\frac{45}{2} \left(z^2g'''(x) + \left[-g'(x) + 2\frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] g''(x) \right)^2 \\
 &\quad - 15zg''(x) \left(2z^3g^{(4)}(x) + 6zg'''(x) \left[-g'(x) + \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \right. \\
 &\quad \quad \left. + 2g''(x) \left[3z\frac{\partial^3}{\partial \sigma^2 \partial x} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + \frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \right. \\
 &\quad \quad \left. + 2g''(x)\frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right) \\
 &\quad + 6g'(x)A^{(5)} - zA^{(6)},
 \end{aligned}$$

with

$$\begin{aligned}
 A^{(5)} &= -\frac{5}{2}g^{(5)}(x)z^4 - \frac{5}{2}g''(x)\frac{\partial^4}{\partial\sigma^4}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad - 15g^{(4)}(x)z^2 \left[-g'(x) + \frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \\
 &\quad - \frac{15}{2}g'''(x) \left[-g'(x) + \frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right]^2 - \frac{15}{2}g'''(x) \left[\frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right]^2 \\
 &\quad - 10zg'''(x) \left(3z\frac{\partial^3}{\partial\sigma^2\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + \frac{\partial^3}{\partial\sigma^3}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right) \\
 &\quad - \frac{5}{2}g''(x) \left(\frac{\partial^4}{\partial\sigma^4}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} - 6g''(x)\frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right. \\
 &\quad \quad - 6g'(x)\frac{\partial^3}{\partial\sigma^2\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad \quad \left. + 6z^2\frac{\partial^4}{\partial\sigma^2\partial x^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + 4z\frac{\partial^4}{\partial\sigma^3\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right), \\
 A^{(6)} &= -3 \left(10g'''(x)\frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \frac{\partial^3}{\partial\sigma^3}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right. \\
 &\quad + g''(x)\frac{\partial^5}{\partial\sigma^5}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + g^{(6)}(x)z^5 \\
 &\quad + 10g^{(5)}(x)z^3 \left[-g'(x) + \frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \\
 &\quad + 15zg^{(4)}(x) \left[-g'(x) + \frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right]^2 \\
 &\quad + 10g^{(4)}(x)z^2 \left[\frac{\partial^3}{\partial\sigma^3}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + 3z\frac{\partial^3}{\partial\sigma^2\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \\
 &\quad + 10g'''(x) \left[-g'(x) + \frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \\
 &\quad \quad \times \left[\frac{\partial^3}{\partial\sigma^3}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + 3z\frac{\partial^3}{\partial\sigma^2\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \\
 &\quad + 5g'''(x)z \left[-6g''(x)\frac{\partial^2}{\partial\sigma^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} - 6g'(x)\frac{\partial^3}{\partial\sigma^2\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right. \\
 &\quad \quad + 6z^2\frac{\partial^4}{\partial\sigma^2\partial x^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + 3z\frac{\partial^4}{\partial\sigma^3\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad \quad \left. + \frac{\partial^4}{\partial\sigma^4}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right] \\
 &\quad + g''(x) \left[-10g''(x)\frac{\partial^3}{\partial\sigma^3}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} - 10g'(x)\frac{\partial^4}{\partial\sigma^3\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right. \\
 &\quad \quad + 5z\frac{\partial^5}{\partial\sigma^4\partial x}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + 10z^2\frac{\partial^5}{\partial\sigma^3\partial x^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad \quad + 10z^3\frac{\partial^5}{\partial\sigma^2\partial x^3}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \\
 &\quad \quad \left. - 30g'(x)z\frac{\partial^4}{\partial\sigma^2\partial x^2}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} + \frac{\partial^5}{\partial\sigma^5}\text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} \right].
 \end{aligned}$$

4.9.3 Derivatives of the Proximity Map for Differentiable Targets

Recall that, in the differentiable case, $\text{prox}_g^{\sigma^{2m}r/2}(x) = -\frac{\sigma^{2m}r}{2}g'(\text{prox}_g^{\sigma^{2m}r/2}(x)) + x$ then

$$\begin{aligned} \frac{\partial}{\partial \sigma} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} &= \begin{cases} -\frac{r}{2}g'(x) & \text{if } m = 1/2 \\ 0 & \text{if } m > 1/2 \\ \infty & \text{otherwise} \end{cases} \\ \frac{\partial^2}{\partial \sigma^2} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} &= \begin{cases} \frac{r^2}{2}g'(x)g''(x) & \text{if } m = 1/2 \\ -rg'(x) & \text{if } m = 1 \\ 0 & \text{if } m > 1 \\ \infty & \text{otherwise} \end{cases} \\ \frac{\partial^3}{\partial \sigma^3} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} &= \begin{cases} -\frac{3r^3}{8}g'''(x)[g'(x)]^2 - \frac{3r^3}{4}g'(x)[g''(x)]^2 & \text{if } m = 1/2 \\ -3rg'(x) & \text{if } m = 3/2 \\ 0 & \text{if } m = 1, m > 3/2 \\ \infty & \text{otherwise} \end{cases} \\ \frac{\partial^4}{\partial \sigma^4} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} &= \begin{cases} < \infty & \text{if } m = 1/2 \\ 6r^2g'(x)g''(x) & \text{if } m = 1 \\ -12rg'(x) & \text{if } m = 2 \\ 0 & \text{if } m = 3/2, m > 2 \\ \infty & \text{otherwise} \end{cases} \\ \frac{\partial^5}{\partial \sigma^5} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} &= \begin{cases} < \infty & \text{if } m = 1/2 \\ -60rg'(x) & \text{if } m = 5/2 \\ 0 & \text{if } m = 1, m = 3/2, m = 2, m > 5/2 \\ \infty & \text{otherwise} \end{cases} \end{aligned}$$

and

$$\frac{\partial}{\partial x} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} = 1, \quad \frac{\partial^{(k)}}{\partial x^{(k)}} \text{prox}_g^{\sigma^{2m}r/2}(x)|_{\sigma=0} = 0,$$

for all integers $k > 1$. For the mixed derivatives we have

$$\begin{aligned} \frac{\partial^2}{\partial \sigma \partial x} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} -\frac{r}{2} g''(x) & \text{if } m = 1/2 \\ 0 & \text{if } m > 1/2 \\ \infty & \text{otherwise} \end{cases} \\ \frac{\partial^3}{\partial \sigma^2 \partial x} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} \frac{r^2}{2} [g''(x)]^2 + \frac{r^2}{2} g'(x) g'''(x) & \text{if } m = 1/2 \\ -r g''(x) & \text{if } m = 1 \\ 0 & \text{if } m > 1/2 \\ \infty & \text{otherwise} \end{cases} \\ \frac{\partial^4}{\partial \sigma^3 \partial x} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} < \infty & \text{if } m = 1/2 \\ -3r g''(x) & \text{if } m = 3/2 \\ 0 & \text{if } m = 1, m > 3/2 \\ \infty & \text{otherwise} \end{cases} \\ \frac{\partial^5}{\partial \sigma^4 \partial x} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} < \infty & \text{if } m = 1/2 \\ 6r^2 (g'(x) g'''(x) + [g''(x)]^2) & \text{if } m = 1 \\ -12r g''(x) & \text{if } m = 2 \\ 0 & \text{if } m = 3/2, m > 2 \\ \infty & \text{otherwise} \end{cases} \end{aligned}$$

and

$$\begin{aligned}
 \frac{\partial^3}{\partial \sigma \partial x^2} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} -\frac{r}{2} g'''(x) & \text{if } m = 1/2 \\ 0 & \text{if } m > 1/2 \\ \infty & \text{otherwise} \end{cases} \\
 \frac{\partial^4}{\partial \sigma \partial x^3} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} < \infty & \text{if } m = 1/2 \\ 0 & \text{if } m > 1/2 \\ \infty & \text{otherwise} \end{cases} \\
 \frac{\partial^4}{\partial \sigma^2 \partial x^2} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} < \infty & \text{if } m = 1/2 \\ -r g'''(x) & \text{if } m = 1 \\ 0 & \text{if } m > 1/2 \\ \infty & \text{otherwise} \end{cases} \\
 \frac{\partial^5}{\partial \sigma^3 \partial x^2} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} < \infty & \text{if } m = 1/2 \\ -3r g'''(x) & \text{if } m = 3/2 \\ 0 & \text{if } m = 1, m > 3/2 \\ \infty & \text{otherwise} \end{cases} \\
 \frac{\partial^5}{\partial \sigma^2 \partial x^3} \text{prox}_g^{\sigma^{2m} r/2}(x)|_{\sigma=0} &= \begin{cases} < \infty & \text{if } m = 1/2 \\ -r g^{(4)}(x) & \text{if } m = 1 \\ 0 & \text{if } m > 1 \\ \infty & \text{otherwise} \end{cases}
 \end{aligned}$$

4.10 Moments and Integrals for the Laplace Distribution

4.10.1 Moments of Acceptance Ratio for the Laplace Distribution

The indicator functions in the definition of ϕ_d identify four different regions:

$$\begin{aligned}
 R_1 &:= \left\{ (x, z) : |x| < \sigma^{2m} r/2 \wedge \left| \left(1 - \frac{1}{\sigma^{2(m-1)r}}\right) x + \sigma z \right| < \sigma^{2m} r/2 \right\}, \\
 R_2 &:= \left\{ (x, z) : |x| \geq \sigma^{2m} r/2 \wedge \left| x - \frac{\sigma^2}{2} \text{sgn}(x) + \sigma z \right| < \sigma^{2m} r/2 \right\}, \\
 R_3 &:= \left\{ (x, z) : |x| < \sigma^{2m} r/2 \wedge \left| \left(1 - \frac{1}{\sigma^{2(m-1)r}}\right) x + \sigma z \right| \geq \sigma^{2m} r/2 \right\}, \\
 R_4 &:= \left\{ (x, z) : |x| > \sigma^{2m} r/2 \wedge \left| x - \frac{\sigma^2}{2} \text{sgn}(x) + \sigma z \right| > \sigma^{2m} r/2 \right\},
 \end{aligned}$$

with corresponding acceptance ratios

$$\begin{aligned}
 \phi_d^1(x, z) &= |x| - \left| \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) x + \sigma z \right| + \frac{z^2}{2} \\
 &\quad - \frac{1}{2\sigma^2} \left(\left(\frac{2}{\sigma^{2(m-1)r}} - \frac{1}{\sigma^{4(m-1)r^2}} \right) x - \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \sigma z \right)^2 \\
 \phi_d^2(x, z) &= |x| - \left| x - \frac{\sigma^2}{2} \operatorname{sgn}(x) + \sigma z \right| + \frac{z^2}{2} \\
 &\quad - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} x + \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \left(\frac{\sigma^2}{2} \operatorname{sgn}(x) - \sigma z \right) \right)^2 \\
 \phi_d^3(x, z) &= |x| - \left| \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) x + \sigma z \right| + \frac{z^2}{2} \\
 &\quad - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} x - \sigma z + \frac{\sigma^2}{2} \operatorname{sgn} \left[\left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) x + \sigma z \right] \right)^2 \\
 \phi_d^4(x, z) &= |x| - \left| x - \frac{\sigma^2}{2} \operatorname{sgn}(x) + \sigma z \right| + \frac{z^2}{2} \\
 &\quad - \frac{1}{2\sigma^2} \left(\frac{\sigma^2}{2} \operatorname{sgn}(x) - \sigma z + \frac{\sigma^2}{2} \operatorname{sgn} \left[x - \frac{\sigma^2}{2} \operatorname{sgn}(x) + \sigma z \right] \right)^2.
 \end{aligned}$$

Let us denote

$$\begin{aligned}
 A_1 &:= \left\{ x : 0 \leq x < \frac{\sigma^{2m}r}{2} \right\}, & A_2 &:= \left\{ x : -\frac{\sigma^{2m}r}{2} < x \leq 0 \right\}, \\
 A_3 &:= \left\{ x : x \geq \frac{\sigma^{2m}r}{2} \right\}, & A_4 &:= \left\{ x : x \leq -\frac{\sigma^{2m}r}{2} \right\},
 \end{aligned}$$

and

$$\begin{aligned}
 B_1 &:= \left\{ z : 0 \leq \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) x + \sigma z < \frac{\sigma^{2m}r}{2} \right\}, \\
 B_2 &:= \left\{ z : -\frac{\sigma^{2m}r}{2} < \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) x + \sigma z \leq 0 \right\}, \\
 B_3 &:= \left\{ z : \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) x + \sigma z \geq \frac{\sigma^{2m}r}{2} \right\}, \\
 B_4 &:= \left\{ z : \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) x + \sigma z \leq -\frac{\sigma^{2m}r}{2} \right\},
 \end{aligned}$$

and

$$\begin{aligned}
 C_1 &:= \left\{ (x, z) : 0 \leq x - \frac{\sigma^2}{2} \operatorname{sgn}(x) + \sigma z < \frac{\sigma^{2m_r}}{2} \right\}, \\
 C_2 &:= \left\{ (x, z) : -\frac{\sigma^{2m_r}}{2} < x - \frac{\sigma^2}{2} \operatorname{sgn}(x) + \sigma z \leq 0 \right\} \\
 C_3 &:= \left\{ (x, z) : x - \frac{\sigma^2}{2} \operatorname{sgn}(x) + \sigma z \geq \frac{\sigma^{2m_r}}{2} \right\}, \\
 C_4 &:= \left\{ (x, z) : x - \frac{\sigma^2}{2} \operatorname{sgn}(x) + \sigma z \leq -\frac{\sigma^{2m_r}}{2} \right\},
 \end{aligned}$$

so that, up to a set of null measure,

$$\begin{aligned}
 R_1 &= (A_1 \cup A_2) \cap (B_1 \cup B_2), & R_2 &= (A_3 \cup A_4) \cap (C_1 \cup C_2), \\
 R_3 &= (A_1 \cup A_2) \cap (B_3 \cup B_4), & R_4 &= (A_3 \cup A_4) \cap (C_3 \cup C_4).
 \end{aligned}$$

Proposition 4.22. *Take X a Laplace random variable and Z a standard normal random variable independent of X , then if $\sigma^2 = \ell^2 d^{-2/3}$, we have*

$$\lim_{d \rightarrow +\infty} d\mathbb{E}[\phi_d(X, Z)] = -\frac{\ell^3}{3\sqrt{2\pi}}.$$

Proof. Taking expectations of $\phi_d^i \mathbb{1}_{R_i}$ for $i = 1, \dots, 4$ and exploiting the symmetry of the

laws of X and Z , we can write

$$\begin{aligned}
 & \mathbb{E} \left[\phi_d^1(X, Z) \mathbb{1}_{R_1}(X, Z) \right] \\
 &= 2\mathbb{E} \left[\left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_1}(X, Z) \right] \\
 &+ 2\mathbb{E} \left[\left(2X - \frac{1}{\sigma^{2(m-1)r}} X + \sigma Z \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_2}(X, Z) \right] \\
 &+ 2\mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\left(\frac{2}{\sigma^{2(m-1)r}} - \frac{1}{\sigma^{4(m-1)r^2}} \right) X - \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \sigma Z \right)^2 \right) \right. \\
 &\quad \left. \times \mathbb{1}_{A_1}(X) \mathbb{1}_{B_1 \cup B_2}(X, Z) \right], \\
 & \mathbb{E} \left[\phi_d^2(X, Z) \mathbb{1}_{R_2}(X, Z) \right] \\
 &= 2\mathbb{E} \left[\left(\frac{\sigma^2}{2} - \sigma Z \right) \mathbb{1}_{A_3}(X) \mathbb{1}_{C_1}(X, Z) \right] \\
 &+ 2\mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z \right) \mathbb{1}_{A_3}(X) \mathbb{1}_{C_2}(X, Z) \right] \\
 &+ 2\mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X + \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \left(\frac{\sigma^2}{2} - \sigma Z \right) \right) \right)^2 \right) \right. \\
 &\quad \left. \times \mathbb{1}_{A_3}(X) \mathbb{1}_{C_1 \cup C_2}(X, Z) \right], \\
 & \mathbb{E} \left[\phi_d^3(X, Z) \mathbb{1}_{R_3}(X, Z) \right] \\
 &= 2\mathbb{E} \left[\left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_3}(X, Z) \right] \\
 &+ 2\mathbb{E} \left[\left(2X - \frac{1}{\sigma^{2(m-1)r}} X + \sigma Z \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z) \right] \\
 &+ 2\mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{\sigma^2}{2} \right)^2 \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_3}(X, Z) \right] \\
 &+ 2\mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z - \frac{\sigma^2}{2} \right)^2 \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z) \right], \\
 & \mathbb{E} \left[\phi_d^4(X, Z) \mathbb{1}_{R_4}(X, Z) \right] = 2\mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z \right) \mathbb{1}_{A_3}(X) \mathbb{1}_{C_4}(X, Z) \right].
 \end{aligned}$$

Using the integrals in Appendix 4.10.4 and Lebesgue's dominated convergence theorem,

we find that for $\alpha = \beta = 1/3$ and $r \geq 0$

$$\begin{aligned}\lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d^1(X, Z)] &= 0 \\ \lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d^2(X, Z)] &= -2 \frac{\ell^3 r}{4\sqrt{2\pi}} \int_{-\infty}^0 e^{-z^2/2} z dz = \frac{\ell^3 r}{2\sqrt{2\pi}} \\ \lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d^3(X, Z)] &= \frac{3\ell^3 r}{8\sqrt{2\pi}} \int_{-\infty}^0 e^{-z^2/2} z dz - \frac{\ell^3 r}{8\sqrt{2\pi}} \int_0^{+\infty} e^{-z^2/2} z dz = -\frac{\ell^3 r}{2\sqrt{2\pi}} \\ \lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d^4(X, Z)] &= \frac{\ell^3}{6\sqrt{2\pi}} \int_{-\infty}^0 e^{-z^2/2} z^3 dz = -\frac{\ell^3}{3\sqrt{2\pi}},\end{aligned}$$

which gives

$$\begin{aligned}\lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d(X, Z)] &= \lim_{d \rightarrow +\infty} d \left(\mathbb{E} [\phi_d^1(X, Z)] + \mathbb{E} [\phi_d^2(X, Z)] + \mathbb{E} [\phi_d^3(X, Z)] + \mathbb{E} [\phi_d^4(X, Z)] \right) \\ &= -\frac{\ell^3}{3\sqrt{2\pi}}.\end{aligned}$$

For $\alpha = 1/3, \beta = m/3$ for $m > 1$ and $r \geq 0$ we have

$$\begin{aligned}\lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d^1(X, Z)] &= 0 \\ \lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d^2(X, Z)] &= 0 \\ \lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d^3(X, Z)] &= 0 \\ \lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d^4(X, Z)] &= \frac{\ell^3}{6\sqrt{2\pi}} \int_{-\infty}^0 e^{-z^2/2} z^3 dz = -\frac{\ell^3}{3\sqrt{2\pi}},\end{aligned}$$

which gives

$$\begin{aligned}\lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d(X, Z)] &= \lim_{d \rightarrow +\infty} d \left(\mathbb{E} [\phi_d^1(X, Z)] + \mathbb{E} [\phi_d^2(X, Z)] + \mathbb{E} [\phi_d^3(X, Z)] + \mathbb{E} [\phi_d^4(X, Z)] \right) \\ &= -\frac{\ell^3}{3\sqrt{2\pi}}.\end{aligned}$$

□

Proposition 4.23. *Take X a Laplace random variable and Z a standard normal random variable independent of X , then if $\sigma^2 = \ell^2 d^{-2/3}$*

$$\lim_{d \rightarrow +\infty} d \text{Var} (\phi_d(X, Z)) = \frac{2\ell^3}{3\sqrt{2\pi}}.$$

Proof. As a consequence of the previous Proposition we have

$$\lim_{d \rightarrow +\infty} d\mathbb{E} [\phi_d(X, Z)]^2 = 0.$$

Then, because $R_j \cap R_i = \emptyset$ for all $j \neq i$, we have that

$$\begin{aligned} \mathbb{E} \left[\phi_d(X, Z)^2 \right] &= \mathbb{E} \left[\phi_d^1(X, Z)^2 R_1(X, Z) \right] + \mathbb{E} \left[\phi_d^2(X, Z)^2 R_2(X, Z) \right] \\ &\quad + \mathbb{E} \left[\phi_d^3(X, Z)^2 R_3(X, Z) \right] + \mathbb{E} \left[\phi_d^4(X, Z)^2 R_4(X, Z) \right], \end{aligned}$$

and, exploiting again the symmetry of the laws of X and Z , we have

$$\begin{aligned} &\mathbb{E} \left[\phi_d^1(X, Z)^2 R_1(X, Z) \right] \\ &= 2\mathbb{E} \left[\left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\left(\frac{2}{\sigma^{2(m-1)r}} - \frac{1}{\sigma^{4(m-1)r^2}} \right) X - \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \sigma Z \right) \right)^2 \right. \\ &\quad \left. \times \mathbb{1}_{A_1}(X) \mathbb{1}_{B_1}(X, Z) \right], \\ &+ 2\mathbb{E} \left[\left(2X - \frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\left(\frac{2}{\sigma^{2(m-1)r}} - \frac{1}{\sigma^{4(m-1)r^2}} \right) X - \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \sigma Z \right) \right)^2 \right. \\ &\quad \left. \times \mathbb{1}_{A_1}(X) \mathbb{1}_{B_2}(X, Z) \right], \\ &\mathbb{E} \left[\phi_d^2(X, Z)^2 \mathbb{1}_{R_2}(X, Z) \right] \\ &= 2\mathbb{E} \left[\left(\frac{\sigma^2}{2} - \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X + \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \left(\frac{\sigma^2}{2} - \sigma Z \right) \right) \right)^2 \mathbb{1}_{A_3}(X) \mathbb{1}_{C_1}(X, Z) \right] \\ &+ 2\mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X + \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \left(\frac{\sigma^2}{2} - \sigma Z \right) \right) \right)^2 \right. \\ &\quad \left. \times \mathbb{1}_{A_1}(X) \mathbb{1}_{C_2}(X, Z) \right], \\ &\mathbb{E} \left[\phi_d^3(X, Z)^2 \mathbb{1}_{R_3}(X, Z) \right] \\ &= 2\mathbb{E} \left[\left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{\sigma^2}{2} \right) \right)^2 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_3}(X, Z) \right] \\ &+ 2\mathbb{E} \left[\left(2X - \frac{1}{\sigma^{2(m-1)r}} X + \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z - \frac{\sigma^2}{2} \right) \right)^2 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z) \right], \\ &\mathbb{E} \left[\phi_d^4(X, Z)^2 \mathbb{1}_{R_4}(X, Z) \right] = 2\mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z \right)^2 \mathbb{1}_{A_3}(X) \mathbb{1}_{C_4}(X, Z) \right]. \end{aligned}$$

Proceeding as for Proposition 4.22, using the integrals in Appendix 4.10.4 and Lebesgue's dominated convergence theorem we can then show that for $\alpha = 1/3, \beta = m/3$ for $m \geq 1$ and $r \geq 0$

$$\lim_{d \rightarrow +\infty} d \text{Var}(\phi_d(X, Z)) = \frac{2\ell^3}{3\sqrt{2\pi}}.$$

□

Proposition 4.24. *Take X a Laplace random variable and Z a standard normal random variable independent of X , then if $\sigma^2 = \ell^2 d^{-2/3}$ we have*

$$\lim_{d \rightarrow +\infty} d\mathbb{E} \left[\phi_d(X, Z)^3 \right] = 0.$$

Proof. Following the same structure of the previous propositions we have that

$$\begin{aligned} \mathbb{E} \left[\phi(X, Z)^3 \right] &= \mathbb{E} \left[\phi_d^1(X, Z)^3 R_1(X, Z) \right] + \mathbb{E} \left[\phi_d^3(X, Z)^2 R_2(X, Z) \right] \\ &\quad + \mathbb{E} \left[\phi_d^3(X, Z)^3 R_3(X, Z) \right] + \mathbb{E} \left[\phi_d^4(X, Z)^3 R_4(X, Z) \right], \end{aligned}$$

exploiting again the symmetry of the laws of X and Z , using the integrals in Appendix 4.10.4, the dominated convergence theorem we can then show that

$$\lim_{d \rightarrow +\infty} d\mathbb{E} \left[\phi_d(X, Z)^3 \right] = 0.$$

□

4.10.2 Bound on Second Moment of Acceptance Ratio for the Laplace Distribution

Lemma 3. *Let Z be a standard normal random variable and $\sigma = \ell/d^\alpha$ for $\alpha = 1/3$. Then, there exists a constant $C > 0$ such that for all $a \in \mathbb{R}$ and $d \in \mathbb{N}$:*

$$\mathbb{E} \left[\phi_d(a, Z)^2 \right] \leq \frac{C}{d^{2\alpha}}.$$

Proof. We consider the case $a \geq 0$ and $r \geq \sigma^{2(m-1)}$ only, all the other cases follow from identical arguments. As in the derivation of the moments of ϕ_d in Appendix 4.10.1, we distinguish four regions. We recall that $\sigma = \ell/d^\alpha$ for $\alpha = 1/3$ and thus $\sigma^{p+1} \leq \sigma^p$ for all

