



HAL
open science

Convergence Analysis of Riemannian Stochastic Approximation Schemes

Alain Durmus, Pablo Jiménez, Eric Moulines, Salem Said, Hoi-To Wai

► **To cite this version:**

Alain Durmus, Pablo Jiménez, Eric Moulines, Salem Said, Hoi-To Wai. Convergence Analysis of Riemannian Stochastic Approximation Schemes. 2020. hal-02626383v1

HAL Id: hal-02626383

<https://hal.science/hal-02626383v1>

Preprint submitted on 26 May 2020 (v1), last revised 27 May 2020 (v2)

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.

Convergence Analysis of Riemannian Stochastic Approximation Schemes

Alain Durmus¹, Pablo Jiménez², Éric Moulines², Salem Said³, and Hoi-To Wai⁴

¹*Centre Borelli, UMR 9010, École Normale Supérieure Paris-Saclay, CNRS*

²*CMAP, UMR 7641, École Polytechnique, CNRS*

³*Laboratoire IMS, UMR 5218, CNRS, Université de Bordeaux*

⁴*Department of SEEM, The Chinese University of Hong Kong*

May 26, 2020

Abstract

This paper analyzes the convergence for a large class of Riemannian stochastic approximation (SA) schemes, which aim at tackling stochastic optimization problems. In particular, the recursions we study use either the exponential map of the considered manifold (geodesic schemes) or more general retraction functions (retraction schemes) used as a proxy for the exponential map. Such approximations are of great interest since they are low complexity alternatives to geodesic schemes. Under the assumption that the mean field of the SA is correlated with the gradient of a smooth Lyapunov function (possibly non-convex), we show that the above Riemannian SA schemes find an $\mathcal{O}(b_\infty + \log n/\sqrt{n})$ -stationary point (in expectation) within $\mathcal{O}(n)$ iterations, where $b_\infty \geq 0$ is the asymptotic bias. Compared to previous works, the conditions we derive are considerably milder. First, all our analysis are global as we do not assume iterates to be a-priori bounded. Second, we study biased SA schemes. To be more specific, we consider the case where the mean-field function can only be estimated up to a small bias, and/or the case in which the samples are drawn from a controlled Markov chain. Third, the conditions on retractions required to ensure convergence of the related SA schemes are weak and hold for well-known examples. We illustrate our results on three machine learning problems.

Email: alain.durmus@cmla.ens-cachan.fr, pablo.jimenez-moreno@polytechnique.edu
eric.moulines@polytechnique.edu, salem.said@u-bordeaux.fr, htwai@se.cuhk.edu.hk

1 Introduction

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

$$\text{find } \theta \in \Theta \text{ satisfying } h(\theta) = 0_\theta, \quad \text{where } h(\theta) = \int_{\mathcal{X}} H_\theta(x) d\pi_\theta(x), \quad (1)$$

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. In particular, note that this framework includes stochastic optimization problems where the objective function $f : \Theta \rightarrow \mathbb{R}^*$ is smooth (but possibly non-convex even in the geodesic sense), taking $h = \text{grad } f$, where grad is the Riemannian gradient operator, see Appendix A.8. Problem (1) is motivated by applications in principal component analysis, combinatorial optimization, and geometric barycenter.

Our aim is to study the stochastic approximation (SA) schemes adapted to a Riemannian setting to compute an *approximate stationary point* of (1). In the case where the exponential map on Θ can be evaluated, so geodesic curves can be explicitly calculated, the Riemannian SA scheme that we consider is a natural adaptation of the standard Robbins-Monro algorithm [19], namely the iterates are updated via the exponential map applied to Monte Carlo estimates of $h(\theta)$ for $\theta \in \Theta$. Furthermore, the present paper considers the generalizations of the above methodology where (a) the considered estimators are potentially *biased*, *i.e.* they are either based on a Markov chain targeting π_θ and more generally the right-hand side of (1) is not satisfied *i.e.* $\int_{\mathcal{X}} H_\theta(x) d\pi_\theta(x) = h(\theta) + b(\theta)$, where $b : \Theta \rightarrow T\Theta$ is a smooth vector field modeling the bias introduced by the scheme; (b) the updates are computed using *retraction* functions which are computationally cheaper to evaluate than the exponential map but provide good numerical approximation of the latter. These generalizations are important as we adapt the Riemannian SA scheme to instances of (1) with various constraints in data acquisitions and computation.

A lot of effort has been paid in the analysis of SA in the last decade due to its applications in machine learning and signal processing. As mentioned previously, SA encompasses stochastic gradient (SG) methods but is not limited to the case of $h = \text{grad } f$. Most common examples for which the latter condition is not satisfied are second order methods, online Expectation Maximization algorithms [7], Q-learning [14] or policy gradient [3]. In this sense, our work is in line with recent studies of SA in the Euclidean setting under mild conditions on the mean-field h and which can be applied to stochastic optimization seeking at minimizing non-convex objective function, see e.g. [12, 15] and the references therein. In addition, we also extend results regarding biased SA [21, 15] to the Riemannian setting.

SA on Riemannian manifolds has attracted attentions recently. The pioneering work [5] proved the asymptotic convergence of Riemannian SA to a critical point using martingale techniques usually used in the analysis of SA in Euclidean spaces, see [4, 10]. Under the assumption that the objective function f is geodesically convex (g-convex), [25] showed that the Riemannian SGD scheme finds an ϵ -optimal solution in $\mathcal{O}(1/\epsilon^2)$ iterations. This result

has been improved since then in the authors' follow up works [24, 26]. However, note that these papers focus on geodesic schemes which can be computationally expensive. Another related work is [23] which considers a *retraction* based averaging scheme, and improves the convergence rate to $\mathcal{O}(1/\epsilon)$ by assuming f is geodesically strongly convex and the SA iterates stay in a compact set. In addition, this paper also considers non-convex settings and shows a central limit theorem under strong assumptions for a sequence $(\psi(\bar{\theta}_n))_{n \in \mathbb{N}}$ where $(\bar{\theta}_n)_{n \in \mathbb{N}}$ is an appropriate averaged sequence based on $(\theta_n)_{n \in \mathbb{N}}$ and for some function $\psi : \Theta \rightarrow \mathbb{T}_{\theta_*} \Theta$, θ_* is the minimum of f . To deal with relatively general conditions on the mean-field h , [20] analyzes stochastic recursion schemes adapting the well-known ODE method [4, 16] to the Riemannian setting to obtain asymptotic convergence. Note however that this study is limited to the case where Θ is compact. Analysis of Riemannian SGD for non-convex objective has been addressed in [13] which provides a similar result to one of ours in the case of unbiased Riemannian SA. However, they require a strong assumption on the retraction map, the objective function and the manifold. We should also mention that a few other papers have studied deterministic optimization on Riemannian manifold and the convergence to local minimum in non-convex settings, see e.g. [6, 22, 8].

Our contributions are three-fold. (A) In Section 3, 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 strong g -convexity type of assumptions. Under these relaxed settings, the considered geodesic SA scheme finds an $\mathcal{O}(b_\infty + \log n / \sqrt{n})$ -stationary point in $\mathcal{O}(n)$ iterations in expectation, if the gradient samples are asymptotically $\mathcal{O}(b_\infty)$ -biased. We cover settings when the noise is a sequence of martingale increments or stems from a Markovian dynamics, see Theorem 2 & 3. (B) In Section 4, we consider retraction schemes with first or second order retraction. Similar to the geodesic schemes, the methods we introduce find an $\mathcal{O}(b_\infty + \log n / \sqrt{n})$ -stationary point in $\mathcal{O}(n)$ iterations in expectation, see Theorem 5 & 6. The conditions we consider are stronger than in the case of geodesic schemes but relax the ones required in previous works. Further, we illustrate on several examples that the required conditions on the retraction function hold. (C) In Section 5, we consider three example applications and show that the required convergence conditions are satisfied. These applications are novel as they cannot be handled using previous analysis.

Notations 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 $u_n = \Omega(v_n)$. We denote the tangent space of Θ at θ by $\mathbb{T}_\theta \Theta$ and its tangent bundle $\mathbb{T}\Theta$. 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 $Df : \mathbb{T}\Theta_0 \rightarrow \mathbb{T}\Theta_1$. The Riemannian metric on Θ is denoted by g but for ease of notation and if there is no risk of confusion, for any $\theta \in \Theta$, $u, v \in \mathbb{T}_\theta \Theta$, we should denote $g_\theta(u, v) = \langle u, v \rangle_\theta$ and $\|u\|_\theta^2 = g_\theta(u, u)$. $\mathbb{T}_{t_0 t_1}^\gamma : \mathbb{T}_{\gamma(t_0)} \Theta \rightarrow \mathbb{T}_{\gamma(t_1)} \Theta$ 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 Appendix A.

2 Riemannian Stochastic Approximation Schemes

Let $(X_n)_{n \in \mathbb{N}^*}$ be a stochastic process defined on the filtered probability space $(X, \mathcal{X}, \mathbb{P}, (\mathcal{F}_n)_{n \in \mathbb{N}})$, $(\mathcal{F}_n)_{n \in \mathbb{N}}$ -adapted. The function $\text{Ret} : T\Theta \rightarrow \Theta$ is a retraction and for any $\theta \in \Theta$, Ret_θ stands for the restriction of Ret to $T_\theta\Theta$. We consider the cases where the function Ret can be either the exponential map or a computationally efficient proxy. The present paper studies stochastic approximation (SA) sequences $(\theta_n)_{n \in \mathbb{N}}$, used as approximate solutions of (1), starting from $\theta_0 \in \Theta$ and defined by the recursion

$$\theta_{n+1} = \text{Ret}_{\theta_n} \{ \eta_{n+1} (H_{\theta_n}(X_{n+1}) + b_{\theta_n}(X_{n+1})) \} , \quad (2)$$

where $H, b : \Theta \times X \rightarrow T\Theta$ are bi-measurable functions, and $(\eta_n)_{n \in \mathbb{N}^*}$ is a sequence of positive stepsizes. For any $n \in \mathbb{N}^*$, $H_{\theta_n}(X_{n+1})$ is a noisy version of $h(\theta_n)$ and $b_{\theta_n}(X_{n+1})$ is an additional bias term. Eq. (2) describes the recursion where θ_n is moved along the direction given by $\eta_{n+1}(H_{\theta_n}(X_{n+1}) + b_{\theta_n}(X_{n+1}))$ on the tangent space $T_{\theta_n}\Theta$ originated at θ_n , and the retraction Ret_{θ_n} “projects” the updated iterate back into Θ . In this sense, the vector $H_{\theta_n}(X_{n+1}) + b_{\theta_n}(X_{n+1}) \in T_{\theta_n}\Theta$ can be interpreted as the stochastic update vector in Euclidean SA.

Let us discuss some basic assumptions to be used throughout this paper. Consider the following condition on the Riemannian manifold.

A 1. Θ is a geodesically complete Riemannian manifold of dimension $d \in \mathbb{N}^*$, i.e. the exponential function Exp is defined over $T\Theta$.

For a definition of the Riemannian exponential mapping, see Appendix A.5. We assume in addition the existence of a Lyapunov function for the mean vector field h .

A 2. There exists a continuously differentiable function $V : \Theta \rightarrow \mathbb{R}^*$ satisfying the conditions below.

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

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

(b) The 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 $\gamma : [0, 1] \rightarrow \Theta$ between θ_0 and θ_1 ,

$$\|\text{grad } V(\theta_1) - T_{\theta_0}^\gamma \text{grad } V(\theta_0)\|_{\theta_1} \leq L \ell(\gamma) , \quad (3)$$

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

The first inequality in (a) strengthens the Lyapunov’s second method for stability and Lasalle’s invariance principle for integral curves or ordinary differential equations; see [17]. In fact, it implies $\|h(\theta)\|_\theta \leq \tilde{c} \|\text{grad } V(\theta)\|_\theta$ for any $\theta \in \Theta$, for $\tilde{c} = \underline{c}^{-1} > 0$. 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 10 in Appendix A.9. Overall, conditions (a) and (b) ensure the stability of the recursion (2) as they imply that $h(\theta)$ 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 , one has $\|h(\theta)\|_\theta \leq C(\ell(\gamma) + 1)$. Importantly, A1, A2 allow us to generalize the well-known descent lemma to the Riemannian setting, as follows:

Lemma 1. *Assume A1, A2-(b) hold. For any $\theta_0, \theta_1 \in \Theta$ and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0, θ_1 ,*

$$\left| V(\theta_1) - V(\theta_0) - \langle \text{grad } V(\theta_0), \dot{\gamma}_0 \rangle_{\theta_0} \right| \leq L\ell(\gamma)^2/2.$$

The proof is postponed to Appendix B. This result was stated in [25] without proof as a consequence of A2-(a). Moreover, compared to [1, Lemma 7.4.7], our result holds for any geodesic curves, and is not limited to the length-minimizing ones.

Finally, we assume that the bias is uniformly bounded.