$p \in \mathbb{N}$. Take $r \geq \sigma^{-2(m-1)}$, for R_1 , we have, using Hölder's inequality multiple times,

$$\begin{aligned}
 \mathbb{E} \left[\phi_d^1(a, Z)^2 \mathbb{1}_{R_1}(a, Z) \right] &= \mathbb{E} \left[\phi_d^1(a, Z)^2 \mathbb{1}_{B_1 \cup B_2}(a, Z) \right] \\
 &\leq C\sigma^2 \int_{-(1-1/\sigma^2(m-1)r)a/\sigma}^{\sigma^{2m-1}r/2-(1-1/\sigma^2(m-1)r)a} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left[\left(\frac{1}{\sigma^{2m-1}r} a - z \right)^2 \right. \\
 &\quad \left. + \left(\frac{z^2}{2\sigma} - \frac{1}{2\sigma^3} \left(\left(\frac{2}{\sigma^{2(m-1)r}} - \frac{1}{\sigma^{4(m-1)r^2}} \right) a - \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \sigma z \right)^2 \right)^2 \right] dz \\
 &+ C\sigma^2 \int_{-\sigma^{2m-1}r/2-(1-1/\sigma^2(m-1)r)a}^{-(1-1/\sigma^2(m-1)r)a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left[\left(\frac{2a}{\sigma} - \frac{1}{\sigma^{2m-1}r} a + z \right)^2 \right. \\
 &\quad \left. + \left(\frac{z^2}{2\sigma} - \frac{1}{2\sigma^3} \left(\left(\frac{2}{\sigma^{2(m-1)r}} - \frac{1}{\sigma^{4(m-1)r^2}} \right) a - \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \sigma z \right)^2 \right)^2 \right] dz \\
 &\leq C\sigma^2 \int_{-\infty}^{+\infty} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(4 \left(\frac{a}{\sigma} \right)^2 + 2 \left(\frac{1}{\sigma^{2m-1}r} a \right)^2 + 2z^2 \right) dz \\
 &+ C\sigma^2 \int_{-\sigma^{2m-1}r/2-(1-1/\sigma^2(m-1)r)a}^{\sigma^{2m-1}r/2-(1-1/\sigma^2(m-1)r)a} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \\
 &\quad \times \left(\frac{z^2}{2\sigma} - \frac{1}{2\sigma^3} \left(\left(\frac{2}{\sigma^{2(m-1)r}} - \frac{1}{\sigma^{4(m-1)r^2}} \right) a - \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \sigma z \right)^2 \right)^2 \\
 &\leq C\sigma^2 + C\sigma^2 \int_{-\sigma^{2m-1}r/2-(1-1/\sigma^2(m-1)r)a}^{\sigma^{2m-1}r/2-(1-1/\sigma^2(m-1)r)a} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \\
 &\quad \times \left(\frac{z^4}{4\sigma^2} + \frac{1}{4\sigma^6} \left(\left(\frac{2}{\sigma^{2(m-1)r}} - \frac{1}{\sigma^{4(m-1)r^2}} \right)^4 a^4 + \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right)^4 \sigma^4 z^4 \right) \right) dz \\
 &\leq C\sigma^2,
 \end{aligned}$$

where we used the fact that the moments of Z are bounded and $a \leq \sigma^{2m}r/2$ for the first term, and the fact that $z \leq \sigma^{2m-1}r/2$ for the second one. Proceeding as above, for R_3 ,

$a > 0$ and $r \geq \sigma^{-2(m-1)}$, we have

$$\begin{aligned}
 \mathbb{E} \left[\phi_d^3(a, Z)^2 \mathbb{1}_{R_3}(a, Z) \right] &= \mathbb{E} \left[\phi_d^3(a, Z)^2 \mathbb{1}_{B_3 \cup B_4}(a, Z) \right] \\
 &\leq C\sigma^2 \int_{\sigma^{2m-1}r/2 - (1-1/\sigma^{2(m-1)}r)a/\sigma}^{+\infty} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left[\left(\frac{1}{\sigma^{2m-1}r} a - z \right)^2 \right. \\
 &\quad \left. + \left(\frac{z^2}{2\sigma} - \frac{1}{2\sigma^3} \left(\frac{1}{\sigma^{2(m-1)}r} a - \sigma z + \frac{\sigma^2}{2} \right)^2 \right)^2 \right] dz \\
 &+ C\sigma^2 \int_{-\infty}^{-\sigma^{2m-1}r/2 - (1-1/\sigma^{2(m-1)}r)a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left[\left(\frac{2a}{\sigma} - \frac{1}{\sigma^{2m-1}r} a + z \right)^2 \right. \\
 &\quad \left. + \left(\frac{z^2}{2\sigma} - \frac{1}{2\sigma^3} \left(\frac{1}{\sigma^{2(m-1)}r} a - \sigma z + \frac{\sigma^2}{2} \right)^2 \right)^2 \right] dz \\
 &\leq C\sigma^2 + C\sigma^2 \int_{-\infty}^{+\infty} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(\frac{1}{2\sigma^3} \left(\frac{a^2}{\sigma^{4(m-1)}r^2} + \frac{\sigma^4}{4} - \sigma^3 z + \frac{a}{\sigma^{2m-4}r} - \frac{2az}{\sigma^{2m-3}r} \right) \right)^2 dz \\
 &\leq C\sigma^2,
 \end{aligned}$$

where we used again the boundedness of the moments of Z , the fact that $a \leq \sigma^{2m}r/2$ and that $\sigma^{p+1} \leq \sigma^p$. For R_2 and $a > 0$, we have

$$\begin{aligned}
 \mathbb{E} \left[\phi_d^2(a, Z)^2 \mathbb{1}_{R_2}(a, Z) \right] &= \mathbb{E} \left[\phi_d^2(a, Z)^2 \mathbb{1}_{C_1 \cup C_2}(a, Z) \right] \\
 &\leq C\sigma^2 \int_{\sigma/2 - a/\sigma}^{\sigma/2 + \sigma^{2m-1}r/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left[\left(\frac{\sigma^2}{2} - \sigma z \right)^2 \right. \\
 &\quad \left. + \left(\frac{z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)}r} a + \left(1 - \frac{1}{\sigma^{2(m-1)}r} \right) \left(\frac{\sigma^2}{2} - \sigma z \right) \right) \right)^2 \right] dz \\
 &+ C\sigma^2 \int_{\sigma/2 - \sigma^{2m-1}r/2 - a/\sigma}^{\sigma/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left[\left(2a - \frac{\sigma^2}{2} + \sigma z \right)^2 \right. \\
 &\quad \left. + \left(\frac{z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)}r} a + \left(1 - \frac{1}{\sigma^{2(m-1)}r} \right) \left(\frac{\sigma^2}{2} - \sigma z \right) \right) \right)^2 \right] dz.
 \end{aligned}$$

The first integral is bounded using the moments of Z , while for the third one let us denote

$\chi(a, \sigma, z) := a - \sigma^2/2 + \sigma z$, then

$$\begin{aligned}
 & \int_{\sigma/2 - \sigma^{2m-1}r/2 - a/\sigma}^{\sigma/2 + \sigma^{2m-1}r/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(\frac{z^2}{2\sigma} - \frac{1}{2\sigma^3} \left(\frac{1}{\sigma^{2(m-1)}r} a + \left(1 - \frac{1}{\sigma^{2(m-1)}r} \right) \left(\frac{\sigma^2}{2} - \sigma z \right) \right) \right)^2 dz \\
 &= \int_{\sigma/2 - \sigma^{2m-1}r/2 - a/\sigma}^{\sigma/2 + \sigma^{2m-1}r/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(\frac{z^2}{2\sigma} - \frac{1}{2\sigma^3} \left(\frac{\chi(a, \sigma, z)}{\sigma^{2(m-1)}r} + \frac{\sigma^2}{2} - \sigma z \right) \right)^2 dz \\
 &\leq C \int_{\sigma/2 - \sigma^{2m-1}r/2 - a/\sigma}^{\sigma/2 + \sigma^{2m-1}r/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(\frac{z^2}{2\sigma} - \frac{1}{2\sigma^3} \left(\frac{\sigma^2}{2} - \sigma z \right) \right)^2 dz \\
 &+ C \int_{\sigma/2 - \sigma^{2m-1}r/2 - a/\sigma}^{\sigma/2 + \sigma^{2m-1}r/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(\frac{\chi(a, \sigma, z)^2}{2\sigma^{4m-1}r^2} \right)^2 + \left(\frac{\chi(a, \sigma, z)}{r\sigma^{2m-1}} \left(\frac{\sigma^2}{2} - \sigma z \right) \right)^2 dz;
 \end{aligned}$$

recalling that in R_2 we have $|\chi(a, \sigma, z)| \leq \sigma^{2m}r/2$, we obtain that this term is also bounded by $C\sigma^2$. For R_4 and $a > 0$, we have

$$\begin{aligned}
 \mathbb{E} \left[\phi_d^4(a, Z)^2 \mathbb{1}_{R_4}(a, Z) \right] &= \mathbb{E} \left[\phi_d^4(a, Z)^2 \mathbb{1}_{C_4}(a, Z) \right] \\
 &= \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(2a - \frac{\sigma^2}{2} + \sigma z \right)^2 dz \\
 &= \sigma^2 \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(\frac{2a}{\sigma} - \frac{\sigma}{2} + z \right)^2 dz
 \end{aligned}$$

Collecting all the terms together, we obtain

$$\begin{aligned}
 \mathbb{E} \left[\phi_d(a, Z)^2 \right] &= \sum_{i=1}^4 \mathbb{E} \left[\phi_d^i(a, Z)^2 \mathbb{1}_{T_i}(a, Z) \right] \\
 &\leq C\sigma^2 + C\sigma^2 \int_{-\infty}^{\sigma/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(\frac{2a}{\sigma} - \frac{\sigma}{2} + z \right)^2 dz.
 \end{aligned}$$

Recall that $\sigma = \ell d^{-1/3}$. To bound the last integral we use Hölder's inequality

$$\begin{aligned}
 \int_{-\infty}^{\sigma/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left(\frac{2a}{\sigma} - \frac{\sigma}{2} + z \right)^2 dz &\leq C \int_{-\infty}^{\sigma/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left[\left(\frac{\sigma}{2} + z \right)^2 + 4 \left(\frac{\sigma}{2} - \frac{a}{\sigma} \right)^2 \right] dz \\
 &\leq C \int_{-\infty}^{\sigma/2 - a/\sigma} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left[\left(\frac{\ell^2}{4} + z^2 \right) + 4 \left(\frac{\sigma}{2} - \frac{a}{\sigma} \right)^2 \right] dz.
 \end{aligned}$$

The first term is bounded since the moments of Z are bounded. For the second term we use an estimate of the Gaussian cumulative distribution function. Let $\kappa(\ell, d, a) := \ell d^{-1/3}/2 - ad^{1/3}/\ell$. When $z < \kappa(\ell, d, a) < 0$, we have $1 < z/\kappa(\ell, d, a)$ and therefore

$$\begin{aligned}
 (2\pi)^{-1/2} \kappa(\ell, d, a)^2 \int_{-\infty}^{\kappa(\ell, d, a)} e^{-z^2/2} dz &\leq (2\pi)^{-1/2} \kappa(\ell, d, a) \int_{-\infty}^{\kappa(\ell, d, a)} z e^{-z^2/2} dz, \\
 &= (2\pi)^{-1/2} \kappa(\ell, d, a) \exp(-\kappa(\ell, d, a)^2/2).
 \end{aligned}$$

However $y \mapsto ye^{-y^2/2}$ is bounded over \mathbb{R} , therefore $(a, d) \mapsto (2\pi)^{-1/2}\kappa(\ell, d, a) \exp(-\kappa(\ell, d, a)^2/2)$ is bounded over $\mathbb{R}_+^* \times \mathbb{N}$. If $\kappa(\ell, d, a) \geq 0$, then we still have $\kappa(\ell, d, a) < \ell$ and thus have the inequality

$$(2\pi)^{-1/2}\kappa(\ell, d, a)^2 \int_{-\infty}^{\kappa(\ell, d, a)} e^{-z^2/2} dz \leq (2\pi)^{-1/2}\ell^2 \int_{-\infty}^{+\infty} e^{-z^2/2} dz = \ell^2.$$

The result then follows since $\sigma = \ell d^{-1/3}$. □

4.10.3 Additional Integrals for the Laplace Distribution

We collect here two auxiliary Lemmata which are used in the proof of Proposition 4.9.

Lemma 4. *Take X a Laplace random variable and Z a standard normal random variable independent of X . Let $\tilde{X} := X - \frac{1}{\sigma^2(m-1)r} X \mathbb{1}\{|X| < \sigma^{2m}r/2\} - \frac{\sigma^2}{2} \operatorname{sgn}(X) \mathbb{1}\{|X| \geq \sigma^{2m}r/2\} + \sigma Z$, then, for $\sigma = \ell d^{-\alpha}$ with $\alpha = 1/3$,*

$$\mathbb{E} \left[\mathbb{1} \left\{ \operatorname{sgn}(X) \neq \operatorname{sgn}(\tilde{X}) \right\} \right] \rightarrow 0$$

if $d \rightarrow \infty$.

Proof. Using the same strategy of Appendix 4.10.1 and the symmetry of the laws of X, Z , we find that

$$\begin{aligned} \mathbb{E} \left[\mathbb{1} \left\{ \operatorname{sgn}(X) \neq \operatorname{sgn}(\tilde{X}) \right\} \right] &= 2\mathbb{E} [\mathbb{1}_{A_1}(X) \mathbb{1}_{B_2}(X, Z)] + 2\mathbb{E} [\mathbb{1}_{A_3}(X) \mathbb{1}_{C_2}(X, Z)] \\ &\quad + 2\mathbb{E} [\mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z)] + 2\mathbb{E} [\mathbb{1}_{A_3}(X) \mathbb{1}_{C_4}(X, Z)]. \end{aligned}$$

Using the same strategy used to obtain the moments of ϕ_d in Appendix 4.10.1, we find that

$$\mathbb{E} [\mathbb{1}_{A_1}(X) \mathbb{1}_{B_2}(X, Z)] = o(1),$$

in addition

$$\begin{aligned} \mathbb{E} [\mathbb{1}_{A_3}(X) \mathbb{1}_{C_2}(X, Z)] &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \int_{\sigma^2/2 - \sigma z}^{\sigma^2/2 - \sigma z + \sigma^{2m}r/2} e^{-x} dx dz + o(1) \\ &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \left[\frac{\sigma^2 r}{2} \delta_{m1} + \dots \right] dz + o(1), \end{aligned}$$

where δ_{m1} is a Dirac's delta, and

$$\begin{aligned} &\mathbb{E} [\mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z)] + \mathbb{E} [\mathbb{1}_{A_3}(X) \mathbb{1}_{C_4}(X, Z)] \\ &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \int_0^{\sigma^2/2 - \sigma z - \sigma^{2m}r/2} e^{-x} dx dz + o(1) \\ &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} [-\sigma z + \dots] dz + o(1). \end{aligned}$$

4.10. Moments and Integrals for the Laplace Distribution

Since $\sigma = \ell d^{-1/3}$ and the remainder terms of the Taylor expansions are bounded, Lebesgue's dominated convergence theorem gives

$$\begin{aligned}\mathbb{E} [\mathbb{1}_{A_3}(X) \mathbb{1}_{C_2}(X, Z)] &\rightarrow 0, \\ \mathbb{E} [\mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z)] + \mathbb{E} [\mathbb{1}_{A_3}(X) \mathbb{1}_{C_4}(X, Z)] &\rightarrow 0\end{aligned}$$

as $d \rightarrow \infty$. □

Lemma 5. *Take X a Laplace random variable and Z a standard normal random variable independent of X . Then,*

$$d^\alpha \mathbb{E} [|Z| |\phi_d(X, Z)|] \rightarrow 0$$

for $\alpha = 1/3$.

Proof. Using Cauchy-Schwarz's inequality we have that

$$\mathbb{E} [|Z| |\phi_d(X, Z)|] \leq \mathbb{E} [Z^2]^{1/2} \mathbb{E} [\phi_d(X, Z)^2]^{1/2};$$

the first expectation is equal to one, and the second one converges to zero at rate $d^{1/2}$ by Proposition 4.23. The result follows straightforwardly. □

4.10.4 Integrals for Moment Computations

We distinguish the case $m = 1/2$ and $m \geq 1$ since the integration bounds significantly differ in these two cases. For values between $1/2$ and 1 the integrals are not finite. The expectations below are obtained by integrating w.r.t. x and using a Taylor expansion about $\sigma = 0$ to obtain the leading order terms. Using the Lagrange form of the remainder for the Taylor expansions, we find that the remainder terms are all of the form $\sigma^{1/\alpha+1} f(\gamma(\sigma, z))/(1/\alpha+1)!$ where $\gamma(\sigma, z)$ is a point between the limits of integration w.r.t. x and $f : x \mapsto p(x)e^{-x}$, where p is a polynomial. Therefore, using the boundedness of the remainder and Lebesgue's dominated convergence theorem, the integrals w.r.t. z of the remainder terms all converge to 0.

4.10.4.1 First Moment

For simplicity, we only consider the case for $r \geq \sigma^{-2(m-1)}$, the other case follows analogously.

Region R_1 Let us consider ϕ_d^1 first. We have

$$A_1 \cap B_1 = \begin{cases} 0 \leq x \leq \frac{\sigma^{2m}r}{2} & \text{if } 0 \leq z \leq \frac{\sigma}{2} \\ 0 \leq x \leq \left(\frac{\sigma^{2m}r}{2} - \sigma z\right) \left(1 - \frac{1}{\sigma^{2(m-1)}r}\right)^{-1} & \text{if } \frac{\sigma}{2} \leq z \leq \frac{\sigma^{2m-1}r}{2} \\ -\sigma z \left(1 - \frac{1}{\sigma^{2(m-1)}r}\right)^{-1} \leq x \leq \frac{\sigma^{2m}r}{2} & \text{if } \frac{\sigma}{2} - \frac{\sigma^{2m-1}r}{2} \leq z \leq 0 \end{cases},$$

$$A_1 \cap B_2 = \begin{cases} 0 \leq x \leq \frac{\sigma^{2m}r}{2} & \text{if } -\frac{\sigma^{2m-1}r}{2} \leq z \leq \frac{\sigma}{2} - \frac{\sigma^{2m-1}r}{2} \\ 0 \leq x \leq -\sigma z \left(1 - \frac{1}{\sigma^{2(m-1)}r}\right)^{-1} & \text{if } \frac{\sigma}{2} - \frac{\sigma^{2m-1}r}{2} \leq z \leq 0 \\ -\left(\frac{\sigma^{2m}r}{2} + \sigma z\right) \left(1 - \frac{1}{\sigma^{2(m-1)}r}\right)^{-1} \leq x \leq \frac{\sigma^{2m}r}{2} & \text{if } \frac{\sigma}{2} - \sigma^{2m-1}r \leq z \leq -\frac{\sigma^{2m-1}r}{2} \end{cases},$$

and

$$A_1 \cap (B_1 \cup B_2) = \begin{cases} 0 \leq x \leq \frac{\sigma^{2m}r}{2} & \text{if } -\frac{\sigma^{2m-1}r}{2} \leq z \leq \frac{\sigma}{2} \\ 0 \leq x \leq \left(\frac{\sigma^{2m}r}{2} - \sigma z\right) \left(1 - \frac{1}{\sigma^{2(m-1)}r}\right)^{-1} & \text{if } \frac{\sigma}{2} \leq z \leq \frac{\sigma^{2m-1}r}{2} \\ -\left(\frac{\sigma^{2m}r}{2} + \sigma z\right) \left(1 - \frac{1}{\sigma^{2(m-1)}r}\right)^{-1} \leq x \leq \frac{\sigma^{2m}r}{2} & \text{if } \frac{\sigma}{2} - \sigma^{2m-1}r \leq z \leq -\frac{\sigma^{2m-1}r}{2} \end{cases}.$$

The corresponding expectations are

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{X}{\sigma^{2(m-1)}r} - \sigma Z \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_1}(X, Z) \right] \\ &= \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2}^{\sigma^{2m-1}r/2} e^{-z^2/2} \left[z^{4-2m} \sigma^{4-2m} \xi(r) + \dots \right] dz \\ &+ \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r/2}^0 e^{-z^2/2} \left[z^{4-2m} \sigma^{4-2m} \xi(r) + \dots \right] dz \\ &+ o(1), \\ & \mathbb{E} \left[\left(2X - \frac{X}{\sigma^{2(m-1)}r} + \sigma Z \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_2}(X, Z) \right] \\ &= \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r/2}^0 e^{-z^2/2} \left[z^{4-2m} \sigma^{4-2m} \xi(r) + \dots \right] dz \\ &+ \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r}^{-\sigma^{2m-1}r/2} e^{-z^2/2} \left[z^{4-2m} \sigma^{4-2m} \xi(r) + \dots \right] dz \\ &+ o(1), \end{aligned}$$

4.10. Moments and Integrals for the Laplace Distribution

where $\xi : [0, +\infty) \rightarrow \mathbb{R}$ is a function of r only which might change from one line to the other, and

$$\begin{aligned} \mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\left(\frac{2}{\sigma^{2(m-1)}r} - \frac{1}{\sigma^{4(m-1)}r^2} \right) X - \left(1 - \frac{1}{\sigma^{2(m-1)}r} \right) \sigma Z \right)^2 \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_1 \cup B_2}(X, Z) \right] \\ = \frac{1}{2\sqrt{2\pi}} \int_{-\sigma^{2m-1}r/2}^{\sigma/2} e^{-z^2/2} [+ \dots] dz \\ + \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2}^{\sigma^{2m-1}r/2} e^{-z^2/2} [+ \dots] dz \\ + \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r}^{-\sigma^{2m-1}r/2} e^{-z^2/2} [z^2 \sigma^2 \xi(r) + \dots] dz, \end{aligned}$$

where $\xi : [0, +\infty) \rightarrow \mathbb{R}$ is a function of r only which might change from one line to the other.

Region R_2 For ϕ_d^2 , we have

$$\begin{aligned} A_3 \cap C_1 &= \begin{cases} \sigma^2/2 - \sigma z \leq x \leq -\sigma z + \sigma^2/2 + \sigma^{2m}r/2 & \text{if } z < \sigma/2 - \sigma^{2m-1}r/2 \\ \sigma^{2m}r/2 < x \leq -\sigma z + \sigma^2/2 + \sigma^{2m}r/2 & \text{if } \sigma/2 - \sigma^{2m-1}r/2 \leq z \leq \sigma/2 \end{cases}, \\ A_3 \cap C_2 &= \begin{cases} \sigma^2/2 - \sigma z - \sigma^{2m}r/2 \leq x \leq \sigma^2/2 - \sigma z & \text{if } z < \sigma/2 - \sigma^{2m-1}r \\ \sigma^{2m}r/2 < x \leq \sigma^2/2 - \sigma z & \text{if } \sigma/2 - \sigma^{2m-1}r < z < \sigma/2 - \sigma^{2m-1}r/2 \end{cases} \end{aligned}$$

and

$$A_3 \cap (C_1 \cup C_2) = \begin{cases} \sigma^{2m}r/2 \leq x \leq \sigma^{2m}r/2 + \sigma^2/2 - \sigma z & \text{if } \sigma/2 - \sigma^{2m-1}r \leq z \leq \sigma/2 \\ -\sigma^{2m}r/2 + \sigma^2/2 - \sigma z < x \leq \sigma^{2m}r/2 + \sigma^2/2 - \sigma z \\ \text{if } z < \sigma/2 - \sigma^{2m-1}r \end{cases}.$$

The corresponding expectations are

$$\begin{aligned} \mathbb{E} \left[\left(\frac{\sigma^2}{2} - \sigma Z \right) \mathbb{1}_{A_3}(X) \mathbb{1}_{C_1}(X, Z) \right] &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \left[-\frac{rz}{2} \sigma^{2m+1} + \dots \right] dz \\ &\quad + \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r/2}^{\sigma/2} e^{-z^2/2} [z^2 \sigma^2 + \dots] dz, \\ \mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z \right) \mathbb{1}_{A_3}(X) \mathbb{1}_{C_2}(X, Z) \right] &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \left[-\frac{rz}{2} \sigma^{2m+1} + \dots \right] dz \\ &\quad + \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \left[-\frac{rz}{2} \sigma^{2m+1} + \dots \right] dz, \end{aligned}$$

and

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X + \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \left(\frac{\sigma^2}{2} - \sigma Z \right) \right) \right)^2 \right) \mathbb{1}_{A_3}(X) \mathbb{1}_{C_1 \cup C_2}(X, Z) \right] \\ &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \left[\frac{rz}{2} \sigma^{2m+1} + \dots \right] dz \\ &+ \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r}^{\sigma/2} e^{-z^2/2} \left[-\frac{z^3}{2r} \sigma^{3-2m} + \dots \right] dz. \end{aligned}$$

Region R_3 For ϕ_d^3 , we have, in the case $r \geq \sigma^{-2(m-1)}$,

$$\begin{aligned} A_1 \cap B_3 &= \begin{cases} 0 \leq x \leq \frac{\sigma^{2m}r}{2} & \text{if } z > \frac{\sigma^{2m-1}r}{2} \\ \left(\frac{\sigma^{2m}r}{2} - \sigma z \right) \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right)^{-1} < x \leq \frac{\sigma^{2m}r}{2} & \text{if } \frac{\sigma}{2} \leq z \leq \frac{\sigma^{2m-1}r}{2} \end{cases}, \\ A_1 \cap B_4 &= \begin{cases} 0 \leq x \leq \frac{\sigma^{2m}r}{2} & \text{if } z < \frac{\sigma}{2} - \sigma^{2m-1}r \\ 0 \leq x < \left(\frac{\sigma^{2m}r}{2} - \sigma z \right) \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right)^{-1} & \text{if } \frac{\sigma}{2} - \sigma^{2m-1}r \leq z \leq -\frac{\sigma^{2m-1}r}{2} \end{cases}. \end{aligned}$$

The corresponding expectations are

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{\sigma^2}{2} \right)^2 \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_3}(X, Z) \right] \\ &= \frac{1}{2\sqrt{2\pi}} \int_{\sigma^{2m-1}r/2}^{+\infty} e^{-z^2/2} \left[-\frac{rz}{8} \sigma^{2m+1} + \dots \right] dz \\ &+ \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2}^{\sigma^{2m-1}r/2} e^{-z^2/2} \left[z^{3-2m} \sigma^{3-2m} \xi(r) + \dots \right] dz, \\ & \mathbb{E} \left[\left(2X - \frac{1}{\sigma^{2(m-1)r}} X + \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z - \frac{\sigma^2}{2} \right)^2 \right) \mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z) \right] \\ &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \left[\frac{3}{8} \sigma^{2m+1} + \dots \right] dz \\ &+ \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r/2}^{-\sigma^{2m-1}r/2} e^{-z^2/2} \left[z^3 \sigma^{3-2m} \xi(r) + \dots \right] dz, \end{aligned}$$

where $\xi : [0, +\infty) \rightarrow \mathbb{R}$ is a function of r only which might change from one line to the other.

Region R_4 Finally, for ϕ_d^4 we have

$$A_3 \cap C_4 = \left\{ z < \frac{\sigma}{2} - \sigma^{2m-1}r, \frac{\sigma^{2m}r}{2} < x \leq \frac{\sigma^2}{2} - \sigma z - \frac{\sigma^{2m}r}{2} \right\},$$

and

$$\mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z \right) \mathbb{1}_{A_3}(X) \mathbb{1}_{C_4}(X, Z) \right] = \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \left[\frac{z^3}{6} \sigma^3 + \dots \right] dz.$$

4.10.4.2 Second Moment

For simplicity, we only consider the case for $r \geq \sigma^{-2(m-1)}$, the other case follows analogously.

Region R_1 For ϕ_d^1 we have

$$\begin{aligned} \mathbb{E} \left[\phi_d^1(X, Z)^2 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_1}(X, Z) \right] &= \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2}^{\sigma^{2m-1}r/2} e^{-z^2/2} \left[z^3 \sigma^3 \xi(r) + \dots \right] dz \\ &\quad + \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - r\sigma^{2m-1}/2}^0 e^{-z^2/2} \left[z^3 \sigma^3 \xi(r) + \dots \right] dz \\ &\quad + o(1), \\ \mathbb{E} \left[\phi_d^1(X, Z)^2 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_2}(X, Z) \right] &= \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r/2}^0 e^{-z^2/2} \left[z^3 \sigma^3 \xi(r) + \dots \right] dz \\ &\quad + \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r}^{-\sigma^{2m-1}r/2} e^{-z^2/2} \left[z^3 \sigma^3 \xi(r) + \dots \right] dz \\ &\quad + o(1), \end{aligned}$$

where $\xi : [0, +\infty) \rightarrow \mathbb{R}$ is a function of r only which might change from one line to the other.

Region R_2 For ϕ_d^2 we have

$$\begin{aligned} \mathbb{E} \left[\phi_d^2(X, Z)^2 \mathbb{1}_{A_3}(X) \mathbb{1}_{C_1}(X, Z) \right] &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \left[\frac{rz^2}{2} \sigma^{2m+2} + \dots \right] dz \\ &\quad + \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r/2}^{\sigma/2} e^{-z^2/2} \left[-z^3 \sigma^3 + \dots \right] dz, \\ \mathbb{E} \left[\phi_d^2(X, Z)^2 \mathbb{1}_{A_3}(X) \mathbb{1}_{C_2}(X, Z) \right] &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \left[z^2 \sigma^{2m+2} + \dots \right] dz \\ &\quad + \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \left[-\frac{rz^2}{2} \sigma^{2m+2} + \dots \right] dz. \end{aligned}$$

Region R_3 For ϕ_d^3 we have

$$\begin{aligned}
 & \mathbb{E} \left[\left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{\sigma^2}{2} \right)^2 \right)^2 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_3}(X, Z) \right] \\
 &= \frac{1}{2\sqrt{2\pi}} \int_{\sigma^{2m-1}r/2}^{+\infty} e^{-z^2/2} \left[\frac{rz^2}{24} \sigma^{2m+2} + \dots \right] dz \\
 &+ \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2}^{\sigma^{2m-1}r/2} e^{-z^2/2} \left[\frac{rz^2}{24} \sigma^{2m+2} + \dots \right] dz, \\
 & \mathbb{E} \left[\left(2X - \frac{1}{\sigma^{2(m-1)r}} X + \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z - \frac{\sigma^2}{2} \right)^2 \right)^2 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z) \right] \\
 &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \left[\frac{7rz^2}{24} \sigma^{2m+2} + \dots \right] dz \\
 &+ \frac{1}{2\sqrt{2\pi}} \int_{\sigma/2 - \sigma^{2m-1}r/2}^{-\sigma^{2m-1}r/2} e^{-z^2/2} \left[\frac{7rz^2}{24} \sigma^{2m+2} + \dots \right] dz,
 \end{aligned}$$

Region R_4 For ϕ_d^4 we have

$$\mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z \right)^2 \mathbb{1}_{A_3}(X) \mathbb{1}_{C_4}(X, Z) \right] = \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \left[-\frac{z^3}{3} \sigma^3 + \dots \right] dz.$$

4.10.4.3 Third Moment

Having established that the only possible scaling is given by $\alpha = 1/3$, $\beta = m/3$ with $m \geq 1$, we now proceed to bound the third moment of ϕ_d in this case. For simplicity, we only consider the case for $r \geq 1$, the other case follows analogously.

Since $m \geq 1$, we find that $\mathbb{E} [\phi_d^1(X, Z)^3 \mathbb{1}_{R_1}(X, Z)] = o(1)$ as $d \rightarrow \infty$ since the limits of

integration all converge to 0. Then, using Hölder's inequality for ϕ_d^2 , we have

$$\begin{aligned}
 \mathbb{E} \left[\phi_d^2(X, Z)^3 \right] &\leq C \mathbb{E} \left[\left(\frac{\sigma^2}{2} - \sigma Z \right)^3 \mathbb{1}_{A_3}(X) \mathbb{1}_{C_1}(X, Z) \right] \\
 &\quad + C \mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z \right)^3 \mathbb{1}_{A_3}(X) \mathbb{1}_{C_2}(X, Z) \right] \\
 &\quad + C \mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X + \left(1 - \frac{1}{\sigma^{2(m-1)r}} \right) \left(\frac{\sigma^2}{2} - \sigma Z \right) \right) \right)^2 \right)^3 \right. \\
 &\quad \quad \left. \times \mathbb{1}_{A_3}(X) \mathbb{1}_{C_1 \cup C_2}(X, Z) \right] \\
 &= \frac{C}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r/2} e^{-z^2/2} \left[-\frac{rz^3}{2} \sigma^5 + \dots \right] dz \\
 &\quad + \frac{C}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \left[-\frac{rz^3}{2} \sigma^5 + \dots \right] dz \\
 &\quad + \frac{C}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \left[z^3 \sigma^5 \xi(r) + \dots \right] dz + o(1),
 \end{aligned}$$

where $\xi : [0, +\infty) \rightarrow \mathbb{R}$ is a function of r only which might change from one line to the

other. For ϕ_d^3 , we have, using again Hölder's inequality,

$$\begin{aligned}
 & \mathbb{E} \left[\left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{\sigma^2}{2} \right)^2 \right)^3 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_3}(X, Z) \right] \\
 & \leq C \mathbb{E} \left[\left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z \right)^3 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_3}(X, Z) \right] \\
 & + C \mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z + \frac{\sigma^2}{2} \right)^2 \right)^3 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_3}(X, Z) \right] \\
 & = \frac{1}{2\sqrt{2\pi}} \int_{\sigma^{2m-1}r/2}^{+\infty} e^{-z^2/2} \left[-\frac{z^3 r}{2} \sigma^5 + \dots \right] dz \\
 & + \frac{1}{2\sqrt{2\pi}} \int_{\sigma^{2m-1}r/2}^{+\infty} e^{-z^2/2} \left[z^3 \sigma^5 \xi(r) + \dots \right] dz + o(1) \\
 & \mathbb{E} \left[\left(2X - \frac{1}{\sigma^{2(m-1)r}} X + \sigma Z + \frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z - \frac{\sigma^2}{2} \right)^2 \right)^3 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z) \right] \\
 & \leq C \mathbb{E} \left[\left(2X - \frac{1}{\sigma^{2(m-1)r}} X + \sigma Z \right)^3 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z) \right] \\
 & + C \mathbb{E} \left[\left(\frac{Z^2}{2} - \frac{1}{2\sigma^2} \left(\frac{1}{\sigma^{2(m-1)r}} X - \sigma Z - \frac{\sigma^2}{2} \right)^2 \right)^3 \mathbb{1}_{A_1}(X) \mathbb{1}_{B_4}(X, Z) \right] \\
 & = \frac{1}{2\sqrt{2\pi}} \int_{\sigma^{2m-1}r/2}^{+\infty} e^{-z^2/2} \left[\frac{z^3 r}{2} \sigma^5 + \dots \right] dz \\
 & + \frac{1}{2\sqrt{2\pi}} \int_{\sigma^{2m-1}r/2}^{+\infty} e^{-z^2/2} \left[z^3 \sigma^5 \xi(r) + \dots \right] dz + o(1),
 \end{aligned}$$

where $\xi : [0, +\infty) \rightarrow \mathbb{R}$ is a function of r only which might change from one line to the other. Finally, for ϕ_d^4 we have

$$\mathbb{E} \left[\left(2X - \frac{\sigma^2}{2} + \sigma Z \right)^3 \mathbb{1}_{A_3}(X) \mathbb{1}_{C_4}(X, Z) \right] = \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\sigma/2 - \sigma^{2m-1}r} e^{-z^2/2} \left[\frac{z^3}{10} \sigma^5 + \dots \right] dz.$$

APPENDIX A

Introduction (en français)

A.1 Introduction à l'approximation stochastique

L'approximation stochastique a été inventée en 1951 lors de la parution de l'article [RM51] dans la revue *Annals of Mathematical Statistics*. Il s'agit d'une famille de méthodes itératives pour trouver les racines d'une fonction $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$ appelée champ moyen, lorsque celle-ci n'est accessible qu'avec un bruit. Le schéma de la méthode est le suivant:

$$\theta_{n+1} = \theta_n + \eta_{n+1} H_{\theta_n}(X_{n+1}), \quad (\text{A.1})$$

où $(\theta_n)_{n \in \mathbb{N}}$ est la suite créée par itération, avec $\theta_n \in \mathbb{R}^d$ pour tout $n \in \mathbb{N}$ et qui est initialisée en θ_0 . La suite $(\eta_n)_{n \in \mathbb{N}^*}$ constitue la taille des pas, qui est déterministe et choisie par l'utilisateur sous la condition $\eta_n > 0$ pour tout $n \in \mathbb{N}^*$. Enfin, $H_{\theta_n}(X_{n+1})$ est un estimateur de la fonction continue h évaluée en θ_n . Le bruit est modélisé par une suite de variables aléatoires $(X_{n+1})_{n \in \mathbb{N}}$ à valeurs dans l'espace probabilisable $(\mathcal{X}, \mathcal{X})$.

Le schéma (A.1) a été étudié depuis des décennies, surtout par les communautés d'approximation stochastique (voir [BMP90; Bor09; BCZ11; KY03; PJ92; Rup88]) et d'optimisation (voir [dST+21; Nem+09; Nes04; NV08]). Pendant les vingt dernières années, la communauté d'apprentissage statistique a trouvé d'intéressantes applications, au cœur de nombreuses réussites du domaine (voir [Bot10; Bot12; BB08b; BCN18; Cra+06; DHS11; Sha+12; SSS07; Xia09]). Aussi bien en optimisation qu'en apprentissage statistique, la principale application du schéma (A.1) est la descente de gradient stochastique, qui s'obtient lorsque $h = -\text{grad } f$ pour une fonction différentiable $f : \mathbb{R}^d \rightarrow \mathbb{R}$.

Présentons quelques exemples d'utilisation de l'approximation stochastique.

A.1.1 Motivation

Example A.1 (Estimation de la moyenne). *Pour $d \in \mathbb{N}^*$, on considère une mesure de probabilité π sur \mathbb{R}^d avec un second moment fini. On se donne une suite i.i.d. d'observations $(X_n)_{n \in \mathbb{N}^*}$ de loi π . On peut estimer l'inconnue $\mu = \int_{\mathbb{R}^d} u \, d\pi(u)$ par l'approximation*

stochastique. Soit \mathcal{L}_2 le risque quadratique défini pour tout $\theta \in \mathbb{R}^d$ par

$$\mathcal{L}_2(\theta) = \frac{1}{2} \int_{\mathbb{R}^d} \|\theta - u\|^2 d\pi(u).$$

En dérivant sous le signe intégral, on obtient que \mathcal{L}_2 est différentiable et que son gradient est donné par $\text{grad } \mathcal{L}_2(\theta) = \theta - \mu$. Comme nous n'avons pas accès à μ , on ne peut pas évaluer $\text{grad } \mathcal{L}_2$, cependant on peut utiliser les observations $(X_n)_{n \in \mathbb{N}}$ comme estimateurs de μ pour construire un estimateur simple mais peu efficace de $-\text{grad } \mathcal{L}_2(\theta)$. Pour tout $n \in \mathbb{N}^*$ et tout $\theta \in \mathbb{R}^d$, soit $H_\theta(X_n) = -(\theta - X_n)$. Résumons les propriétés de cette estimation :

$$\mathbb{E}[H_\theta(X_n)] = -\text{grad } \mathcal{L}_2(\theta), \quad \text{and} \quad \mathbb{E}[\|H_\theta(X_n) + \text{grad } \mathcal{L}_2(\theta)\|^2] = \sigma^2 < +\infty,$$

où $\sigma^2 = \text{Var}(X_1)$. On utilise ensuite cet estimateur dans l'algorithme de Robbins-Monro pour effectuer la descente de gradient stochastique donnée par (A.1), afin de minimiser le risque \mathcal{L}_2 , en initialisant avec $\theta_0 = 0$ et un choix de pas $\eta_n = 1/n$ pour tout $n \in \mathbb{N}$:

$$\theta_{n+1} = \theta_n + \eta_{n+1} H_{\theta_n}(X_{n+1}) = \theta_n - \frac{1}{n+1} (\theta_n - X_{n+1}) = \frac{n}{n+1} \theta_n + \frac{1}{n+1} X_{n+1}.$$

Une simple récurrence permet de voir que pour tout $n \in \mathbb{N}^*$, $\theta_n = n^{-1} \sum_{i=1}^n X_i$. Par la loi forte des grands nombres, $\theta_n \xrightarrow[n \rightarrow \infty]{p.s.} \mu$.

Regardons ensuite deux exemples en apprentissage supervisé où l'on utilise la descente de gradient stochastique.

Apprentissage supervisé Étant donné un ensemble de n observations i.i.d. $(x_i, y_i) \in \mathbb{X} \times \mathbb{Y}$ pour $i = 1, \dots, n$, et ayant une nouvelle observation partielle x , on veut prédire la variable inconnue y en utilisant une fonction de la forme $\theta^\top \Phi(x)$. Celle-ci est linéaire par rapport aux caractéristiques $\Phi(x) \in \mathbb{R}^d$. Supposons que l'on a déjà choisi la fonction d'extraction des caractéristiques $\Phi : \mathbb{X} \rightarrow \mathbb{R}^d$, et que l'on veut apprendre les paramètres $\theta \in \mathbb{R}^d$ tels que $\theta^\top \Phi(x)$ soit aussi proche de y que possible. Pour évaluer la performance du paramètre courant θ , on considère une fonction de perte $\ell : \mathbb{Y} \times \mathbb{R} \rightarrow \mathbb{R}_+$, de façon à avoir $\ell(y, \theta^\top \Phi(x))$ petit si et seulement si y est proche de $\theta^\top \Phi(x)$.

Un choix pour la fonction de perte est le *risque empirique*.

Exemple A.2 (Minimisation du risque empirique). Le risque empirique est donné par

$$\min_{\theta \in \mathbb{R}^d} \mathcal{R}_{emp}(\theta), \quad \text{où } \mathcal{R}_{emp} : \theta \mapsto \frac{1}{n} \sum_{i=1}^n \ell(y_i, \theta^\top \Phi(x_i)). \quad (\text{A.2})$$

Le calcul analytique du minimiseur de (A.2) est coûteux voire impossible en général (le nombre d'observations n est souvent très grand). Ainsi, on utilise un algorithme itératif pour minimiser (A.2).

À l'itération t , on tire une variable aléatoire i_{t+1} , uniforme sur $\{1, \dots, n\}$ et on considère g_{t+1} , le gradient de $\theta \mapsto \ell(y_{i_{t+1}}, \theta^\top \Phi(x_{i_{t+1}}))$. Ensuite, étant donnée une initialisation $\theta_0 \in \mathbb{R}^d$, on construit les itérations de la façon suivante :

$$\theta_{t+1} = \theta_t - \eta_{n+1} g_{t+1}(\theta_t) ,$$

où $(\eta_n)_{n \in \mathbb{N}^*}$ est la suite des pas, qui est positive. Dans cette approche, l'aléa provient du choix des indices. Comme $(i_t)_{t \in \mathbb{N}^*}$ est tirée avec remplacement, on utilisera plusieurs fois la même observation (x_i, y_i) . Sous des hypothèses assez générales, la suite $(\theta_t)_{t \in \mathbb{N}}$ converge lorsque $t \rightarrow \infty$ vers un minimum local θ^* du risque empirique (A.2).

Un autre choix de fonction de risque qui est intéressant à considérer est le *risque en espérance*.

Exemple A.3 (Minimisation du risque en espérance). On s'intéresse maintenant à la minimisation du risque en espérance, c'est-à-dire

$$\min_{\theta \in \mathbb{R}^d} \mathcal{R}_{exp}(\theta) , \quad \text{où } \mathcal{R}_{exp} : \theta \mapsto \mathbb{E} \left[\ell \left(y, \theta^\top \Phi(x) \right) \right] ,$$

où (x, y) est une variable aléatoire de même distribution que les observations $(x_i, y_i)_{i \in \{1, \dots, n\}}$ et indépendante de celles-ci. Cette fonction de risque prend en compte toutes les valeurs possibles de x , même celles qui n'ont pas été observées. Il s'agit ainsi d'un meilleur indicateur de la qualité de la prédiction. Néanmoins, la loi d'une observation (x, y) n'est pas connue en pratique, et le risque en espérance n'est donc pas calculable. On décide alors d'optimiser une approximation de ce risque \mathcal{R}_{exp} de façon itérative.

Comme dans l'exemple A.1, les observations nous fournissent des estimateurs simples mais inefficaces de $\text{grad } \mathcal{R}_{exp}$. Ainsi, si on note g_t le gradient de $\theta \mapsto \ell(y_t, \theta^\top \Phi(x_t))$, alors, en dérivant sous le signe intégral, $g_t(\theta)$ est un estimateur sans biais de $\text{grad } \mathcal{R}_{exp}(\theta)$. On obtient le schéma suivant, initialisé en $\theta_0 \in \mathbb{R}^d$,

$$\theta_{t+1} = \theta_t - \eta_{n+1} g_{t+1}(\theta_t) , \tag{A.3}$$

où t est à valeurs dans $\{0, \dots, n-1\}$ et $(\eta_{n+1})_{n \in \mathbb{N}}$ est encore la suite des pas. Ici, l'aléa provient des observations elles mêmes. Pour préserver l'indépendance des estimateurs $(g_t)_{t \in \{1, \dots, n\}}$, on décide de ne faire qu'une passe sur les données. Si le nombre d'observations n tend vers l'infini, alors le schéma (A.3) converge vers θ^* , un minimum du risque en espérance \mathcal{R}_{exp} .

On peut noter que le risque Exemple A.2 est bien adapté lorsque toutes les données sont accessibles depuis le début, *i.e.* lorsque n est fixé. On dit que ces données sont statiques. Cependant Exemple A.3 est mieux adapté pour des données itératives, où une seule observation (x_t, y_t) est donnée à chaque instant $t \in \mathbb{N}$. On appelle cette différence *batch learning* contre *online learning*.

L'exemple suivant montre comment transformer un algorithme adapté pour des données statiques en un algorithme fait pour les données *en ligne* ou itératives.

L'algorithme Espérance-Maximisation (EM) On introduit d'abord l'algorithme EM (voir [DLR77]) dans le contexte d'un modèle de variable latentes. Ceci modélise aussi bien des cas de données manquantes, d'observations censurées ou encore de données bruitées.

Supposons qu'il existe une variable aléatoire non observable, dite *latente* X qui suit la loi de densité $f(x; \theta)$, où $\{f(x; \theta) : \theta \in \Theta\}$ est une famille paramétrique de densités de probabilité, indexées par le paramètre $\theta \in \Theta \subset \mathbb{R}^d$. On se donne accès à une observation $Y \in \mathcal{Y}$ qui est une fonction déterministe de X . On modélise cette dépendance en disant que $y \mapsto g(y; \theta)$ est la vraisemblance de la variable observée du modèle.

Avec n observations i.i.d. $Y = (Y_1, \dots, Y_n)$, l'algorithme EM maximise de façon itérative la log-vraisemblance $\log g(Y; \theta) = \sum_{i=1}^n \log g(Y_i; \theta)$, malgré la possible formulation compliquée de g , qui résulte du modèle à variables latentes. Chaque itération k consiste en deux étapes. L'étape d'espérance, dite étape E consiste à évaluer l'espérance conditionnelle

$$\theta \mapsto Q_{\theta_k}(Y; \theta) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\theta_k} [\log f(X_i; \theta) | Y_i] ,$$

où θ_k est l'estimateur courant de θ et $\mathbb{E}_{\theta}[\cdot | Y]$ est l'espérance conditionnelle sachant Y pour le modèle paramétré par θ . L'étape de maximisation, dite étape M, consiste à trouver la valeur de θ qui maximise $Q_{\theta_k}(Y; \theta)$. Cette valeur est prise comme le nouvel estimateur θ_{k+1} . Le succès de cet algorithme vient du fait que si $Q_{\theta_k}(Y; \theta)$ croît, alors la log-vraisemblance $\log g(Y; \theta)$ croît aussi (voir [DLR77]).

L'algorithme EM utilise toutes les données disponibles à chaque itération. Il s'agit d'un algorithme *batch*. Il existe plusieurs versions en ligne de l'algorithme EM, nous en présentons celle de [CM09].

Exemple A.4 (Algorithme EM en ligne). *En conservant les notations de l'algorithme EM, l'étape E de la version en ligne devient: à l'itération k , pour $\theta \in \Theta$, évaluer la fonction*

$$\hat{Q}_{k+1}(\theta) = \hat{Q}_k(\theta) + \eta_{k+1} \left(\mathbb{E}_{\theta_k} [\log f(X_{k+1}; \theta) | Y_{k+1}] - \hat{Q}_k(\theta) \right) ,$$

où $(\eta_n)_{n \in \mathbb{N}^*}$ est une suite positive de pas. L'étape M reste inchangée, il s'agit de définir θ_{n+1} comme la valeur dans Θ qui maximise \hat{Q}_{n+1} .

Dans [CM09], les auteurs prouvent que l'algorithme converge à une valeur stationnaire de la divergence de Kullback-Leibler évaluée contre la solution (voir Théorème 1 de [CM09]). Cet algorithme présente aussi une propriété bien utile, qui est que la divergence de Kullback-Leibler évaluée contre la solution décroît à chaque itération, tant que le pas est suffisamment petit (voir proposition 2 de [CM09]).

Ayant présenté différents problèmes que l'approximation stochastique est capable de résoudre, on s'intéresse maintenant aux hypothèses et aux techniques de preuves de cette famille de méthodes.

A.1.2 Hypothèses sur le bruit

L'estimateur $H_{\theta_n}(X_{n+1})$ se réécrit comme suit,

$$H_{\theta_n}(X_{n+1}) = h(\theta_n) + e_{\theta_n}(X_{n+1}) ,$$

où $e_{\theta_n}(X_{n+1})$ est appelée l'erreur d'estimation. On définit la filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$ comme la filtration naturelle associée à la suite $(\theta_n)_{n \in \mathbb{N}}$.

L'hypothèse la plus classique à faire sur le bruit (voir l'hypothèse A4 de [Kar+19], l'hypothèse A1 de [GL13] ou encore l'hypothèse H1 de [MB11]), est de borner les deux premiers moments du bruit. En voici un exemple.

A 22 (Cas différence de martingale (euclidien)). *Il existe $\sigma^2 > 0$ tel que pour tout $n \in \mathbb{N}$,*

$$\mathbb{E}[e_{\theta_n}(X_{n+1})|\mathcal{F}_n] = 0 , \text{ et } \mathbb{E}[\|e_{\theta_n}(X_{n+1})\|^2|\mathcal{F}_n] \leq \sigma^2 .$$

Cette hypothèse sur le bruit veut dire que l'estimateur $H_{\theta}(X)$ est sans biais et que sa variance est uniformément bornée. Ce cas est appelé *différence de martingale* car la suite de variables aléatoires $(\sum_{k=0}^n e_{\theta_k}(X_{k+1}))_{n \in \mathbb{N}}$ est une martingale par rapport à la filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$. Pour généraliser cette hypothèse, on peut remplacer l'inégalité sur la variance par $\mathbb{E}[\|e_{\theta_n}(X_{n+1})\|^2|\mathcal{F}_n] \leq \sigma_0^2 + \sigma_1^2 \|h(\theta_n)\|^2$, ce qui veut dire que la borne sur la variance peut dépendre du point courant θ_n .

L'autre type d'hypothèses sur le bruit qui a intéressé la communauté (voir Cas 2 dans [Kar+19] et [BCN18; GL13]) est le cas markovien, où la suite des vecteur de bruit $(e_{\theta_n}(X_{n+1}))_{n \in \mathbb{N}}$ est une chaîne de Markov qui dépend de l'état courant θ_n .

A 23 (Cas markovien (euclidien)). *Il existe un noyau de Markov P sur $(\Theta \times \mathbb{X}) \times \mathcal{X}$ tel que pour tout $n \in \mathbb{N}$ et toute fonction mesurable $\varphi : \mathbb{X} \rightarrow \mathbb{R}$, $\mathbb{E}[\varphi(X_{n+1})|\mathcal{F}_n] = \int_{\mathbb{X}} \varphi(y)P_{\theta_n}(X_n, dy)$. De plus, pour tout $\theta \in \Theta$, P_{θ} admet une unique mesure de probabilité invariante π_{θ} , qui de plus satisfait $h(\theta) = \int_{\mathbb{X}} H_{\theta}(y)d\pi_{\theta}(y)$.*

Ce cadre est important dans des sujets d'apprentissage par renforcement comme le Q-learning, (voir [JJS93]) et les méthodes policy-gradient (voir [BB01]).