A3. *There exists a constant b_∞ such that $\sup_{(\theta, x) \in \Theta \times \mathcal{X}} \|b_\theta(x)\|_\theta \leq b_\infty$.*

The present paper provides *a priori* non-asymptotic guarantees for the SA scheme (2) to find an (approximate) stationary point where $\|h(\theta)\|_\theta^2 \approx \tilde{c}b_\infty$ for some $\tilde{c} > 0$. Roughly, a non-asymptotic performance guarantee ensures the ability of the scheme (2) to produce an iterate $\theta_{I_n} \in \{\theta_0, \dots, \theta_n\}$ such that $\mathbb{E}[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2] \leq \epsilon + \tilde{c}b_\infty$, for given $n \in \mathbb{N}^*$ and $\epsilon > 0$, where for any $k \in \mathbb{N}^*$, we define the randomized stopping rule I_k independent of $(\theta_0, X_1, X_2, \dots)$ with distribution for any $\ell \in \{0, \dots, k\}$,

$$\mathbb{P}(I_k = \ell) = \left(\sum_{i=0}^k \eta_{i+1} \right)^{-1} \eta_{\ell+1}. \quad (4)$$

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

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

For any $x \in \mathcal{X}$ and $\theta \in \Theta$, we define the error in estimating the mean vector field as:

$$e_\theta(x) = H_\theta(x) - h(\theta). \quad (5)$$

3 Analysis of Geodesic SA Schemes

In this section, the retraction Ret appearing in (2) is the Riemannian exponential map of Θ , $\text{Ret} = \text{Exp}$. Such scheme will also be referred to as the *geodesic SA scheme*.

Consider first the martingale setting formalized by the following conditions.

MD1. *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. for any $n \in \mathbb{N}^*$, $\mathbb{E}[e_{\theta_n}(X_{n+1}) | \mathcal{F}_n] = 0$. Also, there exist $\sigma_0^2, \sigma_1^2 < \infty$ such that for any $n \in \mathbb{N}^*$,*

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

We observe that:

Theorem 2. *Assume **A 1-A 2-A 3-MD 1** hold. Consider $(\theta_n)_{n \in \mathbb{N}}$ defined by (2) with $\text{Ret} = \text{Exp}$. If $\sup_{k \in \mathbb{N}^*} \eta_k \leq \underline{c}/(6L(1 + \sigma_1^2))$, then for any $n \in \mathbb{N}^*$,*

$$\mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq 4(\underline{c}\Gamma_{n+1})^{-1} \left\{ \mathbb{E}[V(\theta_0)] + (3L/2)(\sigma_0^2 + b_\infty)\Gamma_{n+1}^{(2)} \right\} + b_\infty \bar{c}/\underline{c}^2, \quad (6)$$

where $I_n \in \{0, \dots, n\}$ is a discrete random variable with distribution defined by (4).

The proof is postponed to Appendix C.1.

A popular choice for the step size is $\eta_k = \rho/\sqrt{k + k_0}$ for some $\rho, k_0 > 0$. In this case, we have $\Gamma_{n+1} = \Omega(\sqrt{n})$ and $\Gamma_{n+1}^{(2)} = \mathcal{O}(\log n)$. Consequently, Theorem 2 shows that (2) finds an $\mathcal{O}(b_\infty + \log(n)/\sqrt{n})$ -stationary point to (1) in $\mathcal{O}(n)$ iterations.

We now turn to the Markovian setting which is formalized by the following assumptions on the sequence $(X_n)_{n \in \mathbb{N}^*}$ and the stochastic vector field H .

MA 1. *There exists a Markov kernel P on $(\Theta \times \mathbf{X}) \times \mathcal{X}$ such that for any $n \in \mathbb{N}$ and bounded and measurable function $\varphi : \mathbf{X} \rightarrow \mathbb{R}_+$, $\mathbb{E}[\varphi(X_{n+1})|\mathcal{F}_n] = \int_{\mathbf{X}} \varphi(y)P_{\theta_n}(X_n, dy)$. In addition, P satisfies the following conditions.*

- (i) *For any $\theta \in \Theta$, P_θ admits a unique invariant distribution π_θ satisfying $h(\theta) = \int_{\mathbf{X}} H_\theta(y) d\pi_\theta(y)$.*
- (ii) *Consider the error function $e_\theta(x)$ defined in (5). There exists a measurable function $\hat{e} : \Theta \times \mathbf{X} \rightarrow \mathbb{T}\Theta$ satisfying for any $x \in \mathbf{X}$, $\theta \in \Theta$, $\hat{e}_\theta(x) \in \mathbb{T}\Theta$ and*

$$\hat{e}_\theta(x) - \int_{\mathbf{X}} P_\theta(x, dy)\hat{e}_\theta(y) = e_\theta(x). \quad (7)$$

- (iii) *There exist $e_\infty, \hat{e}_\infty \geq 0$ such that $\sup_{\theta \in \Theta, x \in \mathbf{X}} \|e_\theta(x)\|_\theta \leq e_\infty$, $\sup_{\theta \in \Theta, x \in \mathbf{X}} \|\hat{e}_\theta(x)\|_\theta \leq \hat{e}_\infty$.*

- (iv) *There exists $L_{\hat{e}} \geq 0$ such that for any $\theta_0, \theta_1 \in \Theta$, and geodesic curve $\gamma : [0, 1] \rightarrow \Theta$ between θ_0, θ_1 ,*

$$\sup_{x \in \mathbf{X}} \left\| \int_{\mathbf{X}} P_{\theta_1}(x, dy)\hat{e}_{\theta_1}(x) - \mathbb{T}_{\theta_1}^\gamma \left[\int_{\mathbf{X}} P_{\theta_0}(x, dy)\hat{e}_{\theta_0}(y) \right] \right\|_{\theta_1} \leq L_{\hat{e}} \ell(\gamma).$$

Condition (ii) assumes the existence of solutions to the Poisson equation (7), which can in general be established if the Markov kernel P_θ is V -geometrically or uniformly ergodic for any $\theta \in \Theta$, see [9, Section 21.2]. In addition, condition (iii) requires that function e is uniformly bounded and the corresponding solution to the Poisson equation (for fixed θ) as well. The uniform boundedness condition on \hat{e} holds if, for example, the function e is uniformly bounded and for any $\theta \in \Theta$, P_θ is uniformly ergodic with convergence rate independent of θ . The latter could be relaxed using appropriate Lyapunov conditions following [11]. Finally, the last assumption (iv) is implied by smoothness conditions on the Markov kernel with respect to the SA parameter θ .

Under **A1-A2-A3-MA1** and for a sequence of stepsize $(\eta_k)_{k \in \mathbb{N}^*}$, consider the constants

$$\begin{aligned} D_{\hat{e}} &= 1 + e_\infty + b_\infty + \hat{e}_\infty(L + \bar{c}a_2), & C(\eta_1) &= \bar{c}\hat{e}_\infty(\eta_1 + 2), \\ C_{\hat{e}} &= L\bar{c}(e_\infty + b_\infty) + L\hat{e}_\infty(e_\infty + b_\infty + 1) + (3L/2)(b_\infty^2 + e_\infty^2). \end{aligned} \quad (8)$$

We observe that:

Theorem 3. *Assume **A1-A2-A3-MA1** hold. Let $(\eta_k)_{k \in \mathbb{N}^*}$ be a sequence of stepsizes and $a_1, a_2 \geq 0$ satisfying*

$$\begin{aligned} \sup_{k \in \mathbb{N}^*} \{\eta_{k+1}/\eta_k\} &\leq 1, & \sup_{k \in \mathbb{N}^*} \{\eta_k/\eta_{k+1}\} &\leq a_1, \\ \sup_{k \in \mathbb{N}^*} \{|\eta_k - \eta_{k+1}|/\eta_k^2\} &\leq a_2, & \sup_{k \in \mathbb{N}^*} \eta_k &\leq c_1/(4(3L/2 + D_{\hat{e}})). \end{aligned} \quad (9)$$

Consider $(\theta_n)_{n \in \mathbb{N}}$ defined by (2) with $\text{Ret} = \text{Exp}$. Then for any $n \in \mathbb{N}^*$,

$$\mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq 4(\underline{c}\Gamma_{n+1})^{-1} \left\{ \mathbb{E}[V(\theta_0)] + C(\eta_1) + C_{\hat{e}}\Gamma_{n+1}^{(2)} \right\} + b_\infty \bar{c}/\underline{c}^2. \quad (10)$$

where I_n has distribution defined by (4).

The proof is postponed to Appendix C.2.

We observe that the right hand side of (10) is akin to that of (6). Similar to the martingale noise setting, choosing the step size $\eta_k = \rho/\sqrt{k + k_0}$ and using (10) shows that (2) finds an $\mathcal{O}(b_\infty + \log(n)/\sqrt{n})$ -stationary point to (1) in $\mathcal{O}(n)$ iterations.

4 Analysis of General Retraction SA Schemes

The previous section focused on the geodesic schemes that require performing the Riemannian exponential map Exp at each iteration. Evaluating the Exp map is often computationally hard as it requires computing complex functions such as matrix exponential. A popular idea is to use an efficient retraction function Ret that approximates Exp . One basic condition required on Ret is the following *local-rigidity* condition:

R1. *For any $\theta \in \Theta$, $\text{Ret}_\theta(0_\theta) = \theta$, where 0_θ is the zero element of $T_\theta\Theta$ as vector space, and $D\text{Ret}_\theta(0_\theta) = \text{Id}$.*

The following ensures that the inverse exponential map is defined on $\text{Ret}_\theta(T_\theta\Theta)$.

R2. *For any $(\theta, u) \in T\Theta$, $\text{Ret}_\theta(u) \notin \text{Cut}(\theta)$, where $\text{Cut}(\theta) \subset \Theta$ is the cut locus of θ (see the definition in A.5).*

Under **R2**, for any $\theta \in \Theta$, the following function is well defined:

$$\Phi_\theta = \text{Exp}_\theta^{-1} \circ \text{Ret}_\theta : T_\theta\Theta \rightarrow T_\theta\Theta, \quad (11)$$

which defines a bundle map $\Phi : T\Theta \rightarrow T\Theta$. Using Φ_θ , the SA scheme (2) may be written as

$$\theta_{n+1} = \text{Exp}_{\theta_n} \{ \eta_{n+1} (H_{\theta_n}(X_{n+1}) + b_{\theta_n}(X_{n+1}) + \Delta_{\theta_n, \eta_{n+1}}(X_{n+1})) \} , \quad (12)$$

where $\Delta_{\theta_n, \eta_{n+1}}(X_{n+1})$ is the ‘retraction bias’ defined for any $(\theta, x) \in \Theta \times \mathsf{X}$ and $\eta > 0$ by

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

In the special case when $\text{Ret} \equiv \text{Exp}$, we have $\Phi_{\theta_n} \equiv \text{Id}$ and thus $\Delta_{\theta, \eta}(x) = 0$ for any $(\theta, x) \in \Theta \times \mathsf{X}$ and $\eta > 0$. To control the retraction bias, we study the following conditions on Ret :

R3. Θ is a homogeneous Riemannian manifold with isometry group G (see the definition in A.7). For any $g \in \mathsf{G}$, $(\theta, u) \in T\Theta$,

$$g \cdot \text{Ret}_\theta(u) = \text{Ret}_{g \cdot \theta}(g \cdot u) . \quad (13)$$

In **R3**, for an isometry $g : \Theta \rightarrow \Theta$, $g \cdot \theta = g(\theta)$ and $g \cdot u = \text{D}g_\theta(u)$ is a vector in $T_{g \cdot \theta} \Theta$. Example of retractions defined on special matrix manifolds satisfying **R3** is given below. Then, we consider the following set of assumptions which lead to the definitions of regular *first-order* and *second-order* retractions [2]. These conditions will ensure that the first terms in the Taylor expansion of Φ_θ for $\theta \in \Theta$ do vanish.

R4. For any $(\theta, u) \in T\Theta$, there exists $\mathcal{L}^{(1)}(\theta) \geq 0$, such that $\sup_{t \in [0,1]} \| \text{D}^2 \Phi_\theta(tu)[u, u] \|_\theta \leq \mathcal{L}^{(1)}(\theta) \|u\|_\theta^2$ where the function $\Phi_\theta : T_\theta \Theta \rightarrow T_\theta \Theta$ is defined by (11).

Let ∇ be the Levi-Civita connection of the metric g on Θ (see Appendix A.2). We consider:

R5. For any $(\theta, u) \in T\Theta$, the following hold.

(i) Setting $\gamma(t) = \text{Ret}_\theta(tu)$ for $t \in \mathbb{R}$, the initial acceleration of the curve γ satisfies $\text{D}_t \dot{\gamma}(0) = 0_\theta$, where D_t stands for the covariant derivative along γ (see [18, Theorem 4.24]).

(ii) There exists $\mathcal{L}^{(2)}(\theta) \geq 0$, such that $\sup_{t \in [0,1]} \| \text{D}^3 \Phi_\theta(tu)[u, u, u] \|_\theta \leq \mathcal{L}^{(2)}(\theta) \|u\|_\theta^3$ where the function $\Phi_\theta : T_\theta \Theta \rightarrow T_\theta \Theta$ is defined by (11).

A retraction Ret satisfying **R4** (respectively **R5**) is called a regular first-order (respectively second-order) retraction. Note that **R5** does not imply **R4** nor vice versa. Intuitively, a first-order retraction approximates the Riemannian exponential map only to the first order, while a second-order retraction approximates the latter to the second order. These facts are formally established in the following result.

Lemma 4. Assume **R1**, **R2**, **R3**.

(a) Under **R4**, there exists $\mathcal{L}_\infty^{(1)} \geq 0$ such that $\|\Phi_\theta(u) - u\|_\theta \leq \mathcal{L}_\infty^{(1)} \|u\|_\theta^2$ for any $(\theta, u) \in \text{T}\Theta$.

(b) Under **R5**, there exists $\mathcal{L}_\infty^{(2)} \geq 0$ such that $\|\Phi_\theta(u) - u\|_\theta \leq \mathcal{L}_\infty^{(2)} \|u\|_\theta^3$ for any $(\theta, u) \in \text{T}\Theta$.

The proof is postponed to Appendix **D.1**.

Lemma 4 bounds the retraction bias. With this lemma in mind, our strategy for analyzing (2) with a general retraction is to incorporate the retraction bias in the analysis done for the geodesic scheme. Before doing so, we strengthen **A2** and **MD1**. Let $a > 0$.

A4. There exist $V_\infty, h_\infty \geq 0$ such that $\sup_{\theta \in \Theta} \|\text{grad} V(\theta)\|_\theta \leq V_\infty$ and $\sup_{\theta \in \Theta} \|h(\theta)\|_\theta \leq h_\infty$.

MD2 (a). Consider the noise function given by (5). There exists $\sigma_{(a)} \geq 0$ such that almost surely, it holds that

$$\mathbb{E} [\|e_{\theta_n}(X_{n+1})\|_{\theta_n}^a | \mathcal{F}_n] \leq \sigma_{(a)},$$

If **MD2(a)** holds for $a \in \mathbb{N}^*$, then **MD2(\tilde{a})** holds for any $\tilde{a} \in \{0, 1, \dots, a\}$ and $\sigma_{(\tilde{a})}$ will then stand for a constant such that almost surely, $\mathbb{E}[\|e_{\theta_n}(X_{n+1})\|_{\theta_n}^{\tilde{a}} | \mathcal{F}_n] \leq \sigma_{(\tilde{a})}$.

Depending on the order of retraction used and moment bound on the sequence $(e_{\theta_n}(X_{n+1}))_{n \in \mathbb{N}}$, we get the following convergence result for $(\theta_n)_{n \in \mathbb{N}}$.

Theorem 5. Assume **A1-A2-A3-A4**, **R1-R2-R3**, **MD1** hold. Consider $(\theta_n)_{n \in \mathbb{N}}$ defined by (2). For any $n \in \mathbb{N}^*$ and let $I_n \in \{0, \dots, n\}$ be distributed according to (4).

(a) Assume **R4**, **MD2(4)**. If $\sup_{k \in \mathbb{N}^*} \eta_k \leq (4L + 6\mathcal{L}_\infty^{(1)}V_\infty)^{-1}\underline{c}$, then

$$\mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq 2(\underline{c}\Gamma_{n+1})^{-1} \left\{ \mathbb{E}[V(\theta_0)] + c_0\Gamma_{n+1}^{(2)} + c_1\Gamma_{n+1}^{(4)} \right\} + 2b_\infty V_\infty / \underline{c},$$

where $c_0 = (3\mathcal{L}_\infty^{(1)}V_\infty + 2L)\{\sigma_{(2)} + b_\infty^2\}$, $c_1 = 54L(\mathcal{L}_\infty^{(1)})^2\{h_\infty^4 + b_\infty^4 + \sigma_{(4)}\}$.

(b) Assume **R5**, **MD2(6)**. If $\sup_{k \in \mathbb{N}^*} \eta_k \leq (4L)^{-1}\underline{c}$, then

$$\mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq 2(\underline{c}\Gamma_{n+1})^{-1} \left\{ \mathbb{E}[V(\theta_0)] + c_0\Gamma_{n+1}^{(2)} + c_1\Gamma_{n+1}^{(3)} + c_2\Gamma_{n+1}^{(6)} \right\} + 2b_\infty V_\infty / \underline{c},$$

where $c_0 = 2L\{b_\infty^2 + \sigma_{(2)}\}$, $c_1 = 9\mathcal{L}_\infty^{(2)}V_\infty\{h_\infty^3 + b_\infty^3 + \sigma_{(3)}\}$, $c_2 = 486L(\mathcal{L}_\infty^{(2)})^2\{h_\infty^6 + b_\infty^6 + \sigma_{(6)}\}$.

The proof is postponed to Appendix **D.2**.

For first-order retractions, one can relax the assumption of a bounded *fourth-order moment* for the noise term of sub-case (a), for a bounded first-order derivative for Φ_θ and a bounded *second-order moment* for the noise term instead, see Theorem 14 in the Appendix. For second-order retractions (b), the effects due to retraction are absorbed into the high-order terms $\Gamma_{n+1}^{(3)}/\Gamma_{n+1}$, $\Gamma_{n+1}^{(6)}/\Gamma_{n+1}$ which decay to zero faster than the ‘standard’ term

$\Gamma_{n+1}^{(2)}/\Gamma_{n+1}$. In other words, the effects of applying a second-order retraction will diminish asymptotically.

Nevertheless, the above results show that similar to the geodesic schemes, if we set the step sizes $\eta_k = \rho/\sqrt{k+k_0}$, then the SA scheme (2) using a first and/or second-order retraction finds an $\mathcal{O}(b_\infty + \log(n)/\sqrt{n})$ -stationary point to (1) in $\mathcal{O}(n)$ iterations.

We consider now the Markovian setting and we adopt the same set of assumptions on the Markov chain as in MA1. Define the constants:

$$\begin{aligned} C_{\hat{e}}^{\text{Ret}} &= \{L_{\hat{e}}\bar{c} + L\hat{e}_\infty\} (e_\infty + b_\infty + 1) + 2L(b_\infty^2 + e_\infty^2), \quad C^{\text{Ret}}(\eta_1) = \bar{c}\hat{e}_\infty(2 + \eta_1), \\ D^{\text{Ret}} &= \bar{c}\hat{e}_\infty(a_2 + 1) + L_{\hat{e}}\bar{c}(a_1(b_\infty + e_\infty) + 1) + L\hat{e}_\infty, \quad \tilde{e}^{(i)} = h_\infty^i + b_\infty^i + e_\infty^i, \quad i \in \mathbb{N}^*. \end{aligned}$$

Theorem 6. *Assume A1-A2-A3-A4, R1-R2-R3, MA1 hold. Assume that $(\eta_k)_{k \in \mathbb{N}^*}$ is a sequence of stepsizes and $a_1, a_2 \geq 0$ satisfying*

$$\sup_{k \in \mathbb{N}^*} \{\eta_{k+1}/\eta_k\} \leq 1, \quad \sup_{k \in \mathbb{N}^*} \{\eta_k/\eta_{k+1}\} \leq a_1, \quad \sup_{k \in \mathbb{N}^*} \{|\eta_k - \eta_{k+1}|/\eta_k^2\} \leq a_2, \quad (14)$$

Consider $(\theta_n)_{n \in \mathbb{N}}$ defined by (2). For any $n \in \mathbb{N}^*$, let $I_n \in \{0, \dots, n\}$ be distributed with (4).

(a) Assume R4. If in addition $\sup_{k \in \mathbb{N}^*} \eta_k \leq (4L + 2D^{\text{Ret}} + 6V_\infty \mathcal{L}_\infty^{(1)})^{-1} \underline{c}$, then

$$\begin{aligned} \mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] &\leq 2(\underline{c}\Gamma_{n+1})^{-1} \left\{ \mathbb{E}[V(\theta_0)] + C^{\text{Ret}}(\eta_1) + c_0\Gamma_{n+1}^{(2)} + c_1\Gamma_{n+1}^{(3)} + c_2\Gamma_{n+1}^{(4)} \right\} \\ &\quad + 2V_\infty b_\infty / \underline{c}, \quad (15) \end{aligned}$$

where $c_0 = C_{\hat{e}}^{\text{Ret}} + (3\mathcal{L}_\infty^{(1)}V_\infty)(b_\infty^2 + e_\infty^2)$, $c_1 = 3\hat{e}_\infty \mathcal{L}_\infty^{(1)}(L_{\hat{e}}h_\infty + L)\tilde{e}^{(2)}$, $c_2 = 54L(\mathcal{L}_\infty^{(1)})^2\tilde{e}^{(4)}$.

(b) Assume R5. If in addition $\sup_{k \in \mathbb{N}^*} \eta_k \leq (4L + 2D^{\text{Ret}})^{-1} \underline{c}$, then

$$\begin{aligned} \mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] &\leq 2(\underline{c}\Gamma_{n+1})^{-1} \left\{ \mathbb{E}[V(\theta_0)] + C^{\text{Ret}}(\eta_1) + c_0\Gamma_{n+1}^{(2)} + c_1\Gamma_{n+1}^{(3)} + c_2\Gamma_{n+1}^{(6)} \right\} \\ &\quad + 2V_\infty b_\infty / \underline{c}, \quad (16) \end{aligned}$$

where $c_0 = C_{\hat{e}}^{\text{Ret}}$, $c_1 = 9\mathcal{L}_\infty^{(2)}(\hat{e}_\infty(L_{\hat{e}}h_\infty + L) + V_\infty)\tilde{e}^{(3)}$, $c_2 = 486L(\mathcal{L}_\infty^{(2)})^2\tilde{e}^{(6)}$.

The proof is postponed to Appendix D.3.

As expected, the convergence of the retraction scheme with Markov noise demonstrate an analogous behavior as in the martingale noise setting. Particularly, using a step size of $\eta_k = \rho/\sqrt{k+k_0}$, the retraction scheme finds an $\mathcal{O}(b_\infty + \log(n)/\sqrt{n})$ -stationary point to (1) in $\mathcal{O}(n)$ iterations. Furthermore, we observe that the step size conditions and constants are improved when a second-order retraction is used in lieu of a first-order one.

Examples of Retraction Maps We present retraction maps on two common Riemannian manifolds. Importantly, we show that they satisfy the assumptions stated in Theorems 5 and 6.

Projective Retraction on the Sphere – Consider the Euclidean unit sphere manifold $\Theta = \mathbb{S}^d = \{x \in \mathbb{R}^{d+1} : \|x\| = 1\}$, where $\|\cdot\|$ stands for the standard Euclidean norm on \mathbb{R}^{d+1} . By [1, Example 3.5.1] for any $\theta \in \mathbb{S}^d$, $T_\theta \mathbb{S}^d = \{u \in \mathbb{R}^{d+1} : u^\top \theta = 0\}$. 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:

$$\text{Exp}_\theta(u) = \cos(\|u\|)\theta + \sin(\|u\|)(u/\|u\|). \quad (17)$$

The following result holds.

Proposition 7. *The projective retraction Ret_θ defined for any $(\theta, u) \in \text{TS}^d$ by*

$$\text{Ret}_\theta(u) = (\theta + u)/\|\theta + u\| \quad (18)$$

*satisfies **R1–R5** and for any $\theta \in \Theta$, Φ_θ defined by (11) has a bounded first derivative.*

The proof is postponed to Appendix D.4.

The retraction (18) is *both* a first-order and second-order retraction, and Φ_θ also has a bounded first-order derivative. Consequently, Theorem 5, & 6 & 14 can be applied according to conditions on the noise properties. We remark that by comparing (17) with (18), the retraction map Ret_θ has a better numerical stability as it does not involve evaluating the trigonometric functions.

Projective Retraction on the Grassmannian – Consider $\Theta = \text{Gr}_r(\mathbb{R}^d)$ as the real Grassmann manifold with its quotient space Riemannian metric [?, Section 2.3.2]. The manifold $\text{Gr}_r(\mathbb{R}^d)$ is the set of r -dimensional subspaces of \mathbb{R}^d . Each $\theta \in \text{Gr}_r(\mathbb{R}^d)$ can be represented by some $B \in \text{St}_r(\mathbb{R}^d) = \{B \in \mathbb{R}^{d \times r} : B^\top B = I_r\}$, where $\text{St}_r(\mathbb{R}^d)$ is the Stiefel manifold [1, Section 3.3.2]. Each element $\theta \in \text{Gr}_r(\mathbb{R}^d)$ is then denoted as $\theta = [B] = \text{Span}(B)$. By [?, Section 2.5], for any $\theta \in \text{Gr}_r(\mathbb{R}^d)$, with representative $B \in \text{St}_r(\mathbb{R}^d)$, the tangent space $T_\theta \text{Gr}_r(\mathbb{R}^d)$ is given by

$$T_\theta \text{Gr}_r(\mathbb{R}^d) = \left\{ u = B_\perp C : C \in \mathbb{R}^{(d-r) \times r} \right\}, \quad (19)$$

where $B_\perp \in \text{St}_{d-r}(\mathbb{R}^d)$ is such that $[B_\perp]$ is the orthogonal complement of $\theta = [B]$. The exponential map on $\text{Gr}_r(\mathbb{R}^d)$ corresponding to the metric $g_\theta(u, v) = \text{Tr}(u^\top v)$ for any $\theta \in \text{Gr}_r(\mathbb{R}^d)$, $u, v \in T_\theta \text{Gr}_r(\mathbb{R}^d)$, is

$$\text{Exp}_\theta(u) = \left[(B, B_\perp) \exp \left(\begin{array}{cc} 0 & -C^\top \\ C & 0 \end{array} \right) \left(\begin{array}{c} I_r \\ 0_{d-r \times r} \end{array} \right) \right]. \quad (20)$$

The following results hold.

Proposition 8. *The projective retraction Ret defined for any $(\theta, u) \in \text{T}\theta\text{Gr}_r(\mathbb{R}^d)$, by*

$$\text{Ret}_\theta(u) = [B + u] \quad (21)$$

*satisfies **R1–R5**. On the other hand, if $r > 1$, Φ_θ does not have bounded first derivative.*

The proof is postponed to Appendix **D.5**.

Similar to the projective retraction on spheres, the retraction (21) is *both* a first-order and second-order retraction, and the results Theorem **5** & **6** follow. We remark that by comparing (20) with (21), the retraction map is simpler to compute as it involves only a simple matrix addition.

5 Applications

In this section, we illustrate our convergence analysis results on three application examples of Riemannian SA scheme. They are subspace tracking method, stochastic semidefinite programming via low rank reparameterization, and robust barycenter problem.