A.1.3 Techniques de preuve pour la convergence

Une approche pour analyser le schéma (A.1) qui a porté ses fruits consiste à trouver une fonction positive $V : \mathbb{R}^d \rightarrow \mathbb{R}_+$ qui décroît lors de chaque itération. Cette fonction V , qui peut s'interpréter comme une sorte d'énergie potentielle est appelée fonction de Lyapunov, tout comme les fonctions de Lyapunov utilisées dans l'analyse d'équations différentielles (voir [Lya92]). La terminologie provient de la deuxième méthode de Lyapunov, décrite dans [KB60], où les auteurs utilisent cette technique pour démontrer la stabilité asymptotique globale des solutions du schéma (A.1). Ils y procèdent en montrant que toutes les trajectoires des solutions de l'équation différentielle ordinaire (EDO) obtenue comme limite, $\dot{\theta} = h(\theta)$ convergent vers zéro. Cette technique de preuve, nommée l'approche par EDO, a été reprise avec beaucoup de succès dans [BFH13; Bor09; BM00; Kar+19]. La

principale contrainte de celle-ci est la connaissance d'une fonction de Lyapunov adaptée au problème, puisqu'il n'y a pas de fonction de Lyapunov générique pour tous les cas. La plupart des analyses sont de nature asymptotique (voir [Fab68; Pf86; Sac58]), mais la littérature non-asymptotique est de plus en plus prolifique dernièrement (voir [BM11; Kar+19; Pf90; RRT17]), puisque les résultats non-asymptotiques ont des implications palpables à toutes les étapes de l'algorithme.

A.1.4 Hypothèses pour la descente de gradient stochastique

Dans le cas de la descente de gradient stochastique, les conditions de Lyapunov reviennent à vérifier différentes hypothèses de convexité. Il y a deux cas qui méritent d'être soulignés, le cas fortement convexe et le cas convexe. Si f est fortement convexe, sous des hypothèses assez faibles et si pour tout $n \in \mathbb{N}^*$, $\eta_n = c/n$ pour un certain $c > 0$ fixé, alors on connaît le taux de convergence non-asymptotique suivant

$$\mathbb{E} \left[f(\bar{\theta}_n) - f(\theta^*) \right] \leq C/n ,$$

où $C > 0$ est une constante, $\bar{\theta}_n = (1/n) \sum_{k=0}^{n-1} \theta_k$ et θ^* est l'unique solution du problème (voir [BM11]). Ceci correspond à la meilleure vitesse de convergence possible pour cette classe de problèmes (voir [Aga+09]). Lorsque f est uniquement convexe, sous des hypothèses similaires à celles du cas fortement convexe, et pour le même choix de pas, on obtient

$$\mathbb{E} \left[f(\bar{\theta}_n) - f(\theta^*) \right] \leq C/\sqrt{n} .$$

Cette borne est aussi optimale pour cette classe de problèmes (voir [Aga+09]).

A.1.5 Schéma projeté

Il existe de multiples raisons pour souhaiter que les itérations de la suite (A.1) restent dans un ensemble $\Theta \subset \mathbb{R}^d$ plus petit que tout l'espace disponible \mathbb{R}^d : c'est utilisé pour accélérer la convergence dans les SVM projetés [LSB12; SSS07], ou encore pour démontrer la stabilité de l'algorithme (on a souvent besoin que les itérés restent dans un ensemble compact) comme dans [MMM18] et parfois on projette pour s'assurer que $(\theta_n)_{n \in \mathbb{N}}$ satisfait un ensemble de contraintes S comme dans [Pf86]. Pour tous ces cas, on réécrit (A.1) de la façon suivante,

$$\theta_{n+1} = \text{proj}_{\Theta} (\theta_n + \eta_{n+1} H_{\theta_n}(X_{n+1})) , \quad (\text{A.4})$$

où proj_{Θ} est la projection orthogonale sur l'ensemble convexe et compact $\Theta \subset \mathbb{R}^d$.

Regardons l'exemple suivant pour motiver l'étude du schéma projeté.

Exemple A.5 (Estimation de vecteurs propres). *On observe une suite de matrices aléatoires, indépendantes et de taille $d \times d$ $(A_k)_{k \in \mathbb{N}}$, qui ont toutes la même espérance $\mathbb{E}[A_k] = A$, où A est une matrice symétrique. On cherche à approximer un vecteur propre de la matrice moyenne A en utilisant les observations bruitées. L'algorithme proposé par [OK85] est le suivant,*

$$\theta_n = \frac{\theta_{n-1} + \eta_n A_n \theta_{n-1}}{\|\theta_{n-1} + \eta_n A_n \theta_{n-1}\|} , \quad (\text{A.5})$$

où $\|\cdot\|$ est la norme euclidienne. Supposons que $\sum_{i \in \mathbb{N}^*} \eta_n = \infty$, $\sum_{i \in \mathbb{N}^*} \eta_n^2 < \infty$ et que l'hypothèse **A22** est satisfaite – ce qui implique que $\sup_{n \in \mathbb{N}} \mathbb{E}[\|A_n\|^2] < \infty$ – alors $(\theta_n)_{n \in \mathbb{N}}$ converge dans un espace propre de A (voir théorème 3 de [Del00]). Si A n'a pas de valeur propre multiple, alors la suite définie plus haut converge vers un vecteur propre de A .

Il est important de remarquer que l'itération (A.5) peut se décomposer en deux étapes : d'abord $\tilde{\theta}_n = \theta_{n-1} + \eta_n A_n \theta_{n-1}$ suivi de $\theta_n = \tilde{\theta}_n / \|\tilde{\theta}_n\| = \text{proj}_{\mathbb{S}^{d-1}}(\tilde{\theta}_n)$, où \mathbb{S}^{d-1} est la sphère unité (de dimension $d - 1$ plongée dans \mathbb{R}^d) pour la norme euclidienne $\|\cdot\|$. La première étape correspond à un schéma de la forme (A.1). Cependant, ne faire que cette étape construirait une suite divergente dans la plupart des cas à cause de $\|\theta_n\| \rightarrow \infty$ – e.g. dans le cas d'une valeur propre positive ou d'un mauvais choix de pas $(\eta_n)_{n \in \mathbb{N}^*}$. Puisque la norme du vecteur n'est pas importante lorsque l'on cherche un vecteur propre, la deuxième étape enlève un degré de liberté superflu en projetant les itérés sur la sphère unité.

Dans ce cas, la projection permet de résoudre l'indétermination de $(\|\theta_n\|)_{n \in \mathbb{N}}$. Néanmoins, puisque $\theta_n - \theta_{n-1}$ n'est plus proportionnel à $A_n \theta_{n-1}$, on ne peut pas savoir si le prochain pas de l'algorithme va être grand ou petit en regardant uniquement $A_n \theta_{n-1}$. On pense qu'il existe de meilleures façons de résoudre l'indétermination sur la norme, en construisant les itérations du type (A.1) dans un sous-espace contraint Θ , dans lequel la distance entre deux itérés consécutifs correspond à la norme du champ moyen. Ce nouveau schéma doit aussi répondre aux nécessités prises en compte par le schéma projeté.

Pour cela, on restreint l'espace des paramètres \mathbb{R}^d à une variété différentiable Θ , que l'on munit d'une métrique riemannienne g .

A.2 Approximation stochastique sur une variété riemannienne

Sur une variété riemannienne, on adapte le schéma itératif (A.1) de façon à ce qu'il reste sur la variété Θ de façon intrinsèque, autrement dit passer par \mathbb{R}^d ni par une étape de projection :

$$\theta_{n+1} = \text{Exp}_{\theta_n}(\eta_{n+1} H_{\theta_n}(X_{n+1})) , \tag{A.6}$$

où $\text{Exp}_\theta : T_\theta \rightarrow \Theta$ est l'exponentielle riemannienne, qui correspond au flot de l'équation différentielle qui définit le chemin le *plus droit* sur une variété riemannienne. On peut approximativement dire que $\text{Exp}_\theta(v) \approx \theta + v$. On peut se demander si (A.6) est fondamentalement différent du schéma itératif projeté (A.4). Intuitivement pour ce nouveau schéma, le chemin qui relie θ_n et θ_{n+1} n'est pas la projection d'une droite sur \mathbb{R}^d comme dans (A.4), mais plutôt le chemin *le plus droit* qui reste dans la variété Θ . Cette différence est cruciale, et elle a pour conséquence que, sous des hypothèses assez faibles¹, la distance riemannienne entre θ_{n+1} et θ_n est exactement la norme riemannienne du vecteur dans l'exponentielle $H_{\theta_n}(X_{n+1})$ multipliée par le pas η_{n+1} .

¹c'est vrai pour des vecteurs suffisamment petits, de façon à ce qu'ils restent à l'intérieur du domaine d'injectivité $\text{TD}(\theta)$ (voir Section 2.11.6)

A.2.1 Schéma géodésique et schéma par rétraction

Le schéma (A.6) est dit géodésique car le chemin entre θ_n et θ_{n+1} est une géodésique. Il présente de nombreux avantages comparé au schéma projeté (A.4), comme l'appartenance intrinsèque de la suite $(\theta_n)_{n \in \mathbb{N}}$ à la variété Θ et la propriété que la distance riemannienne entre deux itérés consécutifs $\rho_\Theta(\theta_n, \theta_{n+1})$ soit égale au pas η_{n+1} multiplié par la norme riemannienne du champ moyen bruité $\|H_{\theta_n}(X_{n+1})\|_{\theta_n}$, comme mentionné ci-dessus.

Cependant, ces bénéfices entraînent des coûts, notamment le calcul de l'exponentielle riemannienne, ou de façon équivalente, le calcul d'une géodésique $\gamma : [0, 1] \rightarrow \Theta$ partant de $\gamma(0) = \theta \in \Theta$ et de vitesse initiale $\dot{\gamma}(0) = u \in T_\theta\Theta$. Dans le meilleur des cas, la géodésique a une forme close qui est simple à calculer effectivement, comme dans ce premier cas.

Exemple A.6 (Géodésiques sur la sphère S^d). Soit $S^d = \{\mathbf{x} \in \mathbb{R}^{d+1} : \|\mathbf{x}\| = 1\}$ la sphère unité, où $\|\cdot\|$ représente la norme euclidienne standard sur \mathbb{R}^{d+1} . La structure classique de variété riemannienne de la sphère est la suivante. Pour tout $\theta \in S^d$, l'espace tangent à un point $\theta \in S^d$ est donné par $T_\theta S^d = \{\mathbf{u} \in \mathbb{R}^{d+1} : \mathbf{u}^\top \theta = 0\}$ – voir [AMS08, Exemple 3.5.1]. La métrique canonique sur la sphère g est définie comme la restriction du produit scalaire euclidien dans \mathbb{R}^{d+1} sur l'espace tangent $T_\theta S^d$. La formule de l'exponentielle riemannienne est donnée dans [Lee19, Proposition 5.27],

$$\text{Exp}_\theta(\mathbf{u}) = \cos(\|\mathbf{u}\|)\theta + \sin(\|\mathbf{u}\|)(\mathbf{u}/\|\mathbf{u}\|) .$$

Dans d'autres cas, la géodésique a une forme close mais elle est bien plus coûteuse à calculer, comme dans cet exemple.

Exemple A.7 (Géodésiques sur la variété de Grassmann $\text{Gr}_r(\mathbb{R}^d)$). Soit $r \in \{1, \dots, d\}$ et soit $\text{Gr}_r(\mathbb{R}^d)$ l'ensemble de sous-espaces de \mathbb{R}^d de dimension r . On peut représenter $\text{Gr}_r(\mathbb{R}^d)$ comme le quotient de la variété de Stiefel non-compacte par l'action du groupe linéaire $\text{GL}_r(\mathbb{R})$. Cette représentation correspond à dire qu'un sous-espace de dimension r de \mathbb{R}^d , noté E , se représente par n'importe quelle famille de r vecteurs linéairement indépendants $\mathbf{v}_1, \dots, \mathbf{v}_r$ de \mathbb{R}^d tels que $\text{Span}(\mathbf{v}_1, \dots, \mathbf{v}_r) = E$. Ensuite, prendre le quotient permet d'identifier les familles de vecteurs qui engendrent le même sous-espace. On effectue une description plus détaillée des variétés de Stiefel et de Grassmann dans Exemple 2.9 et Section 2.7.2. Pour \mathbf{B} dans la variété de Stiefel, on note $[\mathbf{B}]$ la classe qui lui correspond dans $\text{Gr}_r(\mathbb{R}^d)$. Pour tout $[\mathbf{B}] \in \text{Gr}_r(\mathbb{R}^d)$ et $\mathbf{D} \in T_{[\mathbf{B}]} \text{Gr}_r(\mathbb{R}^d)$, une formule tirée de [AMS04, Theorem 3.6] nous donne l'expression de l'exponentielle riemannienne

$$\text{Exp}_{[\mathbf{B}]}(\mathbf{D}) = \left[\left(\mathbf{B} (\mathbf{B}^\top \mathbf{B})^{-1/2} \mathbf{V} \quad \mathbf{U} \right) \begin{pmatrix} \cos(\boldsymbol{\Sigma}) \\ \sin(\boldsymbol{\Sigma}) \end{pmatrix} \right] ,$$

où \cos et \sin sont donnés par leurs définitions en série entières et $\mathbf{D}(\mathbf{B}^\top \mathbf{B})^{-1/2} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ la décomposition en valeurs singulières compacte (que l'on appelle SVD par la suite, voir [HP14, Theorem 1.46]), avec \mathbf{U} une matrice de taille $d \times r$ avec des colonnes orthogonales, i.e. $\mathbf{U}^\top \mathbf{U} = \mathbf{I}_r$, \mathbf{V} une matrice orthogonale de taille $r \times r$, et $\boldsymbol{\Sigma}$ une matrice diagonale à des coefficients positifs de taille $r \times r$.

Dans ce cas, en plus de calculs rapides comme des sommes, des multiplications et des fonctions trigonométriques, le calcul d'une géodésique exige une SVD qui coûte environ $\mathcal{O}(dr^2)$ calculs (voir [HGI07]). De tels calculs sont souvent trop coûteux en temps pour des problèmes en grande dimension, puisque la SVD doit se faire à chaque itération du schéma (A.6).

Le dernier cas est lorsqu'il n'y a pas de forme close pour les solutions de l'équation géodésique – voir Section 2.11.3. Dans ce cas, on peut approcher les solutions de cette équation numériquement, mais le schéma n'est plus exactement géodésique.

Pour se livrer du problème du calcul des géodésiques, présent dans deux des trois cas ci-dessus, la communauté s'est très vite intéressée aux schémas itératifs dits de rétraction. L'idée c'est de remplacer l'exponentielle riemannienne par une approximation convenable, plus facile à calculer en pratique. De telles approximations de l'exponentielle riemannienne sont appelées *rétractions*.

Definition A.8 (Rétraction). *Une fonction continument dérivable $\text{Ret} : T\Theta \rightarrow \Theta$, est dite rétraction lorsque pour tout $\theta \in \Theta$, $\text{Ret}_\theta(0_\theta) = \theta$ et $(D\text{Ret})_\theta(0_\theta) = \text{Id}$, où 0_θ est le zéro de l'espace vectoriel $T_\theta\Theta$.*

Le schéma itératif qui en résulte est le suivant. En commençant en $\theta_0 \in \Theta$, on définit récursivement pour tout $n \in \mathbb{N}^*$,

$$\theta_{n+1} = \text{Ret}_{\theta_n}(\eta_{n+1}H_{\theta_n}(X_{n+1})) . \tag{A.7}$$

Pour des exemples de rétractions, voir Exemples 2.8 and 2.9. Les conditions de la définition impliquent que Exp_θ et Ret_θ ont le même développement de Taylor à l'ordre un autour de 0_θ , ce qui confirme bien l'idée intuitive de la ressemblance entre une rétraction et l'exponentielle riemannienne.

Avant de parler des travaux de la communauté d'approximation stochastique sur une variété riemannienne, il est important de souligner le caractère essentiel des travaux des communautés d'optimisation déterministe et de méthodes numériques sur les variétés riemanniennes. En particulier [AMS08; EAS98] ont ouvert la porte des variétés riemanniennes à la communauté d'optimisation. En comparaison avec les autres branches de la géométrie riemannienne, les introductions faites dans ces articles pour les communautés d'optimisation ou de méthodes numériques font d'énormes efforts de sobriété, en construisant uniquement les notions nécessaires pour les algorithmes d'optimisation. De plus, les exemples traités dans les articles ci-dessus sont devenus de véritables études de cas pour la communauté d'optimisation stochastique. Pour modéliser les variables aléatoires sur les variétés riemanniennes, des travaux comme ceux de [Pen99] ont ouvert la voie pour les statistiques sur les variétés riemanniennes. L'effort combiné de ses trois communautés a mis en place le contexte nécessaire pour rendre possible l'approximation stochastique sur une variété riemannienne.

A.2.2 Premier résultat asymptotique de convergence

En 2013, le premier travail d'approximation stochastique a été paru, il s'agit de l'article [Bon13]. L'auteur y prouve le résultat suivant.

Theorem A.9 ([Bon13], théorème 1). *Considérons la suite $(\theta_n)_{n \in \mathbb{N}}$ définie par la récurrence (A.6) avec la fonction de champ moyen $h = -\text{grad } f$, où $f : \Theta \rightarrow \mathbb{R}$ est une fonction trois fois différentiable sur une variété riemannienne Θ connexe et dont le rayon d'injectivité est minoré uniformément (voir Section 2.11.6). On suppose que la suite des pas $(\eta_n)_{n \in \mathbb{N}^*}$ vérifie la condition standard $\sum_{n \in \mathbb{N}^*} \eta_n = \infty$ ainsi que $\sum_{n \in \mathbb{N}^*} \eta_n^2 < \infty$. Supposons de plus qu'il existe un ensemble compact K tel que pour tout $n \in \mathbb{N}$, $\theta_n \in K$ et qu'il existe $C > 0$ tel que pour tout $\theta \in K$ et tout $x \in X$ on a $\|H_\theta(x)\|_\theta \leq C$. Alors $(f(\theta_n))_{n \in \mathbb{N}}$ converge presque sûrement et $\lim_{n \rightarrow \infty} \text{grad } f(\theta_n) = 0$ p.s.*

Pour l'état de l'art, les hypothèses faites sur la variété et la suite de pas restent très pertinentes. L'ensemble compact K existe si la variété est compacte, ce qui est le cas pour la variété compacte de Stiefel et la variété de Grassmann, voir Exemples 2.8 and 2.9. Cependant, en général, il n'y a pas de garantie en pratique, et en général on s'en convainc en vérifiant empiriquement que les trajectoires simulées restent dans un compact de leur choix. Cette hypothèse est très problématique parce que (1) elle restreint l'initialisation à l'intérieur du compact, ce qui veut dire que l'analyse n'est faite que localement et jamais sur toute la variété, (2) la vérification empirique semble douteuse et ne constitue pas une preuve et (3) l'introduction de cette hypothèse par [Bon13] a motivé une partie importante de la communauté à l'adopter, ce qui conduit à des résultats dans lesquels le diamètre du compact K apparaît explicitement dans les majorations non-asymptotiques. L'un de nos buts par la suite est d'enlever complètement cette hypothèse. Ceci nous poussera sûrement à envisager d'autres techniques de preuve, comme on le verra par la suite. L'hypothèse sur le bruit est assez forte, parce qu'elle demande que le support du bruit soit borné. Des équivalents riemanniens de A22, où l'on ne nécessite des bornes que sur les moments du bruit sont bien plus communs actuellement.

Ce résultat est en lui-même un tremplin pour le reste de la communauté, et il est important de mentionner que l'impact de [Bon13] vient surtout de leur étude du schéma de rétraction ainsi que d'une pléthore d'exemples. En effet, dans [Bon13, Theorem 2], l'auteur obtient le même résultat que Theorem A.9 dans le cas de la rétraction (A.7), sous l'hypothèse supplémentaire que pour tout $\theta \in \Theta$, Ret_θ est différentiable deux fois.

Ce deuxième résultat, combiné avec les brillantes applications dans des problèmes de suivi de sous-espaces (sur la variété de Grassmann), de calcul de moyenne de Karcher (sur un espace hyperbolique) et d'estimation sur la variété des matrices symétriques définies positives, souligne deux points forts de l'analyse dans le cadre riemannien. Premièrement, l'analyse riemannienne encourage l'apparition de nouveaux algorithmes en choisissant minutieusement une métrique sur la variété adaptée au problème puis en calculant la géodésique qui en résulte ou un schéma de rétraction (A.6)-(A.7). Ces algorithmes sont étudiés assez naturellement dans le cadre général de l'approximation stochastique riemannienne. Dans un second temps, on peut aussi analyser un algorithme existant, qui n'est pas apparu comme un algorithme riemannien, mais qui en fait approxime le schéma géodésique suffisamment pour pouvoir l'analyser comme un schéma à rétraction dans le contexte riemannien. Illustrons ce point par le premier exemple de [Bon13, Section 4.1] et la règle d'Oja.

Exemple A.10 (Estimation de composantes principales en ligne). Soit \mathbf{A} une matrice de covariance de taille $d \times d$, pour laquelle on souhaite estimer les r premières composantes principales, c'est-à-dire l'espace engendré par les r vecteurs propres associés aux r plus grandes valeurs propres (avec multiplicité). Comme d'habitude dans le cadre en ligne, on se donne une suite i.i.d. d'observations $(X_n)_{n \in \mathbb{N}^*}$ à valeurs dans \mathbb{R}^d , et on fait l'hypothèse que $\mathbb{E}[X_1 X_1^\top] = \mathbf{A}$. Pour répondre au problème posé, regardons d'abord l'algorithme appelé la règle d'Oja (see [Oja92]), où θ_0 est une matrice de taille $d \times r$ telle que $\theta_0^\top \theta_0 = \mathbf{I}_r$. Le schéma de l'algorithme est le suivant,

$$\theta_{n+1} = \theta_n + \eta_{n+1} \left(\mathbf{I}_d - \theta_n \theta_n^\top \right) X_{n+1} X_{n+1}^\top \theta_n. \quad (\text{A.8})$$

Pour simplifier la géométrie du problème, regardons le cas particulier $r = 1$. Le cas général est détaillé dans Section 2.4.1. Une façon de comprendre cet algorithme est d'introduire la fonction de coût $f : \mathbb{R}^d \rightarrow \mathbb{R}$, définie pour tout $\theta \in \mathbb{R}^d$ par $f(\theta) = -\theta^\top \mathbf{A} \theta / 2$. En l'évaluant uniquement sur des vecteurs unitaires, on minimise f lorsque θ est un vecteur propre associé à la valeur propre principale de \mathbf{A} , et le gradient euclidien de f est donné par $\text{grad}^{(E)} f(\theta) = -\mathbf{A} \theta$. Cependant, la règle d'Oja n'utilise pas le schéma (A.1) avec ce gradient, mais il précompose le gradient par la matrice de projection $\mathbf{I}_d - \theta_n \theta_n^\top$. L'interprétation géométrique est que, puisqu'on s'intéresse pas à la norme de θ_n , on n'a besoin que de suivre les parties de $\text{grad}^{(E)} f(\theta_n)$ qui sont orthogonales à θ_n . On accomplit ceci en projetant le gradient sur le supplémentaire orthogonal de $\text{Span}(\theta_n)$, ce qui est fait par la matrice $\mathbf{I}_d - \theta_n \theta_n^\top$.

Ensuite, pour interpréter (A.8) comme un schéma de rétraction de la forme (A.7), on choisit comme variété la sphère unité plongée dans \mathbb{R}^d , que l'on note \mathbb{S}^{d-1} , et on calcule le gradient riemannien de f sur la sphère. En utilisant [AMS08, Example 3.6.1] on obtient $\text{grad}^{(S)} f(\theta) = (\mathbf{I}_d - \theta^\top \theta) \text{grad}^{(E)} f(\theta) = -(\mathbf{I}_d - \theta^\top \theta) \mathbf{A} \theta$. La rétraction la plus simple est celle que l'on a utilisé dans Exemple A.5, i.e. $\text{Ret}_\theta^{(S)} : v \in \text{T}_\theta \mathbb{S}^{d-1} \mapsto (\theta + v) / \|\theta + v\| \in \mathbb{S}^{d-1}$, qui produit le schéma suivant

$$\theta_{n+1} = \text{Ret}_{\theta_n}^{(S)} \left(\eta_{n+1} \left(\mathbf{I}_d - \theta_n^\top \theta_n \right) X_{n+1} X_{n+1}^\top \theta_n \right).$$

A.2.3 Analyse non-asymptotique

Depuis le premier article du domaine [Bon13], une communauté s'est formée pour chercher à comprendre quelles propriétés de l'approximation stochastique sur des espaces euclidiens restent vraies sur des variétés riemanniennes. En tant que contrepartie riemannienne de [BM11], dans [ZS16] les auteurs y font une analyse non-asymptotique pour le cas de la descente de gradient stochastique pour des fonctions géodésiquement convexes ou géodésiquement fortement convexes.

Definition A.11 (Fonction géodésiquement convexe). Une fonction $f : \Theta \rightarrow \mathbb{R}$ est dite géodésiquement convexe, ou g -convexe, si pour tout $\theta_0, \theta_1 \in \Theta$, toute géodésique $\gamma : [0, 1] \rightarrow \Theta$ telle que $\gamma(0) = \theta_0$ et $\gamma(1) = \theta_1$ et tout $t \in [0, 1]$,

$$f(\theta_1) \geq f(\theta_0) + \left\langle \text{grad} f(\theta_0), \text{Exp}_{\theta_0}^{-1}(\theta_1) \right\rangle_{\theta_0}.$$

Definition A.12 (Fonction géodésiquement fortement convexe de paramètre λ). *Soit f une fonction g -convexe et soit $\lambda > 0$. On dit que f est géodésiquement fortement convexe de paramètre λ , ou g - λ -fortement convexe si pour tout $\theta_0, \theta_1 \in \Theta$, il existe $\gamma : [0, 1] \rightarrow \Theta$ telle que $\gamma(0) = \theta_0$ et $\gamma(1) = \theta_1$ tel que pour tout $t \in [0, 1]$,*

$$f(\theta_1) \geq f(\theta_0) + \left\langle \text{grad } f(\theta_0), \text{Exp}_{\theta_0}^{-1}(\theta_1) \right\rangle_{\theta_0} + \frac{\lambda}{2} \rho_{\Theta}^2(\theta_0, \theta_1),$$

où $\rho_{\Theta} : \Theta \times \Theta \rightarrow \mathbb{R}_+$ est la distance riemannienne sur Θ (voir Section 2.11.1 et (2.150)).

Les résultats de [ZS16] sont équivalents à ceux sur les espaces euclidiens, sous l'hypothèse que les itérées de l'algorithme $(\theta_n)_{n \in \mathbb{N}}$ restent dans un compact de Θ . Plus précisément, pour une fonction g -fortement convexe f sur Θ , [ZS16, Theorem 12] établit qu'il existe $C > 0$ tel que pour tout $n \in \mathbb{N}$,

$$\mathbb{E} \left[f(\bar{\theta}_n) - f(\theta^*) \right] \leq C/n,$$

où $\bar{\theta}_0 = \theta_0$, $\bar{\theta}_{s+1} = \text{Exp}_{\bar{\theta}_n}^{-1} \{ (2(n+1))^{-1} \text{Exp}_{\bar{\theta}_n}^{-1}(\theta_{s+1}) \}$ et θ^* est l'unique solution au problème $h(\theta^*) = 0$. Pour le cas convexe, [ZS16, Theorem 10] établit sous les mêmes hypothèses qu'il existe $C > 0$ tel que pour tout $n \in \mathbb{N}$,

$$\mathbb{E} \left[f(\bar{\theta}_n) - f(\theta^*) \right] \leq C/\sqrt{n}.$$

A.2.4 Méthodes par quantité de mouvement et la méthode de Nesterov d'accélération du gradient dans le cas déterministe

Il y a eu récemment des efforts pour implémenter différentes versions de la descente de gradient stochastique avec quantité de mouvement pour les fonctions convexes sur des variétés riemanniennes. Bien que la communauté n'ait analysé que le cas déterministe, ce qui veut dire que $H_{\theta}(X) = -\text{grad } f(\theta)$, ces travaux sont pertinents pour le reste de la communauté et pourraient inspirer de nouveaux schémas stochastiques. Sur des espaces vectoriels, l'idée en ajoutant la quantité de mouvement est la suivante

$$v_{n+1} = \beta v_n + \eta H_{\theta_n}(X_{n+1}), \quad \text{and} \quad \theta_{n+1} = \theta_n + v_{n+1}, \quad (\text{A.9})$$

où $(v_n)_{n \in \mathbb{N}}$ est la suite des quantités de mouvement ou vitesses, $\beta \in (0, 1)$ est l'hyperparamètre de décroissance, pris souvent égal à 0.9 et $\theta \mapsto H_{\theta}(X)$ est un estimateur du champ moyen $-\text{grad } f$ (dans le cas déterministe il sera égal à $-\text{grad } f$). L'idée intuitive est qu'en gardant une partie de la vitesse des itérations précédentes, la trajectoire sera plus lisse, ce qui limitera l'oscillation lorsque l'on s'approche de la solution. En pratique, les techniques par quantité de mouvement permettent d'échapper aux points selles, et il a été prouvé qu'elle obtiennent des bornes de convergence de la forme $f(\theta_n) - f(\theta^*) = \mathcal{O}(1/n^2)$ (voir [Nes83]) pour des fonctions convexes. En particulier l'accélération de Nesterov (voir [Nes04; Nes83]) a connu un succès époustouflant sur les espaces vectoriels. En voici

l'algorithme,

$$\begin{aligned}\tilde{\theta}_n &= \theta_n + \beta v_n, \\ v_{n+1} &= \beta v_n + \eta H_{\tilde{\theta}_n}(X_{n+1}), \\ \theta_{n+1} &= \theta_n + v_{n+1}.\end{aligned}$$

La principale différence par rapport au schéma (A.9) est que le gradient est calculé *après* le pas effectué par la vitesse résiduelle. Pour comprendre intuitivement pourquoi cette différence améliore la performance, on peut dire que l'étape de gradient venant après la dérive due à la vitesse va avoir tendance à *corriger* les mouvements indésirables par la vitesse résiduelle, plutôt que de les accentuer.

Dans le cas riemannien, les méthodes par quantité de vitesse demandent plus de calculs que prévu en venant du monde vectoriel parce que les vecteurs vitesse doivent être transportés parallèlement au fur et à mesure que θ se déplace. Dans le cas déterministe, [Ali+21] propose le schéma suivant, initialisé en $\theta_0 = v_0 \in \Theta$,

$$\begin{aligned}\tilde{\theta}_n &= \text{Exp}_{v_n} \left(\beta_n \text{Exp}_{v_n}^{-1}(\theta_n) \right), \\ \theta_{n+1} &= \text{Exp}_{\tilde{\theta}_n} \left(-\eta \text{grad} f(\tilde{\theta}_n) \right), \\ v_{n+1} &= \text{Exp}_{v_n} \left(-a_{n+1} \text{T}_{\beta_n 0}^{\gamma_n} \text{grad} f(\tilde{\theta}_n) \right),\end{aligned}\tag{A.10}$$

où l'on définit $\gamma_n : [0, 1] \rightarrow \Theta$ comme la géodésique qui commence à $\gamma_n(0) = v_n$ et qui termine en $\gamma_n(1) = \theta_n$, β_n est calculé sur $[0, 1]$ pour minimiser $f \circ \gamma_n$ et $(a_n)_{n \in \mathbb{N}}$ est donné par $a_0 = 0$ et a_{n+1} est la solution positive à l'équation $c_0 a_{n+1}^2 - c_1 a_{n+1} - c_1 \sum_{k=0}^n a_k = 0$, avec des constantes positives $c_0, c_1 > 0$ choisies en fonction de la courbure de la variété d'à quel point $\text{grad} f$ est lisse. Finalement, $\text{T}_{t_1 t_0}^{\gamma_n} : \text{T}_{\gamma(t_1)} \Theta \rightarrow \text{T}_{\gamma(t_0)} \Theta$ est le transport parallèle depuis $\text{T}_{\gamma(t_1)} \Theta$ vers $\text{T}_{\gamma(t_0)} \Theta$ le long de la courbe γ (voir Section 2.11.4 pour plus de détails). Certaines des complications du schéma riemannien viennent du fait que les vecteurs tangents ne sont tangents qu'en un point. Ainsi, pour ajouter des vecteurs tangents, il faut d'abord s'assurer qu'ils sont tangents au même point, ce que l'on fait en transportant les vecteurs parallèlement sur le même espace tangent. En effet, si $\text{grad} f(\tilde{\theta}_n)$ est utilisé pour mettre à jour à la fois θ_n et v_n , il faudra le transporter depuis $\text{T}_{\tilde{\theta}_n} \Theta$ vers $\text{T}_{v_n} \Theta$, ce qui justifie l'apparition du transport parallèle. Leur analyse s'applique aussi bien aux fonctions g -convexes qu'aux fonctions faiblement quasi- g -convexes (voir **F3** dans Section 3.4 pour une définition). Ils font aussi l'hypothèse que les itérations de l'algorithme restent dans un compact géodésiquement convexe $K \subset \Theta$ avec un diamètre borné, et que l'exponentielle riemannienne est un difféomorphisme global sur K . Ils obtiennent des bornes de la forme $f(\theta_n) - f(\theta^*) \leq C_0 \max(C_1/k^2, C_2/k) + \varepsilon$, ce qui, dans le pire des cas, fait aussi bien que le schéma (A.1), mais parfois fait même mieux (voir [Ali+20, Theorem 4]). Cependant, bien que leurs résultats soient prometteurs, l'analyse n'a pas encore été faite avec du bruit dans les observations.

Aussi dans le cadre déterministe, [ZS18] et [AS20] ont étudié un schéma qui suit généralement la forme de (A.10) mais sous des hypothèses de forte g -convexité. Les autres hypothèses suivent celles de [Ali+20], sauf que [AS20] réussit à enlever l'hypothèse du

compact K (ce qui veut dire qu'ils font une analyse globale) en introduisant des fonctions de Lyapunov. Les conclusions de [ZS18] sont que l'accélération (*i.e.* une vitesse bornée par $\mathcal{O}(1/n^2)$) est obtenue lorsque le schéma est initialisé près de la solution, ce qui peut prendre du temps dans des problèmes mal conditionnés. Dans [AS20], les auteurs montrent que leur algorithme est strictement plus rapide que la descente de gradient classique en utilisant dans leur preuve des fonctions des Lyapunov et une inégalité de distorsion de la métrique.

A.3 Introduction aux méthodes de Monte Carlo par chaînes de Markov

A.3.1 Le problème de simuler une loi cible

Le second sujet traité dans ce manuscrit vise à simuler selon une distribution de probabilités. Parmi tous les cadres possibles où simuler une loi cible est utile, on va présenter le cas des statistiques bayésiennes. Grossièrement, simuler selon une loi cible est utilisé en statistiques bayésiennes pour explorer la loi postérieure.

La principale hypothèse en statistique bayésienne est qu'il y a deux modèles probabilistes, un pour les données observées – qui est la plupart de temps un modèle paramétrique, disons dépendant d'un paramètre θ – et l'autre modèle parle du paramètre θ . La distribution du paramètre θ , qui est appelée loi *a priori*, résume toute la connaissance sur le paramètre *avant* de considérer les observations en jeu. Le but ensuite est de combiner l'information a priori avec celle apportée par les observations. On fait ceci en calculant la distribution du paramètre *conditionnellement* aux observations, que l'on appelle loi postérieure ou *a posteriori*.

Plus précisément, soit (Θ, \mathcal{F}) un espace mesurable, et soit θ une variable aléatoire à valeurs dans cet espace, de loi ν_θ . Le plus souvent, l'espace pour le paramètre Θ est discret ou un sous-ensemble de \mathbb{R}^d pour $d \geq 1$. Pour cette raison, on suppose que ν_θ admet une densité par rapport à une mesure de référence ϖ , et on note sa densité ν_θ . Dans un second temps, on observe w , une réalisation d'une variable aléatoire W à valeurs dans l'espace mesurable (W, \mathcal{W}) . Ensuite, on précise la famille de densités conditionnelles de W sachant θ ,

$$\mathcal{P}_\Theta = \{w \mapsto \mathcal{L}(w|\theta) \mid \theta \in \Theta\} ,$$

par rapport à une mesure de référence μ sur (W, \mathcal{W}) . On appelle vraisemblance du modèle la fonction $\theta \mapsto \mathcal{L}(W|\theta)$.

Ayant décrit la loi de θ ainsi que celle de W sachant θ , on a accès à la loi jointe, qui admet la densité $(\theta, w) \mapsto \nu_\theta(\theta)\mathcal{L}(w|\theta)$ par rapport à $\varpi \otimes \mu$ et la loi de θ sachant W , la postérieure. Par le théorème de Bayes ([Sch95, Theorem 1.31]), on a l'expression suivante pour la densité de la postérieure $f_{\theta|X}$, définie pour ν_θ -presque tout θ et μ -presque tout w par

$$f_{\theta|X}(w, \theta) = \frac{\mathcal{L}(w|\theta)\nu_\theta(\theta)}{\int_\Theta \mathcal{L}(w|\tilde{\theta})\nu_{\tilde{\theta}}(\tilde{\theta})\varpi(d\tilde{\theta})} , \tag{A.11}$$

où pour tout w ,

$$f_X(w) = \int_{\Theta} \mathcal{L}(w|\theta) \nu_{\theta}(\theta) \varpi(d\theta), \quad (\text{A.12})$$

est la densité de W par rapport à μ .

A.3.2 Algorithmes de Monte Carlo et leurs versions markoviennes

Soit π une loi de probabilité sur un espace mesurable (E, \mathcal{E}) , appelée la loi cible. On cherche à estimer des quantités de la forme $\int_E f(x) \pi(dx)$, où $f : E \rightarrow \mathbb{R}$ est une fonction intégrable par rapport à π . L'idée des méthodes de Monte Carlo est d'utiliser des estimateurs de la forme

$$\hat{f}_N = \frac{1}{N+1} \sum_{i=0}^N f(Y_i),$$

où $(Y_i)_{i \in \mathbb{N}}$ est une suite de variables aléatoires. Les différentes méthodes de Monte Carlo proposent différentes hypothèses sur la suite $(Y_i)_{i \in \mathbb{N}}$ pour s'assurer que \hat{f}_N est un bon estimateur.

La méthode traditionnelle de Monte Carlo repose sur la loi forte des grands nombres, puisqu'elle considère une suite $(Y_i)_{i \in \mathbb{N}}$ de variables aléatoires indépendantes et identiquement distribuées (i.i.d. pour faire court) de loi π . Dans ce cadre, la loi forte des grands nombres établit que \hat{f}_N converge presque sûrement, lorsque N tend vers l'infini, vers $\int_E f(x) \pi(dx)$. De plus, si f^2 est intégrable par rapport à π , on sait que l'erreur asymptotique est de l'ordre de $\mathcal{O}(1/\sqrt{N})$ en faisant appel au théorème central limite. La simplicité de cette méthode entraîne un coût : il faut un tirage de N variables aléatoires i.i.d. de loi π . On peut citer la méthode de *inverse transform sampling* et l'algorithme d'acceptation-rejet parmi les plus connues de la littérature de méthodes classiques (voir [RC10, Section 2.1.2] et [RC10, Section 2.3]).

Un problème de ces méthodes est qu'elles souffrent du fléau de la dimension. Un autre est que parfois on n'a accès à la densité π qu'à une constante multiplicative près. Par exemple, en inférence bayésienne, où π admet pour densité

$$x \mapsto \frac{\mathcal{L}(w|x)}{f_X(w)},$$

par rapport à la loi a priori (A.11), où w sont les observations. En effet, sauf quand la loi a priori et la vraisemblance sont des lois conjuguées, la loi marginale qui apparaît dans (A.12) n'a pas de forme close. Ceci pose un sévère problème, en limitant l'applicabilité des méthodes ci-dessus.

Une solution pour certains de ces problèmes est d'utiliser des méthodes de Monte Carlo par chaînes de Markov (on utilise l'abréviation anglaise MCMC). L'idée est de ne plus recourir à des observations i.i.d., mais plutôt de générer une suite $(Y_i)_{i \in \mathbb{N}}$ par une chaîne de Markov qui a pour loi invariante π . Lorsque les hypothèses classiques sont satisfaites sur π et la chaîne, on obtient un équivalent de la loi forte des grands nombres pour la chaîne de Markov (voir [MT09, Chapter 17]). Plus précisément [MT09, Theorem

17.1.7] montre que si le noyau de Markov P est Harris récurrent et admet π comme loi invariante, alors pour tout $n \in \mathbf{E}$,

$$\lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{i=0}^N f(Y_i) = \int_{\mathbf{E}} f(y) dy, \quad \mathbb{P}_x\text{-almost surely},$$

où \mathbb{P}_x est la loi induite par le noyau de Markov P avec la loi initiale δ_x sur $(\mathbf{E}^{\mathbb{N}}, \mathcal{E}^{\otimes \mathbb{N}})$ et $(Y_i)_{i \in \mathbb{N}}$ est la chaîne canonique sur $(\mathbf{E}^{\mathbb{N}}, \mathcal{E}^{\otimes \mathbb{N}}, \mathbb{P}_x)$. Pour implémenter ces méthodes, on a besoin de construire un noyau de Markov qui admet π comme loi invariante. Une procédure générique pour y arriver est d'utiliser l'algorithme de Metropolis-Hastings, ou une de ses variantes. Présentons cet algorithme.

A.3.3 L'algorithme de Metropolis-Hastings pour simuler une loi cible

La première apparition de cet algorithme dans la littérature remonte à [Met+53], puis l'algorithme a été amélioré dans [Has70] et [Tie98]. Soit ν une mesure sur $(\mathbf{E}, \mathcal{E})$ qui domine π et soit $x, y \in \mathbf{E}^2 \mapsto q(x, y)$ la densité conditionnelle de y sachant x , autrement dit, pour tout $x \in \mathbf{E}$, $y \mapsto q(x, y)$ est une densité par rapport à ν sur \mathbf{E} et pour ν -presque tout $y \in \mathbf{E}$, $x \mapsto q(x, y)$ est mesurable. Ensuite, étant donné une initialisation $X_0 \in \mathbf{E}$ mesurable sur $(\mathbf{E}, \mathcal{E})$, on définit la chaîne de Markov $(X_n)_{n \in \mathbb{N}}$ pour tout $n \in \mathbb{N}$ par

$$X_{n+1} = Y_{n+1} \mathbb{1}\{U_{n+1} < \alpha(X_n, Y_{n+1})\} + X_n \mathbb{1}\{U_{n+1} \geq \alpha(X_n, Y_{n+1})\}, \quad (\text{A.13})$$

où Y_{n+1} est une variable aléatoire qui, conditionnellement à la tribu engendrée par les $n+1$ premiers termes de la chaîne $\mathcal{F}_n = \sigma(X_0, \dots, X_n)$, est distribuée selon la densité de probabilité $y \mapsto q(X_n, y)$ et $(U_n)_{n \in \mathbb{N}^*}$ est une suite i.i.d. de variables aléatoires uniformes sur $[0, 1]$. La fonction $\alpha : \mathbf{E} \times \mathbf{E} \rightarrow [0, 1]$ est supposée mesurable et on l'appelle le taux d'acceptation. Soit $P : \mathbf{E} \times \mathcal{E} \rightarrow [0, 1]$ le noyau de Markov de la chaîne $(X_n)_{n \in \mathbb{N}}$. On remarque que le noyau de Markov P , associé à la chaîne de Markov définie en (A.13) est donné par, pour tout $x \in \mathbf{E}$ et tout $\mathbf{A} \in \mathcal{E}$,

$$P(x, \mathbf{A}) = \int_{\mathbf{A}} \alpha(x, y) q(x, y) \nu(dy) + \delta_x(\mathbf{A}) \int_{\mathbf{E}} (1 - \alpha(x, y)) q(x, y) \nu(dy).$$

Cette procédure d'acceptation-rejet permet de simuler une chaîne de Markov avec noyau de Markov P en utilisant la famille de lois $y \mapsto q(x, y)$ dites lois des propositions. On note aussi π la densité de la loi cible π par rapport à ν . On rappelle que ce qui nous intéresse est de montrer que la chaîne que l'on vient de construire admet π comme loi invariante. Cependant, on a même une propriété plus forte puisque P est réversible par rapport à la loi π si et seulement si pour toute fonction mesurable et bornée $g : \mathbf{E}^2 \rightarrow \mathbb{R}$,

$$\int_{\mathbf{E}^2} g(x, y) \pi(x) P(x, dy) \nu(dx) = \int_{\mathbf{E}^2} g(x, y) \pi(y) P(y, dx) \nu(dy).$$

Mais, en remarquant la présence de la mesure de Dirac sur la deuxième partie de P , ceci équivaut à : pour ν -presque tout $x, y \in \mathbf{E}$,

$$\alpha(x, y) \pi(x) q(x, y) = \alpha(y, x) \pi(y) q(y, x).$$

Cette dernière condition est vérifiée en particulier si le taux d'acceptation est donné par

$$\alpha(x, y) = \begin{cases} \min\left(1, \frac{\pi(y)q(y, x)}{\pi(x)q(x, y)}\right) & \text{si } \pi(x)q(x, y) > 0 \text{ et } \pi(y)q(y, x) > 0 \\ 0 & \text{sinon .} \end{cases} \quad (\text{A.14})$$

L'algorithme de Metropolis-Hastings est particulièrement utile lorsqu'il est facile de simuler la loi de la proposition admettant pour densité $y \mapsto q(x, y)$.

Si on cherche le choix le plus simple pour la proposition, on peut prendre pour tout $x, y \in \mathbb{E}^2$, $q(x, y) = q(y)$, de façon à ce que la proposition ne dépende pas sur l'état courant de la chaîne. Ceci donne la forme suivante pour le taux d'acceptation α ,

$$\alpha(x, y) = 1 \wedge \frac{\pi(y)q(x)}{\pi(x)q(y)} .$$

Cette version de l'algorithme de Metropolis s'appelle la méthode de Metropolis par simulation indépendante, puisque la proposition ne dépend pas du point courant x . Un autre choix très classique pour la loi de proposition repose sur la structure de groupe abélien de \mathbb{E} et sur le fait que ν soit une mesure invariante par translation. Dans ce cas, on considère la marche aléatoire symétrique sur $(\mathbb{E}, \mathcal{E})$, où les incréments ont pour densité \tilde{q} par rapport à ν . Ceci entraîne que la suite de propositions est donnée par

$$Y_{k+1} = X_k + Z_{k+1} ,$$

où la variable aléatoire X_k est l'état courant de la chaîne à l'étape k et les variables aléatoires $(Z_i)_{i \geq 1}$ sont i.i.d. et admettent pour densité \tilde{q} par rapport à ν . De plus, il existe une relation directe entre la densité de proposition q et la densité de l'incrément \tilde{q} : pour tout $x, y \in \mathbb{E}$ $q(x, y) = \tilde{q}(x - y)$. L'exemple par excellence dans ce cadre est l'algorithme de Metropolis par marche aléatoire symétrique. On le définit sur \mathbb{R}^d avec la mesure de Lebesgue comme mesure de référence, et la loi de l'incrément est une loi gaussienne centrée de matrice de covariance $\sigma^2 \text{Id}$, où $\sigma^2 > 0$.

On termine cette partie en parlant de l'algorithme de Langevin ajusté par Metropolis (on utilise l'abréviation anglaise MALA). On considère $\mathbb{E} = \mathbb{R}^d$ muni de sa tribu borélienne $\mathcal{B}(\mathbb{R}^d)$. De plus, supposons que π admet une densité strictement positive et lisse par rapport à la mesure de Lebesgue Leb , on la notera encore π et soit $U = -\log \pi$. MALA est basé sur la diffusion de Langevin qui est une équation différentielle Stochastique (EDS) définissant un semigroupe de Markov réversible par rapport à π . Cependant, simuler une solution exacte de cette EDS en temps continu n'est pas possible en général. Ainsi, on considère plutôt la discrétisation d'Euler Maruyama associée à cette EDS, qui s'écrit

$$Y_{k+1} = Y_k - \gamma \nabla U(Y_k) + \sqrt{2\gamma} Z_{k+1} , \quad (\text{A.15})$$

où $\gamma > 0$ est la taille du pas de la discrétisation et $(Z_k)_{k \geq 1}$ est une suite de variables aléatoires d dimensionnelles, i.i.d. de loi gaussienne centrée réduite. Cette chaîne de Markov est une approximation discrète de la diffusion de Langevin. Il est important de remarquer qu'elle admet une loi invariante π_γ – qui dépend du choix du pas γ – et que sous des

hypothèses assez faibles sur U , la chaîne est aussi ergodique. Et même si π_γ est différent de π la plupart des fois, la différence est parfois suffisamment petite comme pour prendre ceci comme chaîne approchant π . Cet algorithme a été introduit par [Erm75] et [Par81] pour des application en dynamique moléculaire. Il a reçu de plus de plus d'attention par la communauté d'apprentissage statistique grâce à [Gre83], [GM94a] et dans la communauté de statistique computationnelle grâce à [Nea93] et [RT96a].

Néanmoins, il existe des cas où simuler une loi *proche* de π n'est pas suffisant, puisque cela va biaiser les estimateurs $(\hat{f}_N)_{N \in \mathbb{N}}$, et même ça leur fera perdre la propriété de consistance puisqu'ils ne convergeront plus vers $\int f d\pi$ en général. L'idée de [RDF78] et de [RT96a] est d'utiliser la chaîne de Markov $(Y_n)_{n \in \mathbb{N}}$ définie en (A.15) comme la proposition dans l'algorithme de Metropolis-Hastings. On l'appelle, d'après [RT96a], Algorithme de Langevin ajusté par Metropolis (MALA).

A.4 Analyse asymptotique des algorithmes de Metropolis-Hastings : échelonnement optimal

On s'intéresse ici à l'analyse d'algorithmes de Metropolis-Hastings lorsque la dimension tend vers l'infini. L'échelonnement optimal, inventé par [RGG97a], fournit des outils pour comprendre les propriétés de convergence des méthodes MCMC en grande dimension.

Lorsqu'on utilise l'algorithme de Metropolis par marche aléatoire symétrique ou MALA sur \mathbb{R}^d , il faut bien choisir le paramètre dans la loi de proposition. Dans le cas du premier algorithme, il s'agit de la variance de la loi gaussienne, et dans le cas de MALA c'est le pas de la discrétisation. Le choix du paramètre va avoir une influence sur la variance de la proposition et donc sur le comportement de la chaîne de Markov : de petits incréments exploreront lentement le support de π avec un taux d'acceptation élevé, tandis que des incréments plus grands se déplaceront plus vite sur tout le support de π mais ils seront rejetés plus souvent. L'échelonnement optimal construit un cadre pour formuler la question "Quel est le choix optimal pour ce paramètre σ^2 ?" et y répond en fonction de la dimension d . Nous présentons maintenant les résultats classique de ce domaine.

A.4.1 Introduction à l'échelonnement optimal : l'algorithme de Metropolis par marche aléatoire symétrique

Comme on vient de dire, l'étude de l'échelonnement optimal a commencé avec les travaux de [RGG97a], où ils ont étudié l'algorithme de Metropolis par marche aléatoire symétrique (on prendra l'abréviation anglaise RWM) sur \mathbb{R}^d . Pour une densité cible strictement positive π^d sur \mathbb{R}^d , rappelons l'expression de la chaîne de Markov créée par l'algorithme RWM, qui est de la forme (A.13). Pour toute valeur initiale $X_0^d \in \mathbb{R}^d$ et pour tout $k \in \mathbb{N}$,

$$X_{k+1}^d = X_k^d + \sigma_d Z_{k+1}^d \mathbb{1}\{U_{k+1}^d < \alpha_d(X_k^d, X_k^d + \sigma_d Z_{k+1}^d)\},$$

où $\sigma_d > 0$, $(Z_{k+1}^d)_{k \in \mathbb{N}}$ est une suite de variables aléatoires i.i.d. de loi gaussienne centrée et réduite sur \mathbb{R}^d et $(U_{k+1}^d)_{k \in \mathbb{N}}$ est une suite de variables aléatoires i.i.d. uniformes sur $[0, 1]$,

A.4. Échelonnement optimal

de plus $(Z_{k+1}^d)_{k \in \mathbb{N}}, (U_{k+1}^d)_{k \in \mathbb{N}}$ et X_0 sont indépendantes. De plus, le taux d'acceptation α_d est donné par, pour tout $x, y \in \mathbb{R}^d$,

$$\alpha_d(x, y) = 1 \wedge \frac{\pi^d(y)}{\pi^d(x)}.$$

L'utilisateur peut choisir le paramètre σ_d *i.e.* la variance de la loi de proposition. Une question assez naturelle est “comment *bien* choisir σ_d ?” En effet, une grande variance permet d'explorer plus vite le support de π^d , mais au prix d'un taux d'acceptation faible, ce qui ultimement ralentit l'exploration. Dans le cas d'une distribution initiale différente de π^d , le choix de σ_d aurait une influence aussi sur le temps de mélange de la chaîne. Le point de départ de l'échelonnement optimal consiste à trouver des valeurs de σ_d telles que l'acceptation-rejet ne soit pas asymptotiquement trivial, c'est-à-dire que l'on veut que le taux d'acceptation converge à l'intérieur du segment $]0, 1[$, *i.e.*

$$\lim_{d \rightarrow +\infty} \mathbb{E}_{\pi^d} \left[\alpha_d \left(X_0^d, X_0^d + \sigma_d Z_1^d \right) \right] \in]0, 1[.$$

Pour répondre à cette question, il nous faut un cadre pour chaque $d \in \mathbb{N}$, et le choix de π^d par rapport à la dimension d a une forte influence sur le comportement de la chaîne de Markov. La première hypothèse, faite dans [RGG97a], repose sur une construction de π^d en tant que mesure produit de la forme suivante. Pour tout $x^d = (x_1^d, \dots, x_d^d) \in \mathbb{R}^d$,

$$\pi^d(x^d) = \prod_{i=1}^d \exp \left(-U(x_i^d) \right), \quad (\text{A.16})$$

où $U : \mathbb{R} \rightarrow \mathbb{R}$ est une fonction qui définit une loi unidimensionnelle $\pi(x) = \exp(-U(x))$. On suppose que ces fonctions vérifient les propriétés suivantes,

H 1 U est trois fois continûment différentiable et U' est Lipschitz.