5.1 Principal Component Analysis by Subspace Tracking

We consider a principal component analysis (PCA) problem in which we look for the r principal eigenvectors of a $d \times d$ covariance matrix \mathbf{A} , for which we have access to noisy data $(X_n)_{n \in \mathbb{N}}$ in \mathbb{R}^d . Note that this problem has been tackled in several works, see e.g. [**5**, **25**] and the references therein. PCA corresponds to the following minimization problem on the parameter space $\Theta = \text{Gr}_r(\mathbb{R}^d)$, introduced above,

$$\min_{\theta = [B] \in \text{Gr}_r(\mathbb{R}^d)} \left\{ f(\theta) = -\text{Tr}(B^\top \mathbf{A} B) / 2 \right\} .$$

The mean field corresponds to $h = \text{grad } f : \text{Gr}_r(\mathbb{R}^d) \rightarrow \text{TGr}_r(\mathbb{R}^d)$ and is given by, for any $\theta \in \text{Gr}_r(\mathbb{R}^d)$,

$$\text{grad } f(\theta) = (\mathbf{I}_d - BB^\top) \mathbf{A} B = -B_\perp B_\perp^\top \mathbf{A} B .$$

By writing $C = -B_\perp^\top \mathbf{A} B \in \mathbb{R}^{(d-r) \times r}$, note that we do have for any $\theta \in \text{Gr}_r(\mathbb{R}^d)$, $\text{grad } f(\theta) = B_\perp C \in \text{T}\theta\text{Gr}_r(\mathbb{R}^d)$. Since f is infinitely differentiable and $\text{Gr}_r(\mathbb{R}^d)$ is compact, **A2** is satisfied with $V = -f$. Therefore, using Proposition **8**, if we define for any $\theta \in \text{Gr}_r(\mathbb{R}^d)$ and $x \in \mathbb{R}^d$,

$$H_\theta(x) = B_\perp (-B_\perp^\top x x^\top B) ,$$

we can apply the results of Section **3** & **4** corresponding to either geodesic or retraction schemes for the sequence $(\theta_n)_{n \in \mathbb{N}}$ defined by (2), depending on the conditions of the data $(X_n)_{n \in \mathbb{N}}$.

Compared to the recent line of Riemannian based analysis of PCA, [**5**, **25**], our analysis is more flexible. Particularly, it holds for a general retraction scheme, in lieu of the computationally expensive exponential map.

5.2 Semidefinite Programming

Consider the semidefinite programming problem,

$$\min_{C \in \mathbb{R}^{d \times d}} \text{Tr}(AC), \text{ s.t. } C = C^\top, \text{diag}(C) = \mathbf{1}_d, \mathbf{C} \succeq \mathbf{0}, \quad (22)$$

where $A \in \mathbb{R}^{d \times d}$ is a symmetric (possibly non positive semidefinite) matrix, $\mathbf{1}_d = (\mathbf{1}, \dots, \mathbf{1}) \in \mathbb{R}^d$, and $\text{diag}(B) = (B_{1,1}, \dots, B_{d,d})$ for any matrix $B = (B_{i,j})_{i,j \in \{1, \dots, d\}}$. Problem (22) is a convex relaxation of several NP-hard problems such as MAXCUT [?] and community detection [?]. In these applications, A corresponds to the (weighted) adjacency or Laplacian matrix of a graph with d nodes. When d is large, solving (22) entails a high complexity since the SDP comprises of $\mathcal{O}(d^2)$ unknowns.

As shown in [? ?], for some instances, there exists a low-rank optimal solution C^* to (22) with $\text{rank}(C^*) \leq p$ for some $p \ll d$. This motivates the re-parameterization, $C = \theta\theta^\top$, $\theta \in \mathbb{R}^{d \times p}$ and therefore to consider the non-convex quadratic program,

$$\min_{\theta \in \mathbb{R}^{d \times p}} \left\{ f(\theta) = \text{Tr}(A\theta\theta^\top) \right\} \text{ s.t. } \text{diag}(\theta\theta^\top) = \mathbf{1}_d. \quad (23)$$

Due to the constraints of (23), this problem is therefore equivalent to consider the minimization of f on the Riemannian manifold $\Theta = \prod_{i=1}^d \mathbb{S}^{p-1}$ endowed with the product metric induced by the canonical one on \mathbb{S}^{p-1} , see [18, p. 20]. It is easy to show that gradient $\text{grad } f(\theta) \in \text{T}_\theta\Theta$ is given by

$$[\text{grad } f(\theta)]_{i,:} = c_i^\top \theta (\text{I}_p - [\theta]_{i,:}([\theta]_{i,:})^\top) \quad \text{for } i = 1, \dots, d, \quad (24)$$

where $[\theta]_{i,:}$ denotes the i th row vector of θ . In addition, similarly to the sphere we define the projective retraction Ret on $\text{T}\Theta$ for any $(\theta, u) \in \text{TS}^d$ by

$$[\text{Ret}_\theta(u)]_{i,:} = (\theta_{i,:} + u_{i,:}) / \|\theta_{i,:} + u_{i,:}\|, \quad \text{for } i = 1, \dots, d, \quad (25)$$

Following the same lines as the proof of Proposition 7, Ret defined by (25) satisfies **R1–R5** and for any $\theta \in \Theta$, Φ_θ defined by (11) has a bounded first derivative. Note that [?] analyze a deterministic version of the retraction scheme (2) that we consider to tackle (23). However, an exact gradient $\text{grad } f$ is required in their algorithm.

Evaluating the exact gradient (24) is computationally challenging when d is large and A is a dense matrix. As a remedy, a natural idea is to select $(A_n)_{n \in \mathbb{N}^*}$ as $A_n = A \odot X_n / \delta$, where $X_n \in \{0, 1\}^{d \times d}$ is a binary matrix with i.i.d. elements as $\mathbb{P}((X_n)_{1,1} = 1) = \delta$, for $\delta > 0$, and \odot denotes element-wise product.

Since f is infinitely differentiable and $\text{Gr}_r(\mathbb{R}^d)$ is compact, **A2** is satisfied with $V = -f$. Therefore, if we define for any $\theta \in \text{Gr}_r(\mathbb{R}^d)$ and $x \in \mathbb{R}^d$,

$$[H_\theta(x)]_{i,:} = [x]_{i,:} \theta (\text{I}_p - [\theta]_{i,:}([\theta]_{i,:})^\top), \quad \text{for } i = 1, \dots, d,$$

we can apply the results of Section 3 & 4 corresponding to either geodesic or retraction schemes for the sequence $(\theta_n)_{n \in \mathbb{N}}$ defined by (2), depending on the conditions we assume on the data $(A_n)_{n \in \mathbb{N}^*}$.

5.3 Robust Barycenter in a Hadamard Manifold

Let Θ be a Hadamard manifold – a simply-connected, complete Riemannian manifold of non-positive sectional curvature [18]. We assume that the sectional curvature of Θ is bounded below, $-\kappa^2 \leq \sec \Theta \leq 0$. A common example of this situation is $\Theta = \mathbb{S}_d^{++}$, the space of $d \times d$ symmetric positive-definite matrices, equipped with its affine-invariant metric [?] for which it holds $\kappa^2 = (1/8) d(d-1)(d+2)$ [?].

Consider a set of 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

$$\min_{\theta \in \Theta} \{V(\theta) = \int_{\Theta} \tilde{\rho}(\theta, x) \pi(dx)\}, \quad (26)$$

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 [? ?].

The Riemannian barycentre is known to be sensitive to outliers, which motivated the idea of considering the Riemannian median, obtained by taking $\tilde{\rho}(\theta, x) = \rho_{\Theta}(\theta, x)$ [?]. We consider a *robust barycenter* by using a Huber-like dissimilarity measure, $\tilde{\rho} : \Theta \times \Theta \rightarrow \mathbb{R}_+$,

$$\tilde{\rho}(\theta, \tau) = \delta^2 \left[1 + \{\rho_{\Theta}(\theta, \tau)/\delta\}^2 \right]^{1/2} - \delta^2, \quad (27)$$

where $\delta > 0$ is a cut-off constant. Observe that $\tilde{\rho}(\theta, \tau)$ behaves like $(1/2) \rho_{\Theta}^2(\theta, \tau)$ when $\rho_{\Theta}(\theta, \tau)$ is small compared to δ , and like $\delta \rho_{\Theta}(\theta, \tau)$ when $\rho_{\Theta}(\theta, \tau)$ is large compared to δ . Let π be a probability distribution on Θ . Using (27) in the optimization problem (26) yields a *robust barycenter* problem, and the robust barycenter is a global minimizer of (26).

In the simplest setting where $(X_n)_{n \in \mathbb{N}^*}$ are i.i.d. from π , tackling the problem (26) can be done by consider the following geodesic SA scheme:

$$\theta_{n+1} = \text{Exp}_{\theta_n} \left(\eta_{n+1} \frac{\text{Exp}_{\theta_n}^{-1}(X_{n+1})}{[1 + \{\rho_{\Theta}(\theta_n, X_{n+1})/\delta\}^2]^{1/2}} \right). \quad (28)$$

The above is essentially the “recursive barycenter” scheme proposed by [? ?], 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 confidence is assigned to extreme observations. We show in Appendix E that results from Section 3 can be applied and furthermore (28) finds a unique and global minimizer to (26). Note that our analysis can be easily extended to the case when $(X_n)_{n \in \mathbb{N}^*}$ is a Markov chain with stationary distribution of π and satisfying MA1.

Conclusions In this paper, we have provided a comprehensive study for the convergence analysis of Riemmanian SA schemes. For applications to possibly non-convex optimization

problems, our results generalize a number of recent works in the sense that the SA scheme is potentially biased and we allow the general retraction to be used in lieu of the complex geodesic operation. We also demonstrate our results on three examples motivated by machine learning applications.

References

- [1] P. Absil, R. Mahony, and R. Sepulchre. *Optimisation algorithms on matrix manifolds*. Princeton University Press, 2008.
- [2] P. Absil and J. Malick. Projection-like retractions on matrix manifolds. *SIAM Journal on Optimization*, 22(1):135–158, 2012.
- [3] J. Baxter and P. L. Bartlett. Infinite-horizon policy-gradient estimation. *J. Artif. Int. Res.*, 15(1):319–350, Nov. 2001.
- [4] A. Benveniste, M. Métivier, and P. Priouret. *Adaptive algorithms and stochastic approximations*, volume 22 of *Applications of Mathematics (New York)*. Springer-Verlag, Berlin, 1990. Translated from the French by Stephen S. Wilson.
- [5] S. Bonnabel. Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229, 2013.
- [6] N. Boumal, P.-A. Absil, and C. Cartis. Global rates of convergence for nonconvex optimization on manifolds. *IMA Journal of Numerical Analysis*, 39(1):1–33, 2018.
- [7] O. Cappé and E. Moulines. On-line expectation-maximization algorithm for latent data models. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 71(3):593–613, 2009.
- [8] C. Criscitiello and N. Boumal. Efficiently escaping saddle points on manifolds. In *NeurIPS*, 2019.
- [9] R. Douc, E. Moulines, P. Priouret, and P. Soulier. *Markov chains*. Springer, 2018.
- [10] M. Duflo. *Random Iterative Models*. Springer-Verlag, Berlin, Heidelberg, 1st edition, 1997.
- [11] G. Fort, E. Moulines, and P. Priouret. Convergence of adaptive and interacting Markov chain Monte Carlo algorithms. *Ann. Statist.*, 39(6):3262–3289, 2011.
- [12] S. Ghadimi and G. Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [13] R. Hosseini and S. Sra. An alternative to em for gaussian mixture models: Batch and stochastic riemannian optimization. *Mathematical Programming*, pages 1–37, 2019.

- [14] T. Jaakkola, M. I. Jordan, and S. P. Singh. Convergence of stochastic iterative dynamic programming algorithms. In *Proceedings of the 6th International Conference on Neural Information Processing Systems, NIPS'93*, pages 703–710, San Francisco, CA, USA, 1993. Morgan Kaufmann Publishers Inc.
- [15] B. Karimi, B. Miasojedow, E. Moulines, and H. Wai. Non-asymptotic analysis of biased stochastic approximation scheme. In *Conference on Learning Theory, COLT*, pages 1944–1974, 2019.
- [16] H. J. Kushner and G. G. Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, second edition, 2003. Stochastic Modelling and Applied Probability.
- [17] J. LaSalle. Some extensions of liapunov’s second method. *IRE Transactions on circuit theory*, 7(4):520–527, 1960.
- [18] J. Lee. *Introduction to Riemannian Manifolds*. Springer International Publishing, 2019.
- [19] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of mathematical Statistics*, 22(3):400–407, 1951.
- [20] S. M. Shah. Stochastic approximation on riemannian manifolds. *Applied Mathematics & Optimization*, May 2019.
- [21] T. Sun, Y. Sun, and W. Yin. On markov chain gradient descent. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9896–9905. Curran Associates, Inc., 2018.
- [22] Y. Sun, N. Flammarion, and M. Fazel. Escaping from saddle points on riemannian manifolds. In *NeurIPS*, 2019.
- [23] N. Tripuraneni, N. Flammarion, F. Bach, and M. I. Jordan. Averaging stochastic gradient descent on riemannian manifolds. In *Conference On Learning Theory, COLT*, pages 650–687, 2018.
- [24] H. Zhang, S. J. Reddi, and S. Sra. Riemannian svrg: Fast stochastic optimization on riemannian manifolds. In *Advances in Neural Information Processing Systems*, pages 4592–4600, 2016.
- [25] H. Zhang and S. Sra. First-order methods for geodesically convex optimization. In *Conference on Learning Theory, COLT*, pages 1617–1638, 2016.
- [26] J. Zhang, H. Zhang, and S. Sra. R-spider: A fast riemannian stochastic optimization algorithm with curvature independent rate. *arXiv preprint arXiv:1811.04194*, 2018.