H 2 $\mathbb{E}[(U'(W))^8]$ et $\mathbb{E}[(U''(W))^4]$ sont finis, où W est une variable aléatoire réelle de loi π .

[RGG97a] propose de prendre $\sigma_d = \ell d^{-\alpha/2}$ comme forme générale pour σ_d avec $\ell, \alpha \in \mathbb{R}_+^*$. Ensuite, [RGG97a, corollary 1.2] démontre que si pour tout $d \in \mathbb{N}$, X_0^d est de loi π^d , alors on a pour $\alpha = 1$

$$\lim_{d \rightarrow +\infty} \mathbb{E}_{\pi^d} \left[\alpha_d \left(X_0^d, X_0^d + \ell d^{-1/2} Z_1^d \right) \right] = 2\Phi(-\ell\sqrt{K}/2),$$

où Φ est la fonction de répartition d'une loi gaussienne centrée réduite et

$$K = \int_{\mathbb{R}} (U')^2(x) \pi(x) dx. \quad (\text{A.17})$$

Ayant trouvé l'échelonnement de $\sigma_d = \ell d^{-\alpha/2}$ tel que l'algorithme a une acceptation limite non-dégénérée, on peut maintenant s'intéresser au choix de la constante ℓ . [RGG97a] continue dans cette direction en étudiant le processus à temps continu $(V_t^d)_{t \geq 0}$ défini par

$$V_t^d = X_{\Gamma_t^d, 1}^d, \quad (\text{A.18})$$

où $(\Gamma_t^d)_{t \geq 0}$ est un processus de Poisson de paramètre d . Une explication intuitive de ce choix est que, puisque le processus de Poisson saute td fois en moyenne sur un intervalle de temps t , V_t^d est approximativement la première composante de la chaîne de Markov X^d à l'instant $[td]$. Dans [RGG97a, Theorem 1.1], les auteurs prouvent que si U satisfait les conditions H1 et H2, alors la suite de processus $(V^d)_{d \in \mathbb{N}^*}$ converge faiblement, dans la topologie de Skorokhod, vers un processus V qui satisfait les hypothèses suivantes : V_0 est de loi π et pour tout $t \geq 0$, V est solution de l'équation différentielle stochastique suramortie de Langevin,

$$dV_t = \sqrt{h(\ell)} dB_t^1 - \frac{1}{2} h(\ell) U'(V_t) dt ,$$

où $(B_t^1)_{t \geq 0}$ est un mouvement brownien unidimensionnel, $h(\ell)$ est souvent appelée la vitesse de la diffusion et est donnée pour tout $\ell > 0$ par

$$h(\ell) = 2\ell^2 \Phi \left(-\frac{\ell\sqrt{K}}{2} \right) ,$$

et K est défini par (A.17). Pour comprendre comment ce résultat nous aide à choisir le paramètre ℓ , on commence par définir un rééchelonnement du processus limite. Soit pour tout $t \in \mathbb{R}_+$, $\tilde{V}_t = V_{h(\ell)t}$. Alors \tilde{V} est la diffusion de Langevin avec vitesse unitaire :

$$d\tilde{V}_t = dB_t^1 - \frac{1}{2} U'(\tilde{V}_t) dt .$$

On comprend maintenant qu'en augmentant la vitesse h , la chaîne de Markov explore plus vite les trajectoires de \tilde{V} et donc \mathbb{R}^d , ce qui finit par être plus efficace pour notre problème initial de simulation de lois. Ainsi, on veut trouver la valeur de ℓ qui maximise h . On trouve dans [RGG97a] un critère assez direct pour trouver un tel ℓ . Soit $a(\ell) = 2\Phi(-\ell\sqrt{K}/2)$ la limite du taux d'acceptation, alors, par [RGG97a, Corollary 1.2] on peut exprimer $h(\ell)$ comme une fonction de $a(\ell)$ en inversant l'expression de $a(\ell)$. Leur conclusion est que choisir ℓ_0 tel que $a(\ell_0) = 0.234$ (précision à quatre chiffres significatifs) maximise la vitesse de la diffusion V . Une autre façon d'interpréter la vitesse de la diffusion est par le fait que $\lim_{d \rightarrow \infty} (d/2) \mathbb{E}[(X_0^d - X_1^d)^2] = h(\ell)$, où $\mathbb{E}[(X_0^d - X_1^d)^2]$ exprime la distance moyenne au carré faite à chaque itération.

Dans le même cadre, l'article [Dur+17] considère des hypothèses plus générales pour U , en particulier ils considèrent que la fonction est différentiable en norme L -p.

H 3 Il existe une fonction mesurable $\dot{U} : \mathbb{R} \rightarrow \mathbb{R}$ telle que :

(i) Il existe $p > 4$, $C > 0$ et $\beta > 1$ tels que pour tout $\theta \in \mathbb{R}$,

$$\left(\int_{\mathbb{R}} |U(x + \theta) - U(x) - \theta \dot{U}(x)|^p \pi(x) dx \right)^{1/p} \leq C |\theta|^\beta .$$

(ii) La fonction \dot{U} vérifie

$$\int_{\mathbb{R}} (\dot{U}(x))^6 \pi(x) dx < \infty .$$

La fonction \dot{U} est appelée la dérivée faible de U . Avec ces hypothèses, les auteurs démontrent dans [Dur+17, Theorem 2] la convergence du taux d'acceptation vers $a(\ell) = 2\Phi(-\ell\sqrt{K}/2)$ lorsque $\sigma_d = \ell d^{-1/2}$. Ensuite, dans [Dur+17, Theorem 3] ils démontrent qu'avec ces mêmes hypothèses, le processus V est solution de la même EDS que l'on a vu plus haut, avec en particulier la même expression pour la vitesse de la diffusion $h(\ell)$. Ces résultats sont intéressants pour ce manuscrit puisque les contributions faites au chapitre 4 sont un parallèle de ces résultats, mais pour l'algorithme MALA.

A.4.2 Échelonnement optimal de MALA

On peut se poser à nouveau la question “qu'est-ce qu'un bon choix pour σ_d ?” pour MALA. Rappelons les notations pour cet algorithme. On a une densité cible π^d sur \mathbb{R}^d qui s'écrit comme un produit de la forme (A.16) avec $U : \mathbb{R} \rightarrow \mathbb{R}$ qui est continûment différentiable et étant donnée une initialisation $X_0^d \in \mathbb{R}^d$, la chaîne de Markov est donnée, pour tout $k \in \mathbb{N}$, par

$$X_{k+1}^d = X_k^d + \tilde{Z}_{k+1}^d \mathbb{1}\{U_{k+1}^d < \alpha_d(X_k^d, X_k^d + \tilde{Z}_{k+1}^d)\},$$

où $\tilde{Z}_{k+1}^d = \sigma_d^2 \nabla \log \pi^d(X_k^d) + \sqrt{2}\sigma_d Z_{k+1}^d$ est l'incrément de la proposition, $\sigma_d \in \mathbb{R}_+^*$ et $(Z_i^d)_{i \geq 1}$ est une suite de variables aléatoires i.i.d. de loi gaussienne centrée réduite sur \mathbb{R}^d . En remarquant que conditionnellement à X_k^d , la loi de \tilde{Z}_{k+1}^d est une gaussienne d'espérance $\sigma_d^2 \nabla \log \pi^d(X_k^d)$ et de variance $2\sigma_d^2$, on peut simplifier α_d , défini en (A.14). Pour tout $x, y \in \mathbb{R}^d$, on écrit α_d ainsi,

$$\alpha_d(x, y) = 1 \wedge \frac{\pi^d(y) \exp\left(-\|x - y - \sigma_d^2 \nabla \log \pi^d(y)\|^2 / (4\sigma_d^2)\right)}{\pi^d(x) \exp\left(-\|y - x - \sigma_d^2 \nabla \log \pi^d(x)\|^2 / (4\sigma_d^2)\right)}.$$

Les objets étudiés sont très proches de ceux de RWM, où σ_d demeure le seul paramètre à choisir. En écrivant $\sigma_d = \ell d^{\alpha/2}$, on se demande quelle valeur de α permet de faire converger $\mathbb{E}_{\pi^d}[\alpha_d(X_0^d, X_0^d + \tilde{Z}_1^d)]$, lorsque $d \rightarrow +\infty$, dans $]0, 1[$? Le résultat a été apporté par [RR98a], où les auteurs montrent que pour $\alpha = 1/3$, le taux d'acceptation converge dans $]0, 1[$. Voici les hypothèses de ce résultat.

M 1 U est continûment différentiable huit fois sur \mathbb{R} et il existe un polynôme à coefficients réels P tel que pour tout $i \in \{0, \dots, 8\}$ et $x \in \mathbb{R}$, $|U^{(i)}(x)| \leq P(x)$.

M 2 Pour tout $p \geq 1$, $\int_{\mathbb{R}} |x|^p \pi_1(x) dx < \infty$.

Voici le résultat [RR98a, Theorem 1] tel qu'il y est énoncé.

$$\lim_{d \rightarrow +\infty} \mathbb{E}_{\pi^d} \left[\alpha^d \left(X_0^d, X_0^d + \ell^2 d^{-1/3} \nabla \log \pi^d(X_0^d) + \sqrt{2} \ell d^{-1/6} Z_1^d \right) \right] = 2\Phi(-\ell^3 \sqrt{J}/2), \quad (\text{A.19})$$

où Φ est encore la fonction de répartition empirique de la loi gaussienne centrée réduite et la nouvelle constante J est donnée par

$$J = (48)^{-1} \int_{\mathbb{R}} \left\{ 5(U^{(3)}(x))^2 - 3(U^{(2)}(x))^3 \right\} \pi(x) dx > 0. \quad (\text{A.20})$$

Dans un second temps, les auteurs de [RR98a] démontrent qu'un processus limite satisfait aussi l'EDS de Langevin. Soit encore une fois le processus unidimensionnel $(V_t^d)_{t \geq 0}$ défini par (A.18), mais avec $(\Gamma_t^d)_{t \geq 0}$ un processus de Poisson de paramètre $d^{1/3}$. Le résultat [RR98a, Theorem 2] énonce que si U vérifie M1 et M2, alors la suite de processus $(V^d)_{d \in \mathbb{N}^*}$ converge faiblement, dans la topologie de Skorokhod, vers un processus V qui vérifie les propriétés suivantes : la loi de V_0 est π et pour tout $t \geq 0$, V est solution de l'EDS suramortie de Langevin,

$$dV_t = \sqrt{h(\ell)} dB_t^1 - \frac{1}{2} h(\ell) U'(V_t) dt,$$

où $(B_t^1)_{t \geq 0}$ est un mouvement brownien réel, $h(\ell)$ est souvent appelée la vitesse de la diffusion et est donnée pour tout $\ell > 0$ par

$$h(\ell) = 2\ell^2 \Phi \left(-\frac{\ell^3 \sqrt{J}}{2} \right),$$

et J est défini par (A.20).

On interprète ce résultat avec les mêmes outils que pour l'algorithme de Metropolis par marche aléatoire symétrique, c'est à dire que la limite diffusiv nous dit que maximiser la vitesse de diffusion h correspond au choix optimal pour ℓ . En utilisant (A.19) et en exprimant h en fonction du taux d'acceptation limite, on trouve que maximiser $\ell \mapsto g(\ell)$ revient à avoir un taux d'acceptation limite (donné en (A.19)) égal à environ 0.574.

Il y a deux différences majeures par rapport à l'algorithme RWM. Les interpréter nous permettra de comparer les deux algorithmes. La première est que pour MALA, il est optimal de prendre $\sigma_d = \ell d^{-1/6}$ tandis que pour RWM, on avait $\sigma_d = \ell d^{-1}$. Cette différence veut dire qu'en général, on peut utiliser des valeurs de σ_d plus grandes pour MALA que pour RWM. Comme on peut construire un processus de saut à partir de ces chaînes de Markov qui convergent faiblement vers la même limite (quitte à échelonner pour avoir la même vitesse de diffusion), on sait que MALA aura besoin d'environ $\mathcal{O}(Td^{1/3})$ itérations pour arriver au temps T de la trajectoire, tandis que RWM en aura besoin de $\mathcal{O}(Td)$. Ainsi, on peut conclure que MALA est plus efficace dans le sens de l'échelonnement optimal. La deuxième différence entre les deux algorithmes est qu'à l'optimum, le taux d'acceptation est plus élevé pour MALA que pour RWM. En résumé, lorsque les hypothèses nécessaires pour la convergence de MALA sont satisfaites, MALA est un meilleur choix que RWM. Dans ce manuscrit, on s'intéresse à l'échelonnement optimal de MALA avec la loi Laplace pour cible, qui a pour densité $x \mapsto e^{-|x|}/2$. C'est intéressant parce que l'hypothèse M1 n'est pas vérifiée puisque $x \mapsto |x|$ n'est pas dérivable en zéro.

A.5 Cadre de la Thèse et Contributions

A.5.1 Schémas riemanniens

Dans le chapitre 2, on étudie le schéma standard géodésique (A.6) avec la possibilité supplémentaire d'avoir un terme de biais en plus de l'estimateur non biaisé H , ce qui veut dire que l'on a

$$\theta_{n+1} = \text{Exp}_{\theta_n} [\eta_{n+1} (H_{\theta_n} (X_{n+1}) + b_{\theta_n} (X_{n+1}))] ,$$

où $b : \Theta \times X \rightarrow T\Theta$ est une fonction bi mesurable. On pense être les premiers à permettre ce terme de biais en approximation stochastique riemannienne. On a choisi d'avoir un biais uniformément borné, *i.e.* $\sup_{(\theta,x) \in \Theta \times X} \|b_{\theta}(x)\|_{\theta} \leq b_{\infty}$ voir A.3. On s'est aussi intéressés à des schéma de rétraction plus généraux, de la forme (A.7). L'analyse de ces schémas est très important parce que les rétractions sont souvent plus simples à calculer en pratique et ils font partie de la littérature depuis longtemps – voir [AM12]. L'exemple le plus commun est d'analyser un schéma de projection en tant que rétraction comme dans (A.5) pour la sphère.

On contribue à l'analyse des schémas de rétraction dans la Section 2.3.1 en considérant des rétraction de premier et second ordre (voir R.2 et R.3 respectivement). Pour ces deux types de rétractions, on quantifie l'erreur qui provient de la rétraction. Comme les rétractions sont des approximations de l'exponentielle riemannienne, on considère la notion de *fonction d'erreur de rétraction* dans (2.13). Il s'agit de la fonction $\Phi_{\theta} = \text{Exp}_{\theta}^{-1} \circ \text{Ret}_{\theta} : T_{\theta}\Theta \rightarrow T_{\theta}\Theta$ définie pour tout $\theta \in \Theta$. Ces applications définissent un morphisme de fibrés $\Phi : T\Theta \rightarrow T\Theta$, que l'on compare à l'identité. Dans les Lemmes 2.7 et 2.17, on borne uniformément $\theta \mapsto \|\Phi_{\theta} - \text{Id}_{\theta}\|_{\theta}$ en bornant les termes du développement de Taylor. On vérifie aussi que les rétractions usuelles sur la sphère et la variété de Grassmann sont des rétractions d'ordre un et deux : voir l'exemple 2.8 et la proposition 2.18 pour la sphère et l'exemple 2.9 et la proposition 2.10 pour la variété de Grassmann.

Dans le chapitre 3, on s'intéresse au cas du pas *constant*, avec l'hypothèse nouvelle de considérer un schéma projectif sur une variété riemannienne sur un ensemble fermé et géodésiquement convexe en utilisant un opérateur de pseudo-projection : voir (3.2) et A.6. Ceci contient le cas du schéma géodésique standard (1.6) sous les hypothèses A.6-(ii).

A.5.2 Hypothèses sur le bruit

L'hypothèse classique sur le bruit est le cas de différence de martingale (que l'on choisit de raccourcir en cas martingale). L'hypothèse A.22 se réécrit agréablement dans le cas riemannien.

A.24 (Martingale difference setting). *La suite $(e_{\theta_n}(X_{n+1}))_{n \in \mathbb{N}}$ est une différence de martingale par rapport à la filtration $(\mathcal{F}_n)_{n \in \mathbb{N}^*}$, *i.e.* $\mathbb{E}[e_{\theta_n}(X_{n+1}) | \mathcal{F}_n] = 0_{\theta_n}$ pour tout $n \in \mathbb{N}$. De plus, il existe $\sigma_0^2, \sigma_1^2 \geq 0$ tel que pour tout $n \in \mathbb{N}$*

$$\mathbb{E} \left[\|e_{\theta_n}(X_{n+1})\|_{\theta_n}^2 \middle| \mathcal{F}_n \right] \leq \sigma_0^2 + \sigma_1^2 \|h(\theta_n)\|_{\theta_n}^2 .$$

Dans les chapitres 2 et 3, on étudie ce cas **A24**. Cependant, dans le chapitre 2, on est les premiers à généraliser l'hypothèse markovienne au cadre riemannien.

A25 (Cas markovien). *Il existe un noyau de Markov P sur $(\Theta \times \mathsf{X}) \times \mathcal{X}$ tel que pour tout $n \in \mathbb{N}$ et toute fonction mesurable et bornée $\varphi : \mathsf{X} \rightarrow \mathbb{R}_+$, $\mathbb{E}[\varphi(X_{n+1}) | \mathcal{F}_n] = \int_{\mathsf{X}} \varphi(y) P_{\theta_n}(X_n, dy)$. De plus, pour tout $\theta \in \Theta$, P_θ admet une unique loi invariante π_θ , qui vérifie en plus $h(\theta) = \int_{\mathsf{X}} H_\theta(y) d\pi_\theta(y)$.*

Dans ce contexte, le bruit $(X_n)_{n \in \mathbb{N}^*}$ est une chaîne de Markov, avec un contrôle sur les transitions à travers le paramètre courant θ_n . Dans beaucoup d'applications concrètes, on ne peut pas tirer des variables aléatoires indépendantes de loi π_θ , définie en (3.1). En tout cas, il reste possible de simuler des échantillons avec une chaîne de Markov ergodique, de loi invariante π_θ . Nos hypothèses dans le cadre markovien sont **MA2**(w), et on présente deux cas pratiques où ces conditions sont vérifiées. On rappelle qu'une partie $C \in \mathcal{X}$ est appelée *small set* pour un noyau de Markov Q sur $\mathsf{X} \times \mathcal{X}$ quand il existe $m > 0$ et une mesure non triviale ν_m sur \mathcal{X} telle que pour tout $x \in C$, $B \in \mathcal{X}$ on a $Q^m(x, B) \geq \nu_m(B)$.

A 26 (w). *La famille de noyaux de Markov $(P_\theta)_{\theta \in \Theta}$ et la fonction mesurable $w : \mathsf{X} \rightarrow [1, +\infty)$ vérifient l'une des deux hypothèses suivantes :*

- (a) *La famille $(P_\theta)_{\theta \in \Theta}$ est uniformément ergodique avec une constante uniforme en θ , i.e. il existe $\varepsilon_P \in (0, 1)$ et $C_P \geq 0$ tel que pour tout $\theta \in \Theta$, P_θ admet une unique loi invariante π_θ et pour tout $x, x' \in \mathsf{X}$ et $k \in \mathbb{N}$, on a $\|\delta_x P_\theta^k - \delta_{x'} P_\theta^k\|_{\text{TV}} \leq C_P (1 - \varepsilon_P)^k$. De plus, il existe $C \geq 0$ tel que pour tout $\theta_0, \theta_1 \in \Theta$, $x \in \mathsf{X}$ et toute courbe géodésique $\gamma : [0, 1] \rightarrow \Theta$ reliant θ_0 et θ_1 , on a $\|\delta_x P_{\theta_1} - \delta_x P_{\theta_0}\|_{\text{TV}} \leq C \ell(\gamma)$.*
- (b) *Pour tout $\theta \in \Theta$, $P_\theta = P$, où P est un noyau de Markov sur $\mathsf{X} \times \mathcal{X}$. De plus, il existe $w : \mathsf{X} \rightarrow [1, +\infty)$, $\lambda \in (0, 1)$, $b_w \geq 0$ tel que pour tout $x \in \mathsf{X}$, $Pw(x) \leq \lambda w(x) + b_w \mathbb{1}_C(x)$ avec $C \in \mathcal{X}$ un *small set* pour P .*

Dans la proposition 2.6, on démontre que sous des hypothèses raisonnables, si **A26**-(a) sont vérifiées – i.e. la chaîne est uniformément ergodique, avec une constante d'ergodicité uniforme en $\theta \in \Theta$ – alors **MA2**(w) est vraie aussi, avec une fonction w constante $w : x \mapsto 1$. Par ailleurs, si **A26**(w)-(b) sont vraies – i.e. le noyau de Markov P ne dépend pas de θ et il y a une condition de drift – alors l'hypothèse **MA2**(w) est vérifiée. On pense que cette hypothèse sur le bruit est intéressante puisqu'elle permet de modéliser des cas de figure en dehors du cas martingale.

Pour chaque hypothèse de bruit, on fait attention à ne pas supposer que les itérations restent dans un compact arbitraire K , contrairement à [Bon13; ZS16], pour les raisons mentionnées dans Appendix A.2.2.

A.5.3 Techniques de convergence

Dans le chapitre 2, on généralise la deuxième méthode de Lyapunov – voir [LaS60] – dans le cadre riemannien. Notre hypothèse **A 2** est vérifiée si la fonction de Lyapunov V est

gradient Lipschitz ou elle a une hessienne riemannienne bornée en norme d'opérateur – voir le lemme 2.38.

On utilise les mêmes hypothèses dans le chapitre 3, où l'on échange **A7** avec **A6-A7**. Le problème avec les fonctions de Lyapunov sur les variétés riemanniennes et que le candidat par défaut $\theta \mapsto \rho_{\Theta}^2(\theta, \theta^*)$ – la distance à la solution au carré – n'est pas gradient Lipschitz lorsque Θ n'a pas une courbure nulle ou n'est pas compact. On propose une réponse à ce problème dans la Section 3.2.2, où on introduit une nouvelle fonction de Lyapunov, dite "Huberisée" – voir la proposition 3.5.

A.5.4 Résultats des Chapitres 2 et 3

Dans le but d'enseigner la géométrie riemannienne pour l'approximation stochastique à la communauté d'apprentissage statistique, pour reformuler des formules existant dans d'autres contextes et pour utiliser des preuves simples et utilisant un minimum d'outils, mes coauteurs et moi avons écrit des exemples, définitions et théorèmes sur la géométrie riemannienne dans les chapitres 2 et 3. On conseille à tout lecteur qui souhaite se familiariser avec les concepts et les notations de géométrie riemannienne utilisés dans cette thèse de commencer par lire le préliminaire écrit point par point en Section 2.11.

Dans la Section 2.6.2, on détaille le cas de la sphère qui a été exposé dans les exemples A.6 et 2.8. On utilise les notions introduites en Section 2.11 pour vérifier que la rétraction projective sur la sphère vérifie les hypothèses **R1-R2-R3** – see Proposition 2.18.

Deux des concepts que l'on a dû manipuler minutieusement sont les formules de Taylor (voir les théorèmes 2.37 et 3.29) et le calcul du cut locus dans le cas de la variété de Grassmann (voir la Proposition 2.26). La construction de formules de Taylor a été à la fois difficile et précieuse pour notre technique d'analyse. On a écrit la Section 3.11.2 pour démontrer la formule de Taylor (3.96) du théorème 3.29, en particulier parce qu'on avait besoin de la Proposition 3.30 que l'on n'a pas trouvé dans la littérature. Il s'agit d'une formule de Taylor à l'ordre n avec reste intégral, pour des champs de vecteurs lisses le long d'une géodésique. On a utilisé la notion de dérivée totale covariante d'ordre quelconque pour l'établir. Une grande partie de la Section 3.11.2 a été écrite en adaptant les résultats classiques des connexions au cadre particulier des champ de vecteurs *le long d'une géodésique*. Si le lecteur n'est pas habitué à travailler avec des champs de tenseurs et des connexions, on conseille de lire [Lee19, Chapter 4], puisqu'on s'en est inspirés pour écrire la Section 3.11.2. Après avoir démontré cette formule de Taylor, on l'a adaptée au cadre plus simple des formules de Taylor pour des fonctions à valeurs réelles dans la Section 2.11.9 et le théorème 2.37.

Nos dernières contributions qui concernent uniquement la géométrie riemannienne sont dans la présentation auto contenue de la variété de Grassmann. Notre but principal était d'exprimer les formules de l'exponentielle riemannienne décrites dans [AMS04; BZA20; EAS98] et en plus de calculer le cut locus. Pour faciliter la tâche du lecteur, voici le résumé de la Section 2.7.2. On commence par donner la définition de la variété de Grassmann comme un quotient de la variété non-compacte de Stiefel par l'action du groupe linéaire. Ensuite, pour donner la structure différentielle de cette variété, on commence par construire la fonction dite de *cross section* $\sigma_{\mathbf{B}}$ décrite dans le lemme 2.21. L'inverse

de cette application est utilisée pour écrire explicitement la famille des paramétrisations – ou de façon équivalente l’atlas lisse – qui définit la structure différentielle de la variété de Grassmann, voir la Proposition 2.22.

On continue cette partie en décrivant l’espace tangent en utilisant la décomposition canonique des variétés quotient. On rappelle la définition d’une connexion au sens de [KN96, Chapter II, Section 1] dans la définition 2.24, et on vérifie que (2.71) est une connexion dans ce sens et, avec Proposition 2.23, on obtient une description de l’espace tangent. Avec ces outils, on donne une formulation alternative de l’exponentielle riemannienne dans le lemme 2.25. Cette formule est utilisée pour le dernier résultat de la section, la Proposition 2.26, dans laquelle on calcule le cut locus de la variété de Grassmann. Cette preuve est le résultat d’un long travail et malgré la difficulté et la longueur des preuves, il reste plus simple que l’autre preuve que l’on connaît de ce résultat – voir [Sak76]. En effet, la preuve de [Sak76] utilise des outils plus compliqués de géométrie riemannienne.

Dans le chapitre 2, on calcule des taux de convergence non asymptotiques globaux sous différentes hypothèses sur le bruit et la fonction de Lyapunov, données plus haut. Pour le schéma géodésique étudié dans la Section 2.2, on montre les théorèmes 2.4 et 2.5 qui affirment qu’après n itérations, $\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] = \mathcal{O}(b_\infty + \log n/\sqrt{n})$, où b_∞ est la borne uniforme sur le biais de H et I_n est une variable aléatoire discrète à valeurs sur $\{1, \dots, n\}$. Une discussion plus fine sur le choix des pas et leurs conséquences est donnée après les énoncés de ces théorèmes. Pour les schémas de rétraction étudiés dans la Section 2.3.2, on montre dans les théorèmes 2.11, 2.12 et 2.29 qu’après n itérations, comme dans le cas géodésique, $\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] = \mathcal{O}(b_\infty + \log n/\sqrt{n})$. L’intérêt de cette analyse comparée à celle de [ZS16] est que 1) on n’a pas besoin que les itérations restent dans un compact, 2) on montre ces bornes pour l’algorithme d’approximation stochastique, qui est plus général que le cas de descente de gradient stochastique et 3) les hypothèses faites sur la fonction de Lyapunov permettent de traiter des fonctions qui ne sont pas convexes.

Hypothèses			Résultat
de Base	sur le Schéma	sur le Bruit	
A1-A2-A3	Géodésique (2.3)	Cas martingale MD1	Theorem 2.4
A1-A2-A3	Géodésique (2.3)	Cas markovien MA1-MA2(w)	Theorem 2.5
A1-A2-A3 A4	Rétraction (2.12) R1 et (R2 ou R3)	Cas martingale MD1 et (MD2(4) ou MD2(6))	Theorem 2.11
A1-A2-A3 A4	Rétraction (2.12) R1 et (R2 ou R3)	Cas markovien MA1-MA2(w)-MA4(w)	Theorem 2.12

Table A.1: Résumé des résultats non asymptotiques avec toutes les hypothèses possibles.

Finalement, on considère deux exemples d’applications pour ces résultats. Premièrement, on introduit un problème d’ACP en ligne, similaire au problème introduit dans l’exemple 1.10 mais dans le cas général $r \geq 1$. Dans la Section 2.4.1, on montre que les

schémas itératifs (2.19) et (2.20) vérifient les conditions de nos théorèmes – voir le théorème 2.13. On complète cette analyse par des simulations numériques – voir la Section 2.4.3. Deuxièmement, on s’intéresse à un problème d’estimation de barycentre en ligne dans la Section 2.4.2. Dans cette application, on examine et raffine l’idée initiale de l’exemple 1.1 pour définir le concept de barycentre riemannien robuste pour une loi de probabilités sur une variété de Hadamard. On démontre l’existence et l’unicité de ce barycentre dans la Proposition 2.14 et on en déduit un schéma itératif pour le calculer (2.23), qui vérifie les conditions générales de nos théorèmes – voir la Proposition 2.15.

Dans le chapitre 3, on regarde de plus près le cas du pas constant. On commence par trouver des bornes non asymptotiques dans le théorème 3.1, qui ressemble au théorème 2.4. Néanmoins, on trouve une meilleure borne pour le cas particulier (3.6), où l’on observe l’oubli à vitesse exponentielle des conditions initiales, comme pour la plupart des fonctions fortement convexes dans le cadre euclidien. Dans le théorème 3.2, on montre que la suite $(\theta_n)_{n \in \mathbb{N}}$ est une chaîne de Markov géométriquement ergodique, avec une unique loi invariante μ^η qui dépend du choix du pas η . À partir de là, on cherche à comprendre le comportement de μ^η lorsque $\eta \rightarrow 0$, avec des résultats asymptotiques et non asymptotiques. Le premier pas dans cette direction est le théorème Theorem 3.3, où l’on démontre que $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$ converge faiblement vers la mesure de Dirac en θ^* , où θ^* est l’unique solution du problème (3.1). Dans la Section 3.3, en étudiant le schéma non contraint (1.6), on quantifie la convergence trouvée dans le théorème 3.3-(b). Tout d’abord, on montre un développement asymptotique du biais $\int_{\Theta} g(\theta) d\mu^\eta(\theta) - g(\theta^*)$ par rapport au pas η pour toute fonction g dans une certaine classe de fonctions lisses de Θ vers \mathbb{R} – voir le théorème 3.6. En appliquant ce résultat à la descente de gradient stochastique ($h = -\text{grad } f$), on trouve un développement asymptotique de $\int_{\Theta} \|\text{grad } f(\theta)\|_{\theta}^2 d\mu^\eta(\theta)$. Deuxièmement, on montre que la convergence établie dans le théorème 3.3-(b) se fait à une vitesse $\eta^{1/2}$, en démontrant un théorème central limite pour $(\mu^\eta)_{\eta \in (0, \bar{\eta}]}$ – voir le théorème 3.7. Ces deux résultats peuvent être vus comme une décomposition biais-variance dans laquelle les deux termes sont d’ordre η . Il s’agit donc d’une contrepartie plus faible de [DDB17, Proposition 3, Theorem 5] et [Pff86, Theorem 1] dans le cadre riemannien. Ensuite, dans la Section 3.4, on applique ces résultats pour la descente de gradient stochastique. On considère deux hypothèses sur la fonction f . Lorsque f est fortement géodésiquement convexe – voir la définition 1.12 et l’hypothèse **F2** – on a les résultats du corolaire 3.9. On voit que, comme dans le cas euclidien, il faut $\mathcal{O}(\log(1/\varepsilon)\varepsilon^{-1})$ étapes pour avoir une précision de taille ε , *i.e.* $\mathbb{E}[f(\theta_n) - f(\theta^*)] \leq \varepsilon$. Par ailleurs, on considère aussi le cas où f est géodésiquement quasi convexe, ce qui veut dire dans le cas euclidien que $\langle \text{grad } f(\theta), \theta - \theta^* \rangle \geq \tau[f(\theta) - f(\theta^*)]$, comme défini dans [HMR19]. On obtient des résultats plus faibles dans la Proposition 3.11. Ils sont plus faibles non seulement parce que le membre de droite de l’inégalité est $C + \mathcal{O}(n^{-1})$ – ce qui ressemble beaucoup aux résultats du chapitre 2 – mais aussi parce que le membre de gauche s’exprime avec la fonction “pseudo distance” définie en (3.14). L’aspect plus novateur du résultat est dans ce deuxième cas, où nous ne nécessitons pas d’hypothèses de borne sur Θ . Pour finir, on applique ces résultats au problème d’estimation de barycentre. Pour utiliser les deux hypothèses de convexité, on s’intéresse à l’estimation du barycentre d’une loi discrète – ce

qui définit une fonction fortement convexe f – et d’une loi continue – ce qui définit une fonction géodésiquement quasi convexe. Le deuxième cas est nouveau dans le domaine puisqu’il n’y a pas d’hypothèses de borne sur le support de la loi. Pour le deuxième cas, on introduit un nouveau schéma (3.17), pour lequel on montre des bornes non asymptotiques dans le théorème 3.13. L’estimation de barycentres est illustré par des simulations numériques dans la Section 3.5. On utilise les deux jeux d’hypothèses présentés ci-dessus.

A.5.5 Résultats du chapitre 4

Dans le chapitre 4, on s’intéresse à une généralisation de MALA, qui a été étudiée pour la première fois dans [Per16]. Cette généralisation veut affaiblir l’hypothèse de dérivabilité de la loi cible π , lorsqu’elle s’écrit sous la forme (A.16). Si $\pi \propto \exp(-G)$, alors on suppose que la fonction $G : \mathbb{R}^d \rightarrow \mathbb{R}$ est convexe, propre et semi continue inférieurement. Au lieu d’utiliser le vrai gradient de G dans la proposition, on approxime G par son enveloppe de Moreau-Yoshida $G^\lambda : \mathbb{R}^d \rightarrow \mathbb{R}$, définie pour tout $\lambda \geq 0$ par

$$G^\lambda(\mathbf{x}) = \min_{\mathbf{u} \in \mathbb{R}^d} [G(\mathbf{u}) + \|\mathbf{u} - \mathbf{x}\|^2 / (2\lambda)] .$$

En posant $\pi^\lambda = \exp(-G^\lambda)$, on obtient une approximation lisse de π , où λ contrôle la régularité de π^λ . Sous ces hypothèses, G^λ est continûment différentiable et on peut donc utiliser MALA pour la cible π^λ , où la proposition est de la forme

$$\xi_{k+1} = \xi_k - \frac{\sigma^2}{2} \nabla G^\lambda(\xi_k) + \sigma Z_{k+1} ,$$

voir (4.4) pour une expression du gradient de G^λ . Notre but est de produire des résultats théoriques sur cette classe d’algorithmes, que l’on appelle proximal MALA, d’étudier son comportement asymptotique lorsque la dimension d tend vers l’infini. Nos résultats apportent aussi des conseils pour l’implémentation de ces algorithmes proximaux.

On commence par s’intéresser au cas particulier où π est différentiable – voir la Section 4.3.1 – pour comparer les algorithmes qui utilisent le gradient (comme MALA) avec les algorithmes proximaux. Pour notre analyse, on fait dépendre le pas de discrétisation σ^2 et le paramètre de régularisation λ de la dimension d de la cible. Le théorème 4.1 énonce qu’il y a deux possibilités d’échelonnement. Le premier, dans les cas (b) et (c), arrive lorsque λ tend vers 0 au moins aussi vite que σ^2 . Dans ce cas, l’échelonnement est équivalent à celui de MALA, avec σ^2 proportionnel à $d^{-1/3}$. Cependant, lorsque λ tend vers 0 plus lentement que σ^2 , on trouve un nouvel échelonnement, dans lequel σ^2 est proportionnel à $d^{-1/2}$ – voir (a). Dans ce deuxième cas, proximal MALA est moins efficace que MALA. Dans les deux cas, la convergence du taux d’acceptation $a_d(\ell, r) \xrightarrow{d \rightarrow \infty} a(\ell, r)$ est accompagnée de la convergence faible de l’interpolation linéaire de la première composante de la chaîne $(X_k^d)_{k \geq 0}$, définie par

$$L_t^d = \left(\lceil d^{2\alpha} t \rceil - d^{2\alpha} t \right) X_{\lfloor d^{2\alpha} t \rfloor, 1}^d + \left(d^{2\alpha} t - \lfloor d^{2\alpha} t \rfloor \right) X_{\lceil d^{2\alpha} t \rceil, 1}^d , \quad (\text{A.21})$$

vers la diffusion de Langevin

$$dL_t = h(\ell, r)^{1/2} dB_t - \frac{h(\ell, r)}{2} g'(x) dt ,$$

où $(B_t)_{t \geq 0}$ est le mouvement brownien standard sur \mathbb{R} et $h(\ell, r) = \ell^2 a(\ell, r)$ est la vitesse de la diffusion. Bien que proximal MALA soit au mieux aussi efficace que MALA, comme parfois il est plus simple de calculer l'opérateur proximal que de calculer $\nabla \log \pi$ pour certaines lois π [CP11; PB14; Per16], il nous semble qu'il y a des cas où un bon choix de λ permet à proximal MALA d'être un concurrent de MALA en termes de temps de calcul pour une même précision. Ensuite, nous étudions proximal MALA pour une cible particulière qui n'est pas lisse en tout point, la loi de Laplace de densité $x \mapsto 2^{-1} \exp(-|x|)$. Même si nous démontrons nos résultats pour le cas particulier de la loi de Laplace, notre technique de preuve peut s'adapter à d'autres lois, tant que l'on soit capable de calculer certaines intégrales bien choisies – voir la Section 4.10. On choisit cette loi en particulier parce qu'elle est utilisée en pratique, par exemple en déconvolution d'images image deconvolution [DMP18; Per16; VDC19], pour le LASSO, et la régression creuse [Atc15; GSS21; ZCZ22]. On obtient la convergence du taux d'acceptation dans le théorème 4.4. On observe qu'un seul point de non dérivabilité est suffisant pour changer l'échelonnement par rapport à MALA pour des cibles lisses. En effet, pour la loi de Laplace, le pas σ^2 doit être proportionnel à $d^{-2/3}$ (tandis que MALA pour des cibles lisses nécessite $d^{-1/3}$) pour obtenir une limite non triviale du taux d'acceptation. Cet échelonnement optimal est nouveau parmi les algorithmes MCMC, se trouvant entre le taux d^{-1} de RWM et le taux $d^{-1/3}$ de MALA (pour des cibles lisses). Pour cette étude, on s'intéresse uniquement au cas où λ converge vers 0 au moins aussi vite que σ^2 . Comme d'habitude pour l'échelonnement optimal, on poursuit l'étude avec la convergence faible de l'interpolation de la première composante de la chaîne $(L_t^d)_{t \geq 0}$, définie en (A.21), vers la diffusion de Langevin. On procède en deux étapes, en montrant d'abord la tension de la suite de lois des processus $\{(L_t^d)_{t \geq 0} : d \in \mathbb{N}\}$ – voir la 4.5, et ensuite on démontre que la seule limite possible pour cette suite est effectivement solution de la diffusion de Langevin

$$dL_t = [h^L(\ell)]^{1/2} dB_t - \frac{h^L(\ell)}{2} \operatorname{sgn}(L_t) dt .$$

On a trouvé ce résultat en utilisant le problème de martingale associé à la diffusion – voir la Proposition 4.6. Notre dernier résultat, le théorème 4.7, combine ces deux énoncés pour trouver l'échelonnement optimal de l'algorithme proximal MALA avec cible Laplace. Il est important d'insister que l'on peut étendre la définition de proximal MALA au cas $\lambda = 0$, ce qui transforme proximal MALA en MALA sous-gradient (que l'on raccourcit en sG-MALA). Tous nos résultats sur la cible Laplace sont vrais aussi pour cette généralisation.

On complète nos résultats avec une discussion sur les implications pratiques de nos théorèmes – voir la Section 4.4 – ainsi que par des expériences numériques afin d'illustrer l'échelonnement de cibles différentiables et Laplace pour proximal MALA, avec différents choix des taux de convergence vers zéro des paramètres σ^2 et λ .

Bibliography

- [AMS08] P. A. Absil, R. Mahony et al. *Optimisation algorithms on matrix manifolds*. Princeton University Press, 2008.
- [AM12] P. A. Absil and J. Malick. ‘Projection-like retractions on matrix manifolds’. In: *SIAM Journal on Optimization* 22.1 (2012), pp. 135–158.
- [AMS04] P.-A. Absil, R. Mahony et al. ‘Riemannian geometry of Grassmann manifolds with a view on algorithmic computation’. In: *Acta Appl. Math.* 80.2 (2004), pp. 199–220.
- [Aga+09] Alekh Agarwal, Martin J Wainwright et al. ‘Information-theoretic lower bounds on the oracle complexity of convex optimization’. In: *Advances in Neural Information Processing Systems* 22 (2009).
- [Agr+21] Sanket Agrawal, Dootika Vats et al. ‘Optimal Scaling of MCMC Beyond Metropolis’. In: *arXiv preprint arXiv:2104.02020* (2021).
- [AS20] Kwangjun Ahn and Suvrit Sra. ‘From Nesterov’s estimate sequence to Riemannian acceleration’. In: *Conference on Learning Theory*. PMLR, 2020, pp. 84–118.
- [Ali+20] F. Alimisis, A. Orvieto et al. ‘A Continuous-time Perspective for Modeling Acceleration in Riemannian Optimization’. In: ed. by Silvia Chiappa and Roberto Calandra. Vol. 108. *Proceedings of Machine Learning Research*. Online: PMLR, 26–28 Aug 2020, pp. 1297–1307.
- [Ali+21] F. Alimisis, A. Orvieto et al. ‘Momentum Improves Optimization on Riemannian Manifolds’. In: *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*. Ed. by Arindam Banerjee and Kenji Fukumizu. Vol. 130. *Proceedings of Machine Learning Research*. PMLR, 13–15 Apr 2021, pp. 1351–1359.
- [ABY13] M. Arnaudon, F. Barbaresco et al. ‘Riemannian Medians and Means With Applications to Radar Signal Processing’. In: *IEEE Journal of Selected Topics in Signal Processing* 7.4 (2013), pp. 595–604.
- [Arn+12] M. Arnaudon, C. Dombry et al. ‘Stochastic algorithms for computing means of probability measures’. In: *Stochastic Processes and their Applications* 122.4 (2012), pp. 1437–1455.

-
- [Atc15] Yves F Atchadé. ‘A Moreau-Yosida approximation scheme for a class of high-dimensional posterior distributions’. In: *arXiv preprint arXiv:1505.07072* (2015).
- [Bac20] F. Bach. *On the Effectiveness of Richardson Extrapolation in Machine Learning*. 2020. arXiv: 2002.02835 [cs.LG].
- [BM11] F. Bach and E. Moulines. ‘Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning’. In: *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*. 2011, pp. 451–459.
- [BC11] Heinz H Bauschke and Patrick L Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Vol. 408. Springer, 2011.
- [BB01] J. Baxter and P. L. Bartlett. ‘Infinite-Horizon Policy-Gradient Estimation’. In: *J. Artif. Int. Res.* 15.1 (Nov. 2001), pp. 319–350.
- [BR08] Mylene Bédard and Jeffrey S Rosenthal. ‘Optimal scaling of Metropolis algorithms: Heading toward general target distributions’. In: *Canadian Journal of Statistics* 36.4 (2008), pp. 483–503.
- [BZA20] Thomas Bendokat, Ralf Zimmermann et al. *A Grassmann Manifold Handbook: Basic Geometry and Computational Aspects*. 2020.
- [BMP90] A. Benveniste, M. Métivier et al. *Adaptive algorithms and stochastic approximations*. Vol. 22. Applications of Mathematics (New York). Translated from the French by Stephen S. Wilson. Springer-Verlag, Berlin, 1990, pp. xii+365.
- [BRS09] Alexandros Beskos, Gareth Roberts et al. ‘Optimal scalings for local Metropolis–Hastings chains on nonproduct targets in high dimensions’. In: *The Annals of Applied Probability* 19.3 (2009), pp. 863–898.
- [BFH13] Pascal Bianchi, Gersende Fort et al. ‘Performance of a distributed stochastic approximation algorithm’. In: *IEEE Transactions on Information Theory* 59.11 (2013), pp. 7405–7418.
- [BFR19] Joris Bierkens, Paul Fearnhead et al. ‘The zig-zag process and super-efficient sampling for Bayesian analysis of big data’. In: *The Annals of Statistics* 47.3 (2019), pp. 1288–1320.
- [BI13] D. A. Bini and B. Iannazzo. ‘Computing the Karcher mean of symmetric positive definite matrices’. In: *Linear Algebra and its Applications* 438.4 (2013). 16th ILAS Conference Proceedings, Pisa 2010, pp. 1700–1710.
- [Bon13] Silvere Bonnabel. ‘Stochastic gradient descent on Riemannian manifolds’. In: *IEEE Transactions on Automatic Control* 58.9 (2013), pp. 2217–2229.
- [Bor09] Vivek S Borkar. *Stochastic approximation: a dynamical systems viewpoint*. Vol. 48. Springer, 2009.
- [BM00] Vivek S Borkar and Sean P Meyn. ‘The ODE method for convergence of stochastic approximation and reinforcement learning’. In: *SIAM Journal on Control and Optimization* 38.2 (2000), pp. 447–469.
-

- [Bot10] L. Bottou. ‘Large-Scale Machine Learning with Stochastic Gradient Descent’. In: *Proceedings of the 19th International Conference on Computational Statistics (COMPSTAT’2010)*. Ed. by Yves Lechevallier and Gilbert Saporta. Paris, France: Springer, Aug. 2010, pp. 177–187.
- [Bot12] L. Bottou. ‘Stochastic gradient descent tricks’. In: *Neural networks: Tricks of the trade*. Springer, 2012, pp. 421–436.
- [BB08a] L. Bottou and O. Bousquet. ‘The Tradeoffs of Large Scale Learning’. In: *Advances in Neural Information Processing Systems 20 (NIPS 2007)*. Ed. by J.C. Platt, D. Koller et al. NIPS Foundation (<http://books.nips.cc>), 2008, pp. 161–168.
- [BB08b] L. Bottou and O. Bousquet. ‘The tradeoffs of large scale learning’. In: *Advances in Neural Information Processing Systems*. 2008.
- [BCN18] Léon Bottou, Frank E Curtis et al. ‘Optimization methods for large-scale machine learning’. In: *Siam Review* 60.2 (2018), pp. 223–311.
- [BH13] Nawaf Bou-Rabee and Martin Hairer. ‘Nonasymptotic mixing of the MALA algorithm’. In: *IMA Journal of Numerical Analysis* 33.1 (2013), pp. 80–110.
- [BVD18] Alexandre Bouchard-Côté, Sebastian J Vollmer et al. ‘The bouncy particle sampler: A nonreversible rejection-free Markov chain Monte Carlo method’. In: *Journal of the American Statistical Association* 113.522 (2018), pp. 855–867.
- [BA11] N. Boumal and P.-A. Absil. ‘RTRMC: A Riemannian trust-region method for low-rank matrix completion’. In: *Advances in neural information processing systems*. 2011, pp. 406–414.
- [BVB16] N. Boumal, V. Voroninski et al. ‘The non-convex Burer-Monteiro approach works on smooth semidefinite programs’. In: *Advances in Neural Information Processing Systems*. 2016, pp. 2757–2765.
- [Bou20] Nicolas Boumal. *An introduction to optimization on smooth manifolds*. Available online. Aug. 2020.
- [BAC18] Nicolas Boumal, Pierre-Antoine Absil et al. ‘Global rates of convergence for nonconvex optimization on manifolds’. In: *IMA Journal of Numerical Analysis* 39.1 (2018), pp. 1–33.
- [BCZ11] Mark Broadie, Deniz Cicek et al. ‘General bounds and finite-time improvement for the Kiefer-Wolfowitz stochastic approximation algorithm’. In: *Operations Research* 59.5 (2011), pp. 1211–1224.
- [Bro+11] Steve Brooks, Andrew Gelman et al. *Handbook of Markov chain Monte Carlo*. CRC press, 2011.
- [CM09] O. Cappé and E. Moulines. ‘On-line expectation-maximization algorithm for latent data models’. In: *J. R. Stat. Soc. Ser. B Stat. Methodol.* 71.3 (2009), pp. 593–613.