A 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.

A.1 Metric Tensor and Distance

A smooth manifold (at least C^2) Θ is equipped with a smooth metric tensor field $g \in T^2T^*\Theta$, see [18, 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 [18, 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 [18, page 34]

$$\ell(c) = \int_I \|\dot{c}(t)\|_{c(t)} dt. \quad (29)$$

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 [18, Theorem 2.55],

$$\rho_\Theta(\theta, \theta') = \inf \{ \ell(c) \mid c : [0, 1] \rightarrow \Theta; c(0) = \theta, c(1) = \theta' \}, \quad \theta, \theta' \in \Theta. \quad (30)$$

This is called the Riemannian distance induced by the metric tensor g .

The infimum in (30) 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 [18, Corollary 6.21].

A.2 Levi-Civita Connection

In the Riemannian manifold, the curve γ which attains the infimum in (30) 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 ∇ [18, 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 X on Θ , the connection ∇ associates a vector $\nabla_u X \in T_\theta\Theta$. This vector is called the covariant derivative of X in the direction

of u . This is bilinear in u and X , and satisfies the product rule

$$\nabla_u(fX) = (uf)X(\theta) + f(\theta)\nabla_uX, \quad (31)$$

for any differentiable function $f : \Theta \rightarrow \mathbb{R}$, where uf denotes the derivative of f along u . Moreover, one has the following,

$$u \langle X, Y \rangle = \langle \nabla_uX, Y \rangle_\theta + \langle X, \nabla_uY \rangle_\theta, \quad (32)$$

$$\nabla_XY - \nabla_YX = [X, Y], \quad (33)$$

for any differentiable vector fields X, Y on Θ , where $[X, Y]$ is the Lie bracket of the vector fields X and Y , itself a vector field. Here, (32) states that ∇ is compatible with the metric, and (33) states that ∇ is a connection with zero torsion.

The Levi-Civita connection [18, Theorem 5.10] is defined as the unique affine connection ∇ which satisfies (32) and (33). Note that the uniqueness of this connection can be guaranteed by the Koszul's theorem, also known as the fundamental theorem of Riemannian geometry.

A.3 Geodesic Equation

It can be proved using (31) that ∇_uX depends only on the values of X along a curve tangent to the vector u [18, Proposition 4.26]. This motivates the following definition. Consider $c : I \rightarrow \Theta$ as a smooth curve on Θ and X is an extendible vector field along c , this means that $X : I \rightarrow T\Theta$ satisfies $X(t) \in T_{c(t)}\Theta$ for any $t \in I$, see [18, pages 100-101]. The covariant derivative of X along c is defined by

$$D_tX = \nabla_{\tilde{X}}\tilde{X} \circ c, \quad (34)$$

where \tilde{X} is a vector field on Θ satisfying, for any $t \in I$, $X(t) = \tilde{X} \circ c(t)$. 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 [18, page 103],

$$D_t\dot{\gamma}(t) = 0, \quad t \in I. \quad (35)$$

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 $\varepsilon > 0$ [18, 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 (30) [18, Theorem 6.19].

A.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 (34). Then, a vector field X is said to be parallel along c if it satisfies the parallel transport equation

$$D_t X(t) = 0 .$$

This is a first-order linear ordinary differential equation (ODE). Say $c(0) = \theta_0$, then, for a given initial condition $X(0) = u$, where $u \in T_{\theta_0} \Theta$, it follows that $u \mapsto X(t)$ is a linear mapping from $T_{\theta_0} \Theta$ to $T_{c(t)} \Theta$ [18, Theorem 4.32]. This is denoted T_{0t}^c , and by uniqueness of the solution to the ODE, $T_{t_0}^c$ is its linear inverse [18, Equation (4.22)].

It is useful to derive an equivalent condition to (32), which holds for vector fields X and Y along c [18, Proposition 5.5].

$$\frac{d}{dt} \langle X, Y \rangle_{c(t)} = \langle D_t X, Y \rangle_{c(t)} + \langle X, D_t Y \rangle_{c(t)} . \quad (36)$$

This equation yields that $t \mapsto \langle X, Y \rangle_{c(t)}$ is constant if X and 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 [18, 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

$$T_{0t}^c u = \sum_{i=1}^d u^i e_i(t) . \quad (37)$$

In other words, parallel transport is obtained by simply propagating the coordinates u^i of the vector u along a parallel orthonormal frame.

A.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 (30) is called a length-minimizing curve. While this curve is not always unique, it is always a geodesic [18, Theorem 6.4]; in other words, it is a twice differentiable solution of the geodesic equation (35). On the other hand, it is very 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 [18, 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, (author?) [18, page 128] define the *exponential map* $\text{Exp} : T\Theta \rightarrow \Theta$ as:

$$\text{Exp}_\theta(u) = \gamma_u(1) , \quad (38)$$

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 $[0, t]$ [18, page 307]:

$$c(u) = \sup \{t \geq 0 : \rho_\Theta(\theta, \gamma_u(t)) = t\} . \quad (39)$$

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 [18, 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 $\text{Exp}_\theta(u) = -\theta$.

A.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 Θ can be decomposed into the disjoint union of two sets

$$\Theta = D(\theta) \cup \text{Cut}(\theta) , \quad (40)$$

where $D(\theta)$ is the injectivity domain,

$$D(\theta) = \text{Exp}_\theta \{ \text{TD}(\theta) \} , \quad \text{where } \text{TD}(\theta) = \{tu : u \in T_\theta\Theta ; \|u\|_\theta = 1 \text{ and } 0 \leq t < c(u)\} ,$$

where $c(u)$ is given by (92). We observe:

Proposition 9. [18, 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 $D(\theta) = \Theta - \text{Cut}(\theta)$.*

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, (40) states that the topology of Θ is completely determined by $\text{Cut}(\theta)$. This theorem also ensures that $\text{Cut}(\theta)$ is a closed set of measure zero.

A.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_θ 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 [18, 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 [18, 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 (38), it is seen that

$$g \cdot \text{Exp}_\theta(u) = \text{Exp}_{g \cdot \theta}(g \cdot u) . \quad (41)$$

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 [?, 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 [?, page 67].

A.8 Riemannian Gradient, Hessian, and Taylor Formula

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 [18, Equation 2.14]

$$\langle \text{grad } V, u \rangle_\theta = DV(\theta)(u) , \quad u \in T_\theta\Theta , \quad (42)$$

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 $\text{Hess } V(\theta) : \text{T}_\theta \Theta \rightarrow \text{T}_\theta \Theta$,

$$\text{Hess } V(\theta) u = \nabla_u \text{grad } V(\theta) . \quad (43)$$

The Riemannian definition of the Hessian coincides with the covariant Hessian obtained from the Levi-Civita connection [18, Example 4.22]. The covariant characterization gives, for any vector fields X, Y on Θ ,

$$\langle \text{Hess } V X, Y \rangle = \nabla_X \nabla_Y V - \nabla_{\nabla_X Y} V = X(Y V) - (\nabla_X Y) V ,$$

where the last equality comes from the remark regarding directional derivatives. One can see that (33) yields that the linear operator $\text{Hess } V(\theta) : \text{T}_\theta \Theta \rightarrow \text{T}_\theta \Theta$ 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

$$\text{Hess } V(\theta)(u, v) = \langle \text{Hess } V(\theta) u, v \rangle_\theta , \quad u, v \in \text{T}_\theta \Theta .$$

Using the gradient (42) and the Hessian (43), one can derive the following Taylor formula for the function V . 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 ; then we have

$$V(\theta_1) - V(\theta_0) = \langle \text{grad } V(\theta_0), \dot{\gamma}(0) \rangle_{\theta_0} + \text{Hess } V(\gamma(t_*)) (\dot{\gamma}, \dot{\gamma}) / 2 ,$$

for some $t_* \in (0, 1)$.

A.9 Bounded Hessian Implies Lipschitz Gradient

We can now state a result that can be very useful when the Riemannian Hessian is bounded.

Lemma 10. *If V has a continuous Riemannian Hessian $\text{Hess } V(\theta) : \text{T}_\theta \Theta \rightarrow \text{T}_\theta \Theta$, with operator norm upper bounded uniformly for $\theta \in \Theta$, say by N ; then the Riemannian gradient $\text{grad } V$ satisfies the Lipschitz property (3) with Lipschitz constant $L = N$.*

Proof. Let $(e_i; i = 1, \dots, d)$ be a parallel orthonormal frame along γ . Define $\nabla V^i : [0, 1] \rightarrow \text{T}_\theta \Theta$ by

$$\nabla V^i(t) = \langle \text{grad } V(\gamma(t)), e_i(t) \rangle_{\gamma(t)} , \quad t \in [0, 1] . \quad (44)$$

Also, note from (37), applied to $c = \gamma$ and $u = \text{grad } V(\theta_0)$, that

$$\text{T}_{01}^\gamma \text{grad } V(\theta_0) = \sum_{i=1}^d \nabla V^i(0) e_i(1) .$$

Then, since $(e_i(1); i = 1, \dots, d)$ is an orthonormal basis of $T_{\theta_1} \Theta$,

$$\|\text{grad } V(\theta_1) - T_{01}^\gamma \text{grad } V(\theta_0)\|_{\theta_1}^2 = \sum_{i=1}^d (\nabla V^i(1) - \nabla V^i(0))^2 \quad (45)$$

But, by applying (36) to $c = \gamma$ with $X = \text{grad } V \circ \gamma$ and $Y = e_i$, it follows from (34), (43) and (44),

$$\frac{d}{dt} \nabla V^i(t) = \langle \text{Hess } V(\gamma(t)) \dot{\gamma}(t), e_i(t) \rangle_{\gamma(t)} + \langle \text{grad } V(\gamma(t)), D_t e_i(t) \rangle_{\gamma(t)} . \quad (46)$$

Then, since each $e_i(t)$ is parallel along γ , $D_t e_i(t) = 0$. Plugging (46) into (45) yields, by the mean-value theorem, followed by Jensen's inequality,

$$\begin{aligned} \|\text{grad } V(\theta_1) - T_{01}^\gamma \text{grad } V(\theta_0)\|_{\theta_1} &= \left(\sum_{i=1}^d \left[\int_0^1 \langle \text{Hess } V(\gamma(t)) \dot{\gamma}(t), e_i(t) \rangle_{\gamma(t)} dt \right]^2 \right)^{1/2} , \\ &\leq \int_0^1 \left(\sum_{i=1}^d \left[\langle \text{Hess } V(\gamma(t)) \dot{\gamma}(t), e_i(t) \rangle_{\gamma(t)} \right]^2 \right)^{1/2} dt , \\ &= \int_0^1 \|\text{Hess } V(\gamma(t)) \dot{\gamma}(t)\|_{\gamma(t)} dt . \end{aligned}$$

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 (29),

$$\|\text{grad } V(\theta_1) - T_{01}^\gamma \text{grad } V(\theta_0)\|_{\theta_1} \leq \int_0^1 \|\text{Hess } V(\gamma(t))\|_{op, \gamma(t)} \|\dot{\gamma}(t)\|_{\gamma(t)} dt \leq N \ell(\gamma) .$$

This concludes the proof. \square

B Proof of Lemma 1

Using $\gamma(0) = \theta_0$, $\gamma(1) = \theta_1$, by a Taylor expansion and the definition of the Riemannian gradient, 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 [18, Theorem 4.32] and that γ is a geodesic. Therefore, we obtain, using that the parallel

transport is a linear isometry [18, Proposition 5.5], that

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

where we have used that since γ is a geodesic $\mathbb{T}_{0t}^\gamma \dot{\gamma}(0) = \dot{\gamma}(t)$ and by [18, Corollary 5.6], $\|\dot{\gamma}(t)\|_{\gamma(t)} = \ell(\gamma)$, and $\ell(\gamma|_{[0,t]}) = t\ell(\gamma)$ by [18, Lemma 5.18]. For definitions of parallel transport and geodesic, see Appendix A.1 & A.4.

C Proofs of Section 3

C.1 Proof of Theorem 2

We begin the proof by observing the following lemma:

Lemma 11. *Assume A 1, A 2. Then for any $n \in \mathbb{N}^*$ and $\varepsilon > 0$,*

$$\begin{aligned}
& \sum_{k=0}^n \eta_{k+1} (\underline{c} - (3L/2)\eta_{k+1} - \bar{c}\varepsilon) \|h(\theta_k)\|_{\theta_k}^2 \leq V(\theta_0) - V(\theta_{n+1}) + (3L/2) \sum_{k=0}^n \eta_{k+1}^2 \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \\
& + \sum_{k=0}^n \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} + \sum_{k=0}^n \eta_{k+1} \left\{ (4\varepsilon)^{-1} + (3L/2)\eta_{k+1} \right\} \|b_{\theta_k}(X_{k+1})\|_{\theta_k}^2,
\end{aligned}$$

where e is defined by (5).

Proof. For any $k \geq 0$ and $t \in [0, 1]$, consider $\gamma_t^{(k)} = \text{Exp}_{\theta_k} \{ t\eta_{k+1}(H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})) \}$. Note that $\dot{\gamma}_0^{(k)} = \eta_{k+1} \{ H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) \}$ and $\ell(\gamma^{(k)}) = \eta_{k+1} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}$. Then, by Lemma 1, (5) 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 $k \in \mathbb{N}^*$,

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

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

$$\eta_{k+1}(\underline{c} - (3L/2)\eta_{k+1}) \|h(\theta_k)\|_{\theta_k}^2 \leq V(\theta_k) - V(\theta_{k+1}) + \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \\ + \eta_{k+1} \bar{c}^2 \varepsilon \|h(\theta_k)\|_{\theta_k}^2 + (\eta_{k+1}/4\varepsilon) \|b(\theta_k)\|_{\theta_k}^2 + (3L/2)\eta_{k+1}^2 \left\{ \|b(\theta_k)\|_{\theta_k}^2 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \right\}.$$

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

Equipped with the above lemma, the proof for Theorem 2 proceeds as follows. Let $n \in \mathbb{N}^*$. First note that for any $k \in \mathbb{N}$, $\mathbb{E}[\langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k}] = \mathbb{E}[\langle \text{grad } V(\theta_k), \mathbb{E}[e_{\theta_k}(X_{k+1}) | \mathcal{F}_k] \rangle_{\theta_k}] = 0$, using that θ_k is \mathcal{F}_k -measurable and **MD1**. Therefore, taking the expectation in the inequality given by Lemma 11 and using $\|b_{\theta_k}(X_{k+1})\|_{\theta_k} \leq b_\infty$ using **MD1**, we obtain

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

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

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

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

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

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

C.2 Proof of Theorem 3

Likewise in the proof for Theorem 2, we consider:

Lemma 12. *Assume **A 1-A 2-A 3-MA 1** hold. Let $(\eta_k)_{k \in \mathbb{N}^*}$ be a sequence satisfying (9). Then for any $n \in \mathbb{N}$,*

$$\mathbb{E} \left[- \sum_{k=0}^n \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \right] \leq D_{\hat{e}} \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k} \right] + \tilde{C}_{\hat{e}} \Gamma_{n+1}^{(2)} + C(\eta_1),$$

where $\tilde{C}_{\hat{e}} = C_{\hat{e}} - (3L/2)(b_{\infty}^2 + e_{\infty}^2)$, $D_{\hat{e}}, C_{\hat{e}}$ and $C(\eta_1)$ are given by (8).

Proof. Consider the measurable function $\hat{e} : \Theta \times X \rightarrow T\Theta$ which satisfies **MA1-(ii)-(iii)-(iv)** and for any $k \in \mathbb{N}^*$, consider $\gamma^{(k)} : [0, 1] \rightarrow \Theta$ the geodesic between θ_{k-1} and θ_k for any $k \in \mathbb{N}^*$ 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 $k \in \mathbb{N}^*$,

$$\ell(\gamma^{(k)}) = \eta_k \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}. \quad (47)$$

Using that the parallel transport associated with the Levi-Civita connection is a linear isometry [18, Proposition 5.5] and $(T_{01}^{\gamma})^{-1} = T_{10}^{\gamma}$ by uniqueness of parallel transport [18, Theorem 4.32], we obtain the following decomposition

$$\mathbb{E} \left[- \sum_{k=0}^n \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \right] = -\mathbb{E} \left[\sum_{i=1}^5 A_i \right], \quad (48)$$

where

$$\begin{aligned} A_1 &= \sum_{k=1}^n \eta_{k+1} \langle \text{grad } V(\theta_k), \hat{e}_{\theta_k}(X_{k+1}) - P_{\theta_k} \hat{e}_{\theta_k}(X_k) \rangle_{\theta_k}, \\ A_2 &= \sum_{k=1}^n \eta_{k+1} \langle \text{grad } V(\theta_k), P_{\theta_k} \hat{e}_{\theta_k}(X_k) - T_{01}^{\gamma^{(k)}} P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \rangle_{\theta_k}, \\ A_3 &= \sum_{k=1}^n \eta_{k+1} \langle T_{01}^{\gamma^{(k)}} T_{10}^{\gamma^{(k)}} \text{grad } V(\theta_k) - T_{01}^{\gamma^{(k)}} T_{10}^{\gamma^{(k)}} T_{01}^{\gamma^{(k)}} \text{grad } V(\theta_{k-1}), T_{01}^{\gamma^{(k)}} T_{10}^{\gamma^{(k)}} T_{01}^{\gamma^{(k)}} P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \rangle_{\theta_k} \\ &= \sum_{k=1}^n \eta_{k+1} \langle T_{10}^{\gamma^{(k)}} \text{grad } V(\theta_k) - \text{grad } V(\theta_{k-1}), P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \rangle_{\theta_{k-1}}, \\ A_4 &= \sum_{k=1}^n (\eta_{k+1} - \eta_k) \langle \text{grad } V(\theta_{k-1}), P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \rangle_{\theta_{k-1}}, \\ A_5 &= \eta_1 \langle \text{grad } V(\theta_0), \hat{e}_{\theta_0}(X_1) \rangle_{\theta_0} + \eta_{n+1} \langle \text{grad } V(\theta_n), \hat{e}_{\theta_n}(X_{n+1}) \rangle_{\theta_n}. \end{aligned} \quad (49)$$

We now bound each terms of this decomposition. First note that using **MA 1-(i)-(iii)**, we get that for any $k \in \{1, \dots, n\}$, $\mathbb{E}[\langle \text{grad } V(\theta_k), \hat{e}_{\theta_k}(X_{k+1}) - P_{\theta_k} \hat{e}_{\theta_k}(X_k) \rangle_{\theta_k} | \mathcal{F}_k] = 0$ and therefore

$$\mathbb{E}[A_1] = 0. \quad (50)$$

Using Cauchy-Schwarz inequality, **MA 1-(iii)-(iv)**, (47) and **A 2-(a)**, we get

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

Using $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (9), **A 3**, for any $a, b \in \mathbb{R}$, $a \leq a^2 + 1$ and $|ab| \leq (a^2 + b^2)/2$, we get

$$|A_2| \leq L_{\hat{e}} \bar{c} \left\{ \sum_{k=1}^n \eta_k^2 (e_{\infty} + b_{\infty}) + \sum_{k=1}^n \eta_k^2 \|h(\theta_k)\|_{\theta_k}^2 (1 + e_{\infty} + b_{\infty}) \right\}. \quad (51)$$

Using **A 2-(b)**, Lemma 1, **A 3**, **MA 1-(iii)**, we obtain

$$\begin{aligned} |A_3| &\leq L \sum_{k=1}^n \eta_k \eta_{k+1} \left\{ \|H_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} + \|b_{\theta_{k-1}}(X_k)\|_{\theta_{k-1}} \right\} \left\| P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \right\|_{\theta_{k-1}}, \\ &\leq L \hat{e}_{\infty} \sum_{k=1}^n \eta_k \eta_{k+1} (e_{\infty} + b_{\infty} + \|h(\theta_{k-1})\|_{\theta_{k-1}}), \\ &\leq L \hat{e}_{\infty} \left\{ (e_{\infty} + b_{\infty} + 1) \sum_{k=1}^n \eta_k^2 + \sum_{k=1}^n \eta_k^2 \|h(\theta_{k-1})\|_{\theta_{k-1}}^2 \right\}. \end{aligned} \quad (52)$$

Using the Cauchy-Schwarz inequality, **A 2-(a)**, **MA 1-(iii)**, for any $a \in \mathbb{R}$, $a \leq 1 + a^2$ and that $(\eta_k)_{k \in \mathbb{N}^*}$ satisfies (9), we have

$$\begin{aligned} |A_4| &\leq \bar{c} \sum_{k=1}^n |\eta_{k+1} - \eta_k| \|h(\theta_{k-1})\|_{\theta_{k-1}} \left\| P_{\theta_{k-1}} \hat{e}_{\theta_{k-1}}(X_k) \right\|_{\theta_{k-1}}, \\ &\leq \bar{c} \hat{e}_{\infty} \sum_{k=1}^n |\eta_{k+1} - \eta_k| \|h(\theta_{k-1})\|_{\theta_{k-1}} \leq \bar{c} \hat{e}_{\infty} \left\{ \eta_1 + a_2 \sum_{k=1}^n \eta_k^2 \|h(\theta_{k-1})\|_{\theta_{k-1}}^2 \right\}. \end{aligned} \quad (53)$$

Finally, by the Cauchy-Schwarz inequality, **A 2-(a)**, **MA 1-(iii)** and using that $(\eta_k)_{k \in \mathbb{N}^*}$ is nonincreasing, we obtain that

$$\begin{aligned} |A_5| &\leq \bar{c} \hat{e}_{\infty} \left\{ \eta_1 \|h(\theta_0)\|_{\theta_0} + \eta_{n+1} \|h(\theta_n)\|_{\theta_n} \right\}, \\ &\leq \bar{c} \hat{e}_{\infty} \left\{ 2 + \eta_1^2 \|h(\theta_0)\|_{\theta_0}^2 + \eta_{n+1}^2 \|h(\theta_n)\|_{\theta_n}^2 \right\}. \end{aligned} \quad (54)$$

Combining (50)-(51)-(52)-(53)-(54) in (48) completes the proof. \square

Proof of Theorem 3. The proof only consists in applying Lemma 11 taking $\varepsilon = \underline{c}/(2\bar{c})$, Lemma 12 and using the last inequality condition in (9). □

D Omitted Proofs for Section 4

D.1 Proof of Lemma 4

We preface the proof of the Lemma by a preliminary result which does not assume **R3**.

Lemma 13. *Assume **R1**, **R2** hold.*

(a) Under **R4**, $\|\Phi_\theta(u) - u\|_\theta \leq \mathcal{L}^{(1)}(\theta)\|u\|_\theta^2/2$ for any $(\theta, u) \in \text{T}\Theta$.

(b) Under **R5**, $\|\Phi_\theta(u) - u\|_\theta \leq \mathcal{L}^{(2)}(\theta)\|u\|_\theta^3/6$ for any $(\theta, u) \in \text{T}\Theta$.

Proof. (a) Let $\theta \in \Theta$ and for any $u \in \text{T}_\theta\Theta$, consider the first-order Taylor expansion of $\Phi_\theta : \text{T}_\theta\Theta \rightarrow \text{T}_\theta\Theta$, taken at 0_θ ,

$$\Phi_\theta(u) = \Phi_\theta(0_\theta) + \text{D}\Phi_\theta(0_\theta)[u] + \int_0^1 (1-t) \text{D}^2\Phi_\theta(tu)[u, u] dt, \quad (55)$$

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

$$\Phi_\theta(0_\theta) = \text{Exp}_\theta^{-1} \circ \text{Ret}_\theta(0_\theta) = \text{Exp}_\theta^{-1}(\theta) = 0_\theta, \quad (56)$$

where the second equality follows because $\text{Ret}_\theta(0_\theta) = \theta$, by **R1** 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, we obtain since $\text{DExp}_\theta(0_\theta) = \text{Id}$ by definition and using **R1** that

$$\text{D}\Phi_\theta(0_\theta) = \text{DExp}_\theta^{-1}(\text{Ret}_\theta(0_\theta))\text{DRet}_\theta(0_\theta) = \text{DExp}_\theta^{-1}(\theta)\text{DRet}_\theta(0_\theta) = \text{Id}. \quad (57)$$

The proof is then completed using (56), (57) and **R4** in (55).

(b) Let $\theta \in \Theta$ and consider the second-order Taylor expansion of $\Phi_\theta : \text{T}_\theta\Theta \rightarrow \text{T}_\theta\Theta$, taken at 0_θ :

$$\Phi_\theta(u) = u + \text{D}^2\Phi_\theta(0_\theta)[u, u]/2 + 2^{-1} \int_0^1 (1-t)^2 \text{D}^3\Phi_\theta(tu)[u, u, u] dt. \quad (58)$$

where $\text{D}^3\Phi_\theta$ is the third derivative of Φ_θ . The proof relies on the use of normal coordinates with origin at θ [18, Chapter 5]. These coordinates are smooth and simply defined identifying $\text{T}_\theta\Theta$ with \mathbb{R}^d through Exp_θ^{-1} . More precisely, setting an orthonormal basis $\{\mathbf{b}_i : i \in \{1, \dots, d\}\}$ of $\text{T}_\theta\Theta$, define for any $\tilde{\theta} \notin \text{Cut}(\theta)$,

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

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 [18, 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,

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

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 $\{\partial_k : k \in \{1, \dots, d\}\}$ coordinate vector fields on $T\Theta$ corresponding to φ . But using [18, 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,

$$D_t \dot{\gamma}(0) = D^2 \Phi_\theta(0)[u, u] \quad (59)$$

and by **R5**, $D_t \dot{\gamma}(0) = 0$, which implies that $D^2 \Phi_\theta(0)[u, u] = 0_\theta$. Plugging this result into (58) and using the bound on the third derivative of Φ_θ given by **R5** completes the proof. \square

Proof of Lemma 4. (a) It will now be proven that $\mathcal{L}^{(1)}(\theta)$ can be chosen independent of θ . To do so, fix some $o \in \Theta$ and note from Lemma 13-(a),

$$\|\Phi_o(v) - v\|_o \leq \mathcal{L}^{(1)}(o) \|v\|_o^2 / 2, \quad \text{for any } v \in T_o\Theta. \quad (60)$$

By **R3**, Θ is a homogeneous Riemannian manifold under the action of \mathbf{G} . Thus, for any $\theta \in \Theta$ there exists $g \in \mathbf{G}$ such that $g^{-1} \cdot \theta = o$. Taking $v = g^{-1} \cdot u$ in (60), and noting that g is an isometry (so g preserves Riemannian norms), it follows

$$\|\Phi_\theta(u) - u\|_\theta = \left\| g^{-1} \cdot (\Phi_\theta(u) - u) \right\|_o = \left\| g^{-1} \cdot \Phi_\theta(u) - v \right\|_o. \quad (61)$$

However, it will shortly be shown that

$$g^{-1} \cdot \Phi_\theta(u) = \Phi_o(v). \quad (62)$$

Replacing this into (61),

$$\|\Phi_\theta(u) - u\|_\theta = \|\Phi_o(v) - v\|_o \leq \mathcal{L}^{(1)}(o) \|v\|_o^2 / 2 = \mathcal{L}^{(1)}(o) \|u\|_\theta^2 / 2. \quad (63)$$

where the inequality follows from (60), and the final equality because $v = g^{-1} \cdot u$ and g is an isometry. Clearly, (63) is identical to the required inequality, with $\mathcal{L}^{(1)} = \mathcal{L}^{(1)}(o)/2$ (which is independent of θ). Now, to complete the proof it only remains to show (62). To do so, let $\rho = \text{Ret}_\theta(u)$ and note that

$$g^{-1} \cdot \Phi_\theta(u) = g^{-1} \cdot \text{Exp}_\theta^{-1}(\rho). \quad (64)$$

Because g is an isometry, so g maps geodesics to geodesics, it follows that (see (41) in A.7)

$$g^{-1} \cdot \text{Exp}_\theta^{-1}(\rho) = \text{Exp}_o^{-1}(g^{-1} \cdot \rho). \quad (65)$$

However, from (13) in **R3** and the definition of ρ ,

$$g^{-1} \cdot \rho = \text{Ret}_o(g^{-1} \cdot u) = \text{Ret}_o(v) .$$

Replacing this into (65), it is seen than

$$g^{-1} \cdot \text{Exp}_\theta^{-1}(\rho) = \text{Exp}_o^{-1}(\text{Ret}_o(v)) = \Phi_o(v) .$$

Then, (62) follows from (64).