-
- [CD18] Hervé Cardot and David Degras. ‘Online principal component analysis in high dimension: Which algorithm to choose?’ In: *International Statistical Review* 86.1 (2018), pp. 29–50.
- [Cha+20] R. Chakraborty, L. Yang et al. ‘Intrinsic Grassmann Averages for Online Linear, Robust and Nonlinear Subspace Learning’. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020), pp. 1–1.
- [CRR05] Ole F Christensen, Gareth O Roberts et al. ‘Scaling limits for the transient phase of local Metropolis–Hastings algorithms’. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67.2 (2005), pp. 253–268.
- [CP11] Patrick L Combettes and Jean-Christophe Pesquet. ‘Proximal splitting methods in signal processing’. In: *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Ed. by H. H. Bauschke, R. S. Burachik et al. Springer, 2011, pp. 185–212.
- [Cra+06] Koby Crammer, Ofer Dekel et al. ‘Online passive aggressive algorithms’. In: (2006).
- [CB19] C. Criscitiello and N. Boumal. ‘Efficiently escaping saddle points on manifolds’. In: *NeurIPS*. 2019.
- [dST+21] Alexandre d’Aspremont, Damien Scieur et al. ‘Acceleration methods’. In: *Foundations and Trends® in Optimization* 5.1-2 (2021), pp. 1–245.
- [Del00] Bernard Delyon. ‘Stochastic approximation with decreasing gain: Convergence and asymptotic theory’. In: *Unpublished lecture notes, Université de Rennes* 26 (2000).
- [DLR77] Arthur P Dempster, Nan M Laird et al. ‘Maximum likelihood from incomplete data via the EM algorithm’. In: *Journal of the royal statistical society: series B (methodological)* 39.1 (1977), pp. 1–22.
- [DDB17] A. Dieuleveut, A. Durmus et al. *Bridging the Gap between Constant Step Size Stochastic Gradient Descent and Markov Chains*. 2017. arXiv: 1707.06386 [stat.ML].
- [Dou+18] Randal Douc, Eric Moulines et al. *Markov chains*. Springer, 2018.
- [DHS11] J. Duchi, E. Hazan et al. ‘Adaptive Subgradient Methods for Online Learning and Stochastic Optimization’. In: *Journal of Machine Learning Research* 12.61 (2011), pp. 2121–2159.
- [Duf97] M. Duflo. *Random Iterative Models*. 1st. Berlin, Heidelberg: Springer-Verlag, 1997.
- [Dur+20] A. Durmus, P. Jiménez et al. ‘Convergence Analysis of Riemannian Stochastic Approximation Schemes’. In: *arXiv preprint arXiv:2005.13284* (2020).
- [Dur+21] Alain Durmus, Pablo Jiménez et al. ‘On Riemannian Stochastic Approximation Schemes with Fixed Step-Size’. In: *Proceedings of Machine Learning Research*. PMLR, 2021.
-

- [Dur+17] Alain Durmus, Sylvain Le Corff et al. ‘Optimal scaling of the random walk Metropolis algorithm under L_p mean differentiability’. In: *Journal of Applied Probability* 54.4 (2017), pp. 1233–1260.
- [DM22] Alain Durmus and Éric Moulines. *On the geometric convergence for MALA under verifiable conditions*. 2022.
- [DMP18] Alain Durmus, Eric Moulines et al. ‘Efficient Bayesian computation by proximal Markov chain Monte Carlo: when Langevin meets Moreau’. In: *SIAM Journal on Imaging Sciences* 11.1 (2018), pp. 473–506.
- [Dur19] Rick Durrett. *Probability: Theory and Examples*. Vol. 49. Cambridge University Press, 2019.
- [Dwi+18] Raaz Dwivedi, Yuansi Chen et al. ‘Log-concave sampling: Metropolis-Hastings algorithms are fast!’ In: *Proceedings of the 31st Conference On Learning Theory*. Ed. by Sébastien Bubeck, Vianney Perchet et al. Vol. 75. Proceedings of Machine Learning Research. PMLR, June 2018, pp. 793–797.
- [EAS98] A. Edelman, T. A. Arias et al. ‘The Geometry of Algorithms with Orthogonality Constraints’. In: *SIAM journal on Matrix Analysis and Applications* 20.2 (1998), pp. 303–353.
- [Erm75] D. L Ermak. ‘A computer simulation of charged particles in solution. I. Technique and equilibrium properties’. In: *The Journal of Chemical Physics* 62.10 (1975), pp. 4189–4196.
- [EK09] Stewart N Ethier and Thomas G Kurtz. *Markov processes: characterization and convergence*. Vol. 282. John Wiley & Sons, 2009.
- [Fab68] Vaclav Fabian. ‘On asymptotic normality in stochastic approximation’. In: *The Annals of Mathematical Statistics* (1968), pp. 1327–1332.
- [FP99] J.-C. Fort and G. Pagès. ‘Asymptotic Behavior of a Markovian Stochastic Algorithm with Constant Step’. In: *SIAM Journal on Control and Optimization* 37.5 (1999), pp. 1456–1482. eprint: <https://doi.org/10.1137/S0363012997328610>.
- [GHL04] S. Gallot, D. Hulin et al. *Riemannian geometry (Third Edition)*. Springer-Verlag, 2004.
- [GRG96] A Gelman, G O Roberts et al. ‘Efficient Metropolis jumping rules’. In: *Bayesian Statistics* (1996).
- [Geo14] H. Geoffrey. *Lecture 6e RMSprop: Divide the gradient by a running average of its recent magnitude*. 2014.
- [GL13] S. Ghadimi and G. Lan. ‘Stochastic First- and Zeroth-Order Methods for Nonconvex Stochastic Programming’. In: *SIAM Journal on Optimization* 23.4 (2013), pp. 2341–2368. eprint: <https://doi.org/10.1137/120880811>.
- [GV08] Alvina Goh and René Vidal. ‘Clustering and dimensionality reduction on Riemannian manifolds’. In: *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE. 2008, pp. 1–7.

-
- [GSS21] Jacob Vorstrup Goldman, Torben Sell et al. ‘Gradient-Based Markov Chain Monte Carlo for Bayesian Inference With Non-differentiable Priors’. In: *Journal of the American Statistical Association* (2021), pp. 1–12.
- [Gre83] U. Grenander. ‘Tutorial in pattern theory’. Division of Applied Mathematics, Brown University, Providence. 1983.
- [GM94a] U. Grenander and M. I. Miller. ‘Representations of knowledge in complex systems’. In: *J. Roy. Statist. Soc. Ser. B* 56.4 (1994). With discussion and a reply by the authors, pp. 549–603.
- [GM94b] Ulf Grenander and Michael I Miller. ‘Representations of knowledge in complex systems’. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 56.4 (1994), pp. 549–581.
- [HH03] B. Hall and B.C. Hall. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*. Graduate Texts in Mathematics. Springer, 2003.
- [HG20] Andi Han and Junbin Gao. *Variance reduction for Riemannian non-convex optimization with batch size adaptation*. 2020. arXiv: 2007.01494 [math.OC].
- [HMR19] Moritz Hardt, Tengyu Ma et al. *Gradient Descent Learns Linear Dynamical Systems*. 2019. arXiv: 1609.05191 [cs.LG].
- [Has70] W. K. Hastings. ‘Monte Carlo sampling methods using Markov chains and their application’. In: 57 (1970), pp. 97–109.
- [HP14] Fumio Hiai and Dénes Petz. *Introduction to matrix analysis and applications*. Springer Science & Business Media, 2014.
- [HGI07] Michael Holmes, Alexander Gray et al. ‘Fast SVD for large-scale matrices’. In: *Workshop on Efficient Machine Learning at NIPS*. Vol. 58. 2007, pp. 249–252.
- [HJ94] R. A. Horn and C. R. Johnson. *Topics in matrix analysis*. Cambridge university press, 1994.
- [HS19] Reshad Hosseini and Suvrit Sra. ‘An alternative to EM for gaussian mixture models: Batch and stochastic Riemannian optimization’. In: *Mathematical Programming* (2019), pp. 1–37.
- [HS20] Reshad Hosseini and Suvrit Sra. ‘Recent Advances in Stochastic Riemannian Optimization’. In: *Handbook of Variational Methods for Nonlinear Geometric Data* (2020), pp. 527–554.
- [IP18] B. Iannazzo and M. Porcelli. ‘The Riemannian Barzilai–Borwein method with nonmonotone line search and the matrix geometric mean computation’. In: *Ima Journal of Numerical Analysis* 38 (2018), pp. 495–517.
- [Ish+11] M. Ishteva, P-A Absil et al. ‘Best low multilinear rank approximation of higher-order tensors, based on the Riemannian trust-region scheme’. In: *SIAM Journal on Matrix Analysis and Applications* 32.1 (2011), pp. 115–135.
-

- [JJS93] T. Jaakkola, M. I. Jordan et al. ‘Convergence of Stochastic Iterative Dynamic Programming Algorithms’. In: *Proceedings of the 6th International Conference on Neural Information Processing Systems*. NIPS’93. Denver, Colorado: Morgan Kaufmann Publishers Inc., 1993, pp. 703–710.
- [JR07] Søren F Järner and Gareth O Roberts. ‘Convergence of heavy-tailed Monte Carlo Markov chain algorithms’. In: *Scandinavian Journal of Statistics* 34.4 (2007), pp. 781–815.
- [Jos05] J. Jost. *Riemannian Geometry and Geometric Analysis*. Springer Universitat texts. Springer, 2005.
- [JLM14] Benjamin Jourdain, Tony Lelièvre et al. ‘Optimal scaling for the transient phase of Metropolis Hastings algorithms: the longtime behavior’. In: *Bernoulli* 20.4 (2014), pp. 1930–1978.
- [Kal21] Olav Kallenberg. *Foundations of Modern Probability*. Springer, 2021.
- [KB60] Rudolf E Kalman and John E Bertram. ‘Control system analysis and design via the “second method” of Lyapunov: I—Continuous-time systems’. In: (1960).
- [Kar+19] B. Karimi, B. Miasojedow et al. ‘Non-asymptotic Analysis of Biased Stochastic Approximation Scheme’. In: *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*. 2019, pp. 1944–1974.
- [Ken78] J. Kent. ‘Time-reversible diffusions’. In: *Adv. in Appl. Probab.* 10.4 (1978), pp. 819–835.
- [KL17] Masoud Badiie Khuzani and Na Li. ‘Stochastic Primal-Dual Method on Riemannian Manifolds with Bounded Sectional Curvature’. In: *arXiv preprint arXiv:1703.08167* (2017).
- [KB17] D. P. Kingma and J. Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG].
- [KN96] S. Kobayashi and K. Nomizu. *Foundations of Differential Geometry, Volumes 1, 2*. A Wiley Publication in Applied Statistics. Wiley, 1996.
- [KH81] H. J. Kushner and H. Huang. ‘Asymptotic Properties of Stochastic Approximations with Constant Coefficients’. In: *SIAM Journal on Control and Optimization* 19.1 (1981), pp. 87–105. eprint: <https://doi.org/10.1137/0319007>.
- [KY03] H. J. Kushner and G. G. Yin. *Stochastic approximation and recursive algorithms and applications*. Second. Vol. 35. Applications of Mathematics (New York). Stochastic Modelling and Applied Probability. Springer-Verlag, New York, 2003, pp. xxii+474.
- [LSB12] Simon Lacoste-Julien, Mark Schmidt et al. ‘A simpler approach to obtaining an $o(1/t)$ convergence rate for the projected stochastic subgradient method’. In: *arXiv preprint arXiv:1212.2002* (2012).

-
- [LaS60] J. LaSalle. ‘Some extensions of Lyapunov’s second method’. In: *IRE Transactions on circuit theory* 7.4 (1960), pp. 520–527.
- [Le04] H. Le. ‘Estimation of Riemannian Barycentres’. In: *Lms Journal of Computation and Mathematics* 7 (2004), pp. 193–200.
- [Lee19] J. M. Lee. *Introduction to Riemannian Manifolds*. Springer International Publishing, 2019.
- [Lee03] J.M. Lee. *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer, 2003.
- [LZ22] Samuel Livingstone and Giacomo Zanella. ‘The Barker proposal: Combining robustness and efficiency in gradient-based MCMC’. In: *Journal of the Royal Statistical Society Series B* 84.2 (2022), pp. 496–523.
- [Lya92] Aleksandr Mikhailovich Lyapunov. ‘The general problem of the stability of motion’. In: *International journal of control* 55.3 (1992), pp. 531–534.
- [MCF15] Yi-An Ma, Tianqi Chen et al. ‘A complete recipe for stochastic gradient MCMC’. In: *Advances in Neural Information Processing Systems* 28 (2015).
- [MBB18] S. Ma, R. Bassily et al. ‘The Power of Interpolation: Understanding the Effectiveness of SGD in Modern Over-parametrized Learning’. In: *International Conference on Machine Learning*. 2018, pp. 3325–3334.
- [MMM18] Szymon Majewski, Błażej Miasojedow et al. ‘Analysis of nonsmooth stochastic approximation: the differential inclusion approach’. In: *arXiv preprint arXiv:1805.01916* (2018).
- [Mat+19] Emile Mathieu, Charline Le Lan et al. ‘Continuous Hierarchical Representations with Poincaré Variational Auto-Encoders’. In: *Advances in neural information processing systems*. 2019, pp. 12544–12555.
- [MT96] Kerrie L Mengersen and Richard L Tweedie. ‘Rates of convergence of the Hastings and Metropolis algorithms’. In: *Annals of Statistics* 24.1 (1996), pp. 101–121.
- [Met+53] N. Metropolis, A. W. Rosenbluth et al. ‘Equations of state calculations by fast computing machines’. In: *Journal of Chemical Physics* 23.2 (1953), pp. 1087–1092.
- [Mey00] C.D. Meyer. *Matrix Analysis and Applied Linear Algebra*. Other Titles in Applied Mathematics. SIAM, 2000.
- [MT09] S. Meyn and R. Tweedie. *Markov Chains and Stochastic Stability*. 2nd. New York, NY, USA: Cambridge University Press, 2009.
- [MB11] Eric Moulines and Francis Bach. ‘Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning’. In: *Advances in Neural Information Processing Systems*. Ed. by J. Shawe-Taylor, R. Zemel et al. Vol. 24. Curran Associates, Inc., 2011.
-

- [NR08] Peter Neal and Gareth Roberts. ‘Optimal scaling for random walk Metropolis on spherically constrained target densities’. In: *Methodology and Computing in Applied Probability* 10.2 (2008), pp. 277–297.
- [NRY+12] Peter Neal, Gareth Roberts et al. ‘Optimal scaling of random walk Metropolis algorithms with discontinuous target densities’. In: *Annals of Applied Probability* 22.5 (2012), pp. 1880–1927.
- [Nea93] R. M. Neal. ‘Bayesian Learning via Stochastic Dynamics’. In: *Advances in Neural Information Processing Systems 5, [NIPS Conference]*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1993, pp. 475–482.
- [NB01] A. Nedić and D. Bertsekas. ‘Convergence Rate of Incremental Subgradient Algorithms’. In: *Stochastic Optimization: Algorithms and Applications*. Boston, MA: Springer US, 2001, pp. 223–264.
- [NWS14] D. Needell, R. Ward et al. ‘Stochastic gradient descent, weighted sampling, and the randomized Kaczmarz algorithm’. In: *Advances in neural information processing systems*. 2014, pp. 1017–1025.
- [Nem+09] Arkadi Nemirovski, Anatoli Juditsky et al. ‘Robust stochastic approximation approach to stochastic programming’. In: *SIAM Journal on optimization* 19.4 (2009), pp. 1574–1609.
- [Nes04] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Applied Optimization. Springer, 2004.
- [NV08] Y. Nesterov and J-Ph Vial. ‘Confidence level solutions for stochastic programming’. In: *Automatica* 44.6 (2008), pp. 1559–1568.
- [Nes83] Yurii Evgen’evich Nesterov. ‘A method of solving a convex programming problem with convergence rate $O(k^{-2})$ ’. In: *Doklady Akademii Nauk*. Vol. 269. 3. Russian Academy of Sciences. 1983, pp. 543–547.
- [Oja92] Erkki Oja. ‘Principal components, minor components, and linear neural networks’. In: *Neural networks* 5.6 (1992), pp. 927–935.
- [OK85] Erkki Oja and Juha Karhunen. ‘On stochastic approximation of the eigenvectors and eigenvalues of the expectation of a random matrix’. In: *Journal of mathematical analysis and applications* 106.1 (1985), pp. 69–84.
- [PB14] Neal Parikh and Stephen Boyd. ‘Proximal algorithms’. In: *Foundations and Trends in Optimization* 1.3 (2014), pp. 127–239.
- [Par81] G. Parisi. ‘Correlation functions and computer simulations’. In: *Nuclear Physics B* 180 (1981), pp. 378–384.
- [PFA06] X. Pennec, P. Fillard et al. ‘A Riemannian framework for tensor computing’. In: *International Journal of Computer Vision* 66.1 (2006), pp. 41–66.
- [Pen99] Xavier Pennec. ‘Probabilities and statistics on Riemannian manifolds: Basic tools for geometric measurements.’ In: *NSIP*. Vol. 3. 1999, pp. 194–198.

-
- [Per16] Marcelo Pereyra. ‘Proximal Markov chain Monte Carlo algorithms’. In: *Statistics and Computing* 26.4 (2016), pp. 745–760.
- [Pfl86] G. Pflug. ‘Stochastic Minimization with Constant Step-Size: Asymptotic Laws’. In: *SIAM Journal on Control and Optimization* 24.4 (1986), pp. 655–666. eprint: <https://doi.org/10.1137/0324039>.
- [Pfl90] Georg Ch Pflug. ‘Non-asymptotic confidence bounds for stochastic approximation algorithms with constant step size’. In: *Monatshefte für Mathematik* 110 (1990), pp. 297–314.
- [Pil22] Natesh S. Pillai. ‘Optimal Scaling for the Proximal Langevin Algorithm in High Dimensions’. In: *arXiv preprint arXiv:2204.10793* (2022).
- [PST12] Natesh S. Pillai, Andrew M. Stuart et al. ‘Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions’. In: *The Annals of Applied Probability* 22.6 (2012), pp. 2320–2356.
- [PJ92] B. T. Polyak and A. B. Juditsky. ‘Acceleration of stochastic approximation by averaging’. In: *SIAM J. Control Optim.* 30.4 (1992), pp. 838–855.
- [RRT17] M. Raginsky, A. Rakhlin et al. ‘Non-convex learning via Stochastic Gradient Langevin Dynamics: a nonasymptotic analysis’. In: *arXiv preprint arXiv:1702.03849* (2017).
- [RM51] H. Robbins and S. Monro. ‘A stochastic approximation method’. In: *The Annals of mathematical Statistics* 22.3 (1951), pp. 400–407.
- [RC10] C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer New York, 2010.
- [RGG97a] G. O. Roberts, A. Gelman et al. ‘Weak convergence and optimal scaling of random walk Metropolis algorithms’. In: *Ann. Applied Prob.* 7 (1997), pp. 110–120.
- [RR98a] G. O. Roberts and J. S. Rosenthal. ‘Optimal scaling of discrete approximations to Langevin diffusions’. In: *J. R. Stat. Soc. Ser. B Stat. Methodol.* 60.1 (1998), pp. 255–268.
- [RT96a] G. O. Roberts and R. L. Tweedie. ‘Exponential convergence of Langevin distributions and their discrete approximations’. In: *Bernoulli* 2.4 (1996), pp. 341–363.
- [RGG97b] Gareth O Roberts, Andrew Gelman et al. ‘Weak convergence and optimal scaling of random walk Metropolis algorithms’. In: *Annals of Applied Probability* 7.1 (1997), pp. 110–120.
- [RR98b] Gareth O Roberts and Jeffrey S Rosenthal. ‘Optimal scaling of discrete approximations to Langevin diffusions’. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 60.1 (1998), pp. 255–268.
- [RR+01] Gareth O Roberts, Jeffrey S Rosenthal et al. ‘Optimal scaling for various Metropolis-Hastings algorithms’. In: *Statistical Science* 16.4 (2001), pp. 351–367.
-

- [RT96b] Gareth O Roberts and Richard L Tweedie. ‘Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms’. In: *Biometrika* 83.1 (1996), pp. 95–110.
- [RT96c] Gareth O. Roberts and Richard L. Tweedie. ‘Exponential convergence of Langevin distributions and their discrete approximations’. In: *Bernoulli* 2.4 (1996), pp. 341–363.
- [RW98] R Tyrrell Rockafellar and Roger J-B Wets. *Variational Analysis*. Vol. 317. Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer, Berlin, Heidelberg, 1998.
- [RDF78] P. J. Rossky, J. D. Doll et al. ‘Brownian dynamics as smart Monte Carlo simulation’. In: *The Journal of Chemical Physics* 69.10 (1978), pp. 4628–4633.
- [Rup88] David Ruppert. *Efficient estimations from a slowly convergent Robbins-Monro process*. Tech. rep. Cornell University Operations Research and Industrial Engineering, 1988.
- [Sac58] Jerome Sacks. ‘Asymptotic distribution of stochastic approximation procedures’. In: *The Annals of Mathematical Statistics* 29.2 (1958), pp. 373–405.
- [SM19] S. Said and J. H. Manton. ‘The Riemannian Barycentre as a Proxy for Global Optimisation’. In: *Geometric Science of Information, GSI*. 2019, pp. 657–664.
- [Sak76] T. Sakai. ‘On the cut loci of compact symmetric spaces’. In: *Hokkaido Mathematical Journal* 6.3 (1976), pp. 136–161.
- [SKM19] H. Sato, H. Kasai et al. ‘Riemannian stochastic variance reduced gradient algorithm with retraction and vector transport’. In: *SIAM Journal on Optimization* 29.2 (2019), pp. 1444–1472.
- [Sch95] M. J. Schervish. *Theory of statistics*. Springer series in statistics. Corrected second printing : 1997. New York, Berlin, Heidelberg: Springer, 1995.
- [Sha19] S. M. Shah. ‘Stochastic Approximation on Riemannian Manifolds’. In: *Applied Mathematics & Optimization* (May 2019).
- [Sha+12] Shai Shalev-Shwartz et al. ‘Online learning and online convex optimization’. In: *Foundations and Trends® in Machine Learning* 4.2 (2012), pp. 107–194.
- [SSS07] Shai Shalev-Shwartz, Yoram Singer et al. ‘Pegasos: Primal estimated sub-gradient solver for svm’. In: *Proceedings of the 24th international conference on Machine learning*. 2007, pp. 807–814.
- [SR+09] Chris Sherlock, Gareth Roberts et al. ‘Optimal scaling of the random walk Metropolis on elliptically symmetric unimodal targets’. In: *Bernoulli* 15.3 (2009), pp. 774–798.
- [Shi96] Albert N Shiryaev. *Probability*. Vol. 25. Springer, 1996.
- [SV79] Daniel W. Stroock and Srinivasa Varadhan. *Multidimensional Diffusion Processes*. Springer, 1979.

-
- [Stu03] K. T. Sturm. ‘Probability Measures on Metric Spaces of Nonpositive Curvature’. In: *Contemporary Mathematics* 338 (Jan. 2003).
- [SQW17] J. Sun, Q. Qu et al. ‘Complete Dictionary Recovery Over the Sphere II: Recovery by Riemannian Trust-Region Method’. English (US). In: *IEEE Transactions on Information Theory* 63.2 (Feb. 2017), pp. 885–914.
- [SFF19] Yue Sun, Nicolas Flammarion et al. ‘Escaping from saddle points on Riemannian manifolds’. In: *NeurIPS*. 2019.
- [Tie98] L. Tierney. ‘A note on Metropolis-Hastings kernels for general state spaces’. In: 8 (1998), pp. 1–9.
- [Tri+18] N. Tripuraneni, N. Flammarion et al. ‘Averaging Stochastic Gradient Descent on Riemannian Manifolds’. In: *Conference On Learning Theory, COLT*. 2018, pp. 650–687.
- [VBS19] S. Vaswani, F. Bach et al. ‘Fast and Faster Convergence of SGD for Over-Parameterized Models and an Accelerated Perceptron’. In: *The 22nd International Conference on Artificial Intelligence and Statistics*. 2019, pp. 1195–1204.
- [VDC19] Maxime Vono, Nicolas Dobigeon et al. ‘Bayesian image restoration under Poisson noise and log-concave prior’. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2019, pp. 1712–1716.
- [VSS21] Jacob Vorstrup Goldman, Torben Sell et al. ‘Gradient-Based Markov Chain Monte Carlo for Bayesian Inference With Non-Differentiable Priors’. In: *Journal of the American Statistical Association* 0.0 (2021), pp. 1–12.
- [Xia09] Lin Xiao. ‘Dual averaging method for regularized stochastic learning and online optimization’. In: *Advances in Neural Information Processing Systems 22* (2009).
- [YRR20] Jun Yang, Gareth O Roberts et al. ‘Optimal scaling of random-walk Metropolis algorithms on general target distributions’. In: *Stochastic Processes and their Applications* 130.10 (2020), pp. 6094–6132.
- [ZS16] H. Zhang and S. Sra. ‘First-order Methods for Geodesically Convex Optimization’. In: *Conference on Learning Theory, COLT*. 2016, pp. 1617–1638.
- [ZRS16] Hongyi Zhang, Sashank J Reddi et al. ‘Riemannian SVRG: Fast stochastic optimization on Riemannian manifolds’. In: *Advances in Neural Information Processing Systems*. 2016, pp. 4592–4600.
- [ZS18] Hongyi Zhang and Suvrit Sra. ‘Towards Riemannian accelerated gradient methods’. In: *arXiv preprint arXiv:1806.02812* (2018).
- [ZZS18] Jingzhao Zhang, Hongyi Zhang et al. ‘R-spider: A fast Riemannian stochastic optimization algorithm with curvature independent rate’. In: *arXiv preprint arXiv:1811.04194* (2018).
-

-
- [ZCZ22] Xinkai Zhou, Eric C. Chi et al. ‘Proximal MCMC for Bayesian Inference of Constrained and Regularized Estimation’. In: *arXiv preprint arXiv:2205.07378* (2022).

Titre: Analyse d'Algorithmes Stochastiques pour l'Echantillonnage et l'Approximation Riemannienne

Mots clés: Variété riemannienne, optimisation stochastique, approximation stochastique, échantillonnage par chaîne de Markov, échelonnement optimal

Résumé: Cette thèse porte sur l'analyse de deux algorithmes stochastiques : l'approximation stochastique sur une variété riemannienne et une méthode de Monte Carlo par chaîne de Markov. Pour les deux algorithmes, nous étendons la théorie existante pour une meilleure compréhension de l'efficacité de ces méthodes. Dans une première partie, nous nous intéressons à deux cadres pour l'approximation stochastique sur une variété riemannienne. D'une part, nous examinons le taux de convergence, en mettant en place des majorations non-asymptotiques avec des hypothèses plus faibles qu'auparavant, ainsi qu'avec une modélisation plus souple du bruit. D'une autre part, on s'intéresse de près au cas du pas constant, dans lequel on s'intéresse à la convergence de la loi stationnaire

de la chaîne de Markov définie par le schéma stochastique. On effectue une décomposition asymptotique avec un équilibre biais-variance, ainsi qu'un théorème central limite. Nous apportons des applications de nos résultats sur des exemples classiques du domaine. Dans une deuxième partie, on s'intéresse au comportement en grande dimension d'une méthode de Monte Carlo par chaîne de Markov qui approche la distribution cible de façon lisse, afin d'utiliser le gradient pour accélérer la convergence. Dans le contexte classique de l'échelonnement optimal, nous trouvons des comportements similaires à la théorie existante pour des distributions lisses mais nouveaux pour des distributions ayant une dérivée discontinue en un point. Nous illustrons ces résultats par des simulations.

Title: Analysis of Stochastic Algorithms for Sampling and Riemannian Approximation

Keywords: Riemannian manifold, stochastic optimisation, stochastic approximation, Markov chain sampling, optimal scaling

Abstract: This thesis deals with the analysis of two stochastic algorithms: stochastic approximation on Riemannian manifolds and a Markov chain Monte Carlo method. For both algorithms, we extend the existing theory for a better understanding of the efficiency of these methods. In a first part, we focus on two frameworks for stochastic approximation on Riemannian manifolds. On the one hand, we examine the rate of convergence, by proving non-asymptotic bounds with weaker assumptions than before, as well as with a more flexible set of assumptions for the noise. On the other hand, we take a close look at the constant step-size case, in which we are interested in the convergence of the stationary law of

the Markov chain defined by the stochastic scheme. We perform an asymptotic decomposition with a bias-variance equilibrium, and a central limit theorem. We provide applications of our results on classical examples of Riemannian stochastic approximation. In a second part, we focus on the high dimensional behavior of a Markov chain Monte Carlo method which approximates the target distribution smoothly, in order to use its gradient to accelerate convergence. In the classical context of optimal scaling, we find behaviors similar to the existing theory for smooth distributions but new ones for distributions with a discontinuity in their derivative. We illustrate these results with simulations.