(b) the proof is almost identical to item (a). First, fixe some $o \in \Theta$ and note from Lemma 13-(b),

$$\|\Phi_o(v) - v\|_o \leq \mathcal{L}^{(2)}(o) \|v\|_o^3 / 6 , \text{ for any } v \in T_o\Theta . \quad (66)$$

Then, for any $\theta \in \Theta$, write $g^{-1} \cdot \theta = o$. Taking $v = g^{-1} \cdot u$ in (66) and noting that g is an isometry, it follows

$$\|\Phi_\theta(u) - u\|_\theta = \left\| g^{-1} \cdot (\Phi_o(u) - u) \right\|_\theta = \left\| g^{-1} \cdot \Phi_o(u) - v \right\|_o .$$

Then, using (62),

$$\|\Phi_\theta(u) - u\|_\theta = \|\Phi_o(v) - v\|_o \leq \mathcal{L}^{(2)}(o) \|v\|_o^3 / 6 = \mathcal{L}^{(2)}(o) \|u\|_\theta^2 / 6 . \quad (67)$$

where the inequality follows from (66), and the final equality because $v = g^{-1} \cdot u$ and g is an isometry. Clearly, (67) is identical to the required inequality, with $\mathcal{L}^{(1)} = \mathcal{L}^{(1)}(o)/2$ (which is independent of θ). \square

D.2 Proof of Theorem 5

Before we begin proving Theorem 5, let us state a similar result to Theorem 5-(a), with slightly different assumptions.

Theorem 14. *Assume **A1-A2-A3-A4**, **R1-R2-R3-R4**, **MD1-MD2(2)** hold. Consider $(\theta_n)_{n \in \mathbb{N}}$ defined by (2). For any $n \in \mathbb{N}^*$ and let $I_n \in \{0, \dots, n\}$ be distributed according to (4). Assume for any $\theta \in \Theta$, the function Φ_θ defined by (11) has first derivative bounded by $\mathcal{L}_\infty^{(0)}$. If $\sup_{k \in \mathbb{N}^*} \eta_k \leq (4L + 6\mathcal{L}_\infty^{(1)}V_\infty)^{-1}\underline{c}$, then*

$$\mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \leq 2(\underline{c}\Gamma_{n+1})^{-1} \left\{ \mathbb{E}[V(\theta_0)] + c_0\Gamma_{n+1}^{(2)} + c_1\Gamma_{n+1}^{(4)} \right\} + 2b_\infty V_\infty / \underline{c} ,$$

where $c_0 = 2L(5 + 8(\mathcal{L}_\infty^{(0)})^2)(\sigma_{(2)} + b_\infty^2)$, $c_1 = 8L(\mathcal{L}_\infty^{(1)})^2 h_\infty^4$.

The main idea of our proof is to exploit the geodesic expression (12) of the retraction stochastic approximation scheme. We begin by applying Lemma 1 to (12). Note that as

$$\ell(\gamma^{(k)}) = \eta_{k+1} \left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k} ,$$

we have

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

The above implies

$$\begin{aligned} -\eta_{k+1} \langle \text{grad } V(\theta_k), h(\theta_k) \rangle_{\theta_k} & \leq V(\theta_k) - V(\theta_{k+1}) + \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \\ & \quad + \eta_{k+1} \left\langle \text{grad } V(\theta_k), b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\rangle_{\theta_k} \\ & \quad + \left(L\eta_{k+1}^2/2 \right) \left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 . \end{aligned} \tag{68}$$

Applying **A2-(a)** to the left hand side yields the lower bound to the left hand side:

$$-\eta_{k+1} \langle \text{grad } V(\theta_k), h(\theta_k) \rangle_{\theta_k} \geq \eta_{k+1} \underline{c} \|h(\theta_k)\|_{\theta_k}^2 .$$

Therefore taking the conditional expectation $\mathbb{E}[\cdot | \mathcal{F}_k]$ on both sides of (68) shows that:

$$\begin{aligned} & \eta_{k+1} \underline{c} \|h(\theta_k)\|_{\theta_k}^2 \\ & \leq V(\theta_k) - V(\theta_{k+1}) + \eta_{k+1} \mathbb{E} \left[\left\langle \text{grad } V(\theta_k), b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\rangle_{\theta_k} \middle| \mathcal{F}_k \right] \\ & \quad + \left(L\eta_{k+1}^2/2 \right) \mathbb{E} \left[\left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \middle| \mathcal{F}_k \right] . \end{aligned} \tag{69}$$

Next, we consider different cases of retraction and prove their corresponding bounds for the last two terms in (69).

First-order Retraction Consider either case (a) in Theorem 5 or Theorem 14, where **R4**, **MD2**(2 or 4) hold. Applying the Cauchy-Schwarz inequality and **A4** give

$$\begin{aligned} \left\langle \text{grad } V(\theta_k), \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\rangle_{\theta_k} & \leq \|\text{grad } V(\theta_k)\|_{\theta_k} \left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k} , \\ & \leq V_\infty \left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k} . \end{aligned}$$

Invoking Lemma 4-(a) shows that

$$\begin{aligned} \left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k} & \leq \mathcal{L}_\infty^{(1)} \eta_{k+1} \left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) \right\|_{\theta_k}^2 , \\ & \leq 3\mathcal{L}_\infty^{(1)} \eta_{k+1} \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \right\} . \end{aligned} \tag{70}$$

Therefore, the above inequalities give

$$\begin{aligned} & \eta_{k+1} \mathbb{E} \left[\left\langle \text{grad } V(\theta_k), b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\rangle_{\theta_k} \middle| \mathcal{F}_k \right] \\ & \leq \eta_{k+1} V_\infty b_\infty + 3\eta_{k+1}^2 V_\infty \mathcal{L}_\infty^{(1)} \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + \sigma_{(2)} \right\}. \end{aligned} \quad (71)$$

Next, we consider

$$\begin{aligned} & \mathbb{E} \left[\left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \middle| \mathcal{F}_k \right] \\ & \leq 4 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + \sigma_{(2)} + \mathbb{E} \left[\left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \middle| \mathcal{F}_k \right] \right\}. \end{aligned} \quad (72)$$

and focus on the following sub-cases with first order retraction:

Sub-case when MD2(2) holds and Φ_θ has bounded first order derivative — Let $D\Phi_{\theta_k}(x)[u]$ be the first-order derivative of Φ_{θ_k} at $x \in T_{\theta_k} \Theta$ in the direction of $u \in T_{\theta_k} \Theta$. Observe

$$\begin{aligned} & \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \\ & = \eta_{k+1}^{-1} \{ \Phi_{\theta_k}(\eta_{k+1} \{ H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) \}) - \Phi_{\theta_k}(\eta_{k+1} h(\theta_k)) \} \\ & \quad + \eta_{k+1}^{-1} \{ \Phi_{\theta_k}(\eta_{k+1} h(\theta_k)) - \eta_{k+1} \{ H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) \} \}, \\ & = \int_0^1 D\Phi_{\theta_k}(\eta_{k+1}(h(\theta_k) + t(e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})))) [e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})] dt \\ & \quad + \eta_{k+1}^{-1} \{ \Phi_{\theta_k}(\eta_{k+1} h(\theta_k)) - \eta_{k+1} h(\theta_k) - \eta_{k+1} \{ e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) \} \}. \end{aligned}$$

where the last equality is due to the mean value theorem. Together with Lemma 4, the above is bounded by using Hölder's inequality, and bounding each term. The integral is treated as follows

$$\begin{aligned} & \left\| \int_0^1 D\Phi_{\theta_k}(\eta_{k+1}(h(\theta_k) + t(e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})))) [e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})] dt \right\|_{\theta_k}^2 \\ & \leq (\mathcal{L}_\infty^{(0)})^2 \|e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^2 \leq 2(\mathcal{L}_\infty^{(0)})^2 (\sigma_{(2)} + b_\infty^2), \end{aligned}$$

where $\mathcal{L}_\infty^{(0)}$ upper bounds the operator norm of $D\Phi_{\theta_k}(x)$. Dealing with the other terms is easier, yielding

$$\mathbb{E} \left[\left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \middle| \mathcal{F}_k \right] \leq 4 \left\{ (1 + 2(\mathcal{L}_\infty^{(0)})^2)(\sigma_{(2)} + b_\infty^2) + \eta_{k+1}^2 (\mathcal{L}_\infty^{(1)})^2 h_\infty^4 \right\}.$$

Substituting the above into (72), and combining it with (71) into (69) leads to:

$$\begin{aligned} & \eta_{k+1} (\underline{c} - (2L + 3\mathcal{L}_\infty^{(1)} V_\infty) \eta_{k+1}) \|h(\theta_k)\|_{\theta_k}^2 \\ & \leq V(\theta_k) - \mathbb{E}[V(\theta_{k+1}) | \mathcal{F}_k] + \eta_{k+1} V_\infty b_\infty + \eta_{k+1}^2 \{ (3\mathcal{L}_\infty^{(1)} V_\infty + 2L(5 + 8(\mathcal{L}_\infty^{(0)})^2)) (\sigma_{(2)} + b_\infty^2) \} \\ & \quad + \eta_{k+1}^4 8L(\mathcal{L}_\infty^{(1)})^2 h_\infty^4. \end{aligned}$$

As the step size is chosen such that

$$\underline{c} - (2L + 3\mathcal{L}_\infty^{(1)}V_\infty)\eta_{k+1} \geq \underline{c}/2 ,$$

summing up the above inequality from $k = 0$ to $k = n$ gives:

$$\begin{aligned} & (\underline{c}/2) \sum_{k=0}^n \eta_{k+1} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\ & \leq \mathbb{E}[V(\theta_0)] + V_\infty b_\infty \Gamma_{n+1} + \{(3\mathcal{L}_\infty^{(1)}V_\infty + 2L(5 + 8(\mathcal{L}_\infty^{(0)})^2))(\sigma_{(2)} + b_\infty^2)\} \Gamma_{n+1}^{(2)} + 8L(\mathcal{L}_\infty^{(1)})^2 h_\infty^4 \Gamma_{n+1}^{(4)} . \end{aligned}$$

Using that $V(\theta_{n+1}) \geq 0$ yields

$$\begin{aligned} & \mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \\ & \leq (2/\underline{c}) \left\{ \mathbb{E}[V(\theta_0)]/\Gamma_{n+1} + \{(3\mathcal{L}_\infty^{(1)}V_\infty + 2L(5 + 8(\mathcal{L}_\infty^{(0)})^2))(\sigma_{(2)} + b_\infty^2)\} \Gamma_{n+1}^{(2)}/\Gamma_{n+1} \right. \\ & \quad \left. + 8L(\mathcal{L}_\infty^{(1)})^2 h_\infty^4 \Gamma_{n+1}^{(4)}/\Gamma_{n+1} + V_\infty b_\infty \right\} . \end{aligned}$$

The above inequality yields the desired bound in Theorem 14.

Sub-case when MD2(4) holds — invoking Lemma 4-(a) we observe that

$$\begin{aligned} \mathbb{E} \left[\left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \middle| \mathcal{F}_k \right] & \leq \eta_{k+1}^2 (\mathcal{L}_\infty^{(1)})^2 \mathbb{E} \left[\|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^4 \middle| \mathcal{F}_k \right] , \\ & \leq 27\eta_{k+1}^2 (\mathcal{L}_\infty^{(1)})^2 \{h_\infty^4 + b_\infty^4 + \sigma_{(4)}\} . \end{aligned}$$

where we have used $H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) = h(\theta_k) + e_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})$ and **A4** in the last inequality. Substituting the above into (69) leads to:

$$\begin{aligned} & \eta_{k+1} (\underline{c} - (2L + 3\mathcal{L}_\infty^{(1)}V_\infty)\eta_{k+1}) \|h(\theta_k)\|_{\theta_k}^2 \\ & \leq V(\theta_k) - \mathbb{E}[V(\theta_{k+1})|\mathcal{F}_k] + \eta_{k+1} V_\infty b_\infty + \eta_{k+1}^2 (3\mathcal{L}_\infty^{(1)}V_\infty + 2L)\{b_\infty^2 + \sigma_{(2)}\} \\ & \quad + 54L(\mathcal{L}_\infty^{(1)})^2 \eta_{k+1}^4 \{h_\infty^4 + b_\infty^4 + \sigma_{(4)}\} . \end{aligned}$$

As we have set

$$\underline{c} - (2L + 3\mathcal{L}_\infty^{(1)}V_\infty)\eta_{k+1} \geq \underline{c}/2 ,$$

taking the full expectation and summing up the previous inequality from $k = 0$ to $k = n$ leads to

$$\begin{aligned} & (\underline{c}/2) \sum_{k=0}^n \eta_{k+1} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\ & \leq \mathbb{E}[V(\theta_0) - V(\theta_{n+1})] + V_\infty b_\infty \Gamma_{n+1} + \{(3\mathcal{L}_\infty^{(1)}V_\infty + 2L)(\sigma_{(2)} + b_\infty^2)\} \Gamma_{n+1}^{(2)} \\ & \quad + 54L(\mathcal{L}_\infty^{(1)})^2 \{h_\infty^4 + b_\infty^4 + \sigma_{(4)}\} \Gamma_{n+1}^{(4)} . \end{aligned}$$

Hence, using that $V(\theta_{n+1}) \geq 0$ for any $n \in \mathbb{N}$,

$$\begin{aligned} & (\underline{c}/2)\mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \\ & \leq \mathbb{E} [V(\theta_0)]/\Gamma_{n+1} + (3\mathcal{L}_\infty^{(1)}V_\infty + 2L)\{\sigma_{(2)} + b_\infty^2\} \Gamma_{n+1}^{(2)}/\Gamma_{n+1} \\ & \quad + 54L(\mathcal{L}_\infty^{(1)})^2\{h_\infty^4 + b_\infty^4 + \sigma_{(4)}\} \Gamma_{n+1}^{(4)}/\Gamma_{n+1} + V_\infty b_\infty. \end{aligned}$$

The above inequality yields the desirable bound in case (a).

Second-order Retraction We consider the case where **R 5**, **MD 2(6)** hold. Invoking Lemma 4-(b), we have that

$$\begin{aligned} \left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k} & \leq \eta_{k+1}^2 \mathcal{L}_\infty^{(2)} \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^3, \\ & \leq 9\eta_{k+1}^2 \mathcal{L}_\infty^{(2)} \{ \|e_{\theta_k}(X_{k+1})\|_{\theta_k}^3 + h_\infty^3 + b_\infty^3 \}. \end{aligned} \quad (73)$$

Thus,

$$\mathbb{E} \left[\eta_{k+1} \left\langle \text{grad } V(\theta_k), \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\rangle_{\theta_k} \mid \mathcal{F}_k \right] \leq 9\mathcal{L}_\infty^{(2)} V_\infty \eta_{k+1}^3 \{h_\infty^3 + \sigma_{(3)} + b_\infty^3\}.$$

Moreover, since

$$\begin{aligned} \mathbb{E} \left[\left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \mid \mathcal{F}_k \right] & \leq \eta_{k+1}^4 (\mathcal{L}_\infty^{(2)})^2 \mathbb{E} \left[\|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k}^6 \mid \mathcal{F}_k \right], \\ & \leq 3^5 \eta_{k+1}^4 (\mathcal{L}_\infty^{(2)})^2 \{h_\infty^6 + b_\infty^6 + \sigma_{(6)}\}, \end{aligned}$$

we have

$$\begin{aligned} & \mathbb{E} \left[\left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \mid \mathcal{F}_k \right] \\ & \leq 4 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + \sigma_{(2)} + \mathbb{E} \left[\left\| \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \mid \mathcal{F}_k \right] \right\}, \\ & \leq 4 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + \sigma_{(2)} + 243(\mathcal{L}_\infty^{(2)})^2 \eta_{k+1}^4 (h_\infty^6 + b_\infty^6 + \sigma_{(6)}) \right\}. \end{aligned}$$

Substituting the above into (69) leads to:

$$\begin{aligned} & \eta_{k+1} \{ \underline{c} - 2L\eta_{k+1} \} \|h(\theta_k)\|_{\theta_k}^2 \\ & \leq V(\theta_k) - \mathbb{E} [V(\theta_{k+1}) \mid \mathcal{F}_k] + \eta_{k+1} V_\infty b_\infty + 9\mathcal{L}_\infty^{(2)} V_\infty \eta_{k+1}^3 \{h_\infty^3 + \sigma_{(3)} + b_\infty^3\} \\ & \quad + 2L\eta_{k+1}^2 \left\{ b_\infty^2 + \sigma_{(2)} + 3^5 (\mathcal{L}_\infty^{(2)})^2 \eta_{k+1}^4 (h_\infty^6 + b_\infty^6 + \sigma_{(6)}) \right\}. \end{aligned}$$

As we have set

$$\underline{c} - 2L\eta_{k+1} \geq (\underline{c}/2),$$

taking the full expectation and summing up both sides of the previous inequality from $k = 0$ to $k = n$ yields

$$\begin{aligned} & (\underline{c}/2) \sum_{k=0}^n \eta_{k+1} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k} \right] \\ & \leq \mathbb{E} [V(\theta_0) - V(\theta_{n+1})] + V_\infty b_\infty \Gamma_{n+1} + 2L\{b_\infty^2 + \sigma_{(2)}\} \Gamma_{n+1}^{(2)} \\ & \quad + 9\mathcal{L}_\infty^{(2)} V_\infty \{h_\infty^3 + \sigma_{(3)} + b_\infty^3\} \Gamma_{n+1}^{(3)} + 486L(\mathcal{L}_\infty^{(2)})^2 (h_\infty^6 + b_\infty^6 + \sigma_{(6)}) \Gamma_{n+1}^{(6)}. \end{aligned}$$

As such,

$$\begin{aligned} & (\underline{c}/2) \mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \\ & \leq \mathbb{E} [V(\theta_0)] / \Gamma_{n+1} + 2L\{b_\infty^2 + \sigma_{(2)}\} \Gamma_{n+1}^{(2)} / \Gamma_{n+1} + 9\mathcal{L}_\infty^{(2)} V_\infty \{h_\infty^3 + \sigma_{(3)} + b_\infty^3\} \Gamma_{n+1}^{(3)} / \Gamma_{n+1} \\ & \quad + 486L(\mathcal{L}_\infty^{(2)})^2 \{h_\infty^6 + b_\infty^6 + \sigma_{(6)}\} \Gamma_{n+1}^{(6)} / \Gamma_{n+1} + V_\infty b_\infty. \end{aligned}$$

The above inequality yields the desirable bound in case (b).

D.3 Proof of Theorem 6

We begin the proof by showing a similar lemma to Lemma 12 as follows:

Lemma 15. *Assume A 1-A 2-A 3-A 4-MA 1. Let $(\eta_k)_{k \in \mathbb{N}^*}$ be a sequence satisfying (14). We have defined*

$$\begin{aligned} \tilde{C}_\hat{e}^{\text{Ret}} &= \{L_\hat{e}\bar{c} + L\hat{e}_\infty\} (e_\infty + b_\infty + 1), \quad C^{\text{Ret}}(\eta_1) = \bar{c}\hat{e}_\infty(2 + \eta_1), \\ D^{\text{Ret}} &= \bar{c}\hat{e}_\infty(a_2 + 1) + L_\hat{e}\bar{c}(a_1(b_\infty + e_\infty) + 1) + L\hat{e}_\infty, \\ P_n &:= D^{\text{Ret}} \sum_{k=0}^n \eta_{k+1}^2 \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] + \tilde{C}_\hat{e}^{\text{Ret}} \Gamma_{n+1}^{(2)} + C^{\text{Ret}}(\eta_1). \end{aligned}$$

Note that $\tilde{C}_\hat{e}^{\text{Ret}} = C_\hat{e}^{\text{Ret}} - 2L(b_\infty^2 + e_\infty^2)$.

(a) Under **R4**, it holds that for any $n \in \mathbb{N}$,

$$\mathbb{E} \left[- \sum_{k=0}^n \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \right] \leq P_n + E_{(1)}^{\text{Ret}} \Gamma_{n+1}^{(3)}, \quad (74)$$

where $E_{(1)}^{\text{Ret}} = 3\hat{e}_\infty \mathcal{L}_\infty^{(1)} (L_\hat{e}h_\infty + L)(e_\infty^2 + b_\infty^2 + h_\infty^2)$.

(b) Under **R5**, it holds that for any $n \in \mathbb{N}$,

$$\mathbb{E} \left[- \sum_{k=0}^n \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \right] \leq P_n + E_{(2)}^{\text{Ret}} \Gamma_{n+1}^{(4)}, \quad (75)$$

where $E_{(2)}^{\text{Ret}} = 9\hat{e}_\infty \mathcal{L}_\infty^{(2)} (L_\hat{e}h_\infty + L)(e_\infty^3 + b_\infty^3 + h_\infty^3)$.

Proof. The proof is almost identical to that of Lemma 12 with the following modifications: first, we observe that (47) is updated to

$$\begin{aligned} \ell(\gamma^{(k)}) &= \eta_k \left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}, \\ &\leq \eta_k \left\{ \|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1})\|_{\theta_k} + \|\Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k} \right\}. \end{aligned} \quad (76)$$

Next we can consider the same decomposition of the left hand side of (48):

$$\mathbb{E} \left[- \sum_{k=0}^n \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \right] = \mathbb{E} \left[- \sum_{i=1}^5 A_i \right],$$

where A_i are defined to be exactly the same as (49). In particular, we have $\mathbb{E}[A_1] = 0$ and

$$\begin{aligned} |A_4| &\leq \bar{c}\hat{e}_\infty \left\{ \eta_1 + a_2 \sum_{k=1}^n \eta_k^2 \|h(\theta_{k-1})\|_{\theta_{k-1}}^2 \right\}, \\ |A_5| &\leq \bar{c}\hat{e}_\infty \left\{ 2 + \eta_1^2 \|h(\theta_0)\|_{\theta_0}^2 + \eta_{n+1}^2 \|h(\theta_n)\|_{\theta_n}^2 \right\}. \end{aligned}$$

Our remaining task is to bound $|A_2|, |A_3|$ using the updated bound (76). Observe that using Cauchy-Schwarz inequality, MA1-(iii)-(iv), (76) and A2-(a), we get

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

Furthermore, using the inequality $a \leq a^2 + 1$ and $|ab| \leq a^2/2 + b^2/2$ gives:

$$|A_2| \leq L_{\hat{e}} \bar{c} \left\{ (b_\infty + e_\infty) \Gamma_{n+1}^{(2)} + (a_1(b_\infty + e_\infty) + 1) \sum_{k=0}^n \eta_{k+1}^2 \|h(\theta_k)\|_{\theta_k}^2 + \sum_{k=1}^n \eta_k^2 h_\infty \|\Delta_{\theta_{k-1}, \eta_k}(X_k)\|_{\theta_{k-1}} \right\}.$$

For *first order retraction*, i.e., under R4, using (14), (70), A4, we have

$$\begin{aligned} |A_2| &\leq L_{\hat{e}} \bar{c} \left\{ (b_\infty + e_\infty) \Gamma_{n+1}^{(2)} + (a_1(b_\infty + e_\infty) + 1) \sum_{k=0}^n \eta_{k+1}^2 \|h(\theta_k)\|_{\theta_k}^2 \right\} \\ &\quad + 3L_{\hat{e}} \bar{c} \mathcal{L}_\infty^{(1)} \sum_{k=1}^n \eta_k^3 h_\infty \left\{ b_\infty^2 + e_\infty^2 + \|h(\theta_{k-1})\|_{\theta_{k-1}}^2 \right\}, \\ &\leq L_{\hat{e}} \bar{c} \left\{ (b_\infty + e_\infty) \Gamma_{n+1}^{(2)} + (a_1(b_\infty + e_\infty) + 1) \sum_{k=0}^n \eta_{k+1}^2 \|h(\theta_k)\|_{\theta_k}^2 + 3\mathcal{L}_\infty^{(1)} h_\infty (b_\infty^2 + e_\infty^2 + h_\infty^2) \Gamma_{n+1}^{(3)} \right\}. \end{aligned}$$

For second order retraction, i.e., under **R5**, similarly we can use (73) to obtain:

$$|A_2| \leq L\hat{e}\bar{c} \left\{ (b_\infty + e_\infty)\Gamma_{n+1}^{(2)} + (a_1(b_\infty + e_\infty) + 1) \sum_{k=0}^n \eta_{k+1}^2 \|h(\theta_k)\|_{\theta_k}^2 + 9\mathcal{L}^{(2)}h_\infty(b_\infty^3 + e_\infty^3 + h_\infty^3)\Gamma_{n+1}^{(4)} \right\}.$$

On the other hand, we have

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

For first order retraction, i.e., under **R4**, again using (70), **A4**, we have

$$|A_3| \leq L\hat{e}_\infty \left\{ (e_\infty + b_\infty + 1)\Gamma_{n+1}^{(2)} + \sum_{k=0}^n \eta_{k+1}^2 \|h(\theta_k)\|_{\theta_k}^2 + 3\mathcal{L}^{(1)}(e_\infty^2 + b_\infty^2 + h_\infty^2)\Gamma_{n+1}^{(3)} \right\}.$$

For second order retraction, i.e., under **R5**, using (73) yields that

$$|A_3| \leq L\hat{e}_\infty \left\{ (e_\infty + b_\infty + 1)\Gamma_{n+1}^{(2)} + \sum_{k=0}^n \eta_{k+1}^2 \|h(\theta_k)\|_{\theta_k}^2 + 9\mathcal{L}^{(2)}(e_\infty^3 + b_\infty^3 + h_\infty^3)\Gamma_{n+1}^{(4)} \right\}.$$

This concludes the proof. \square

Proof of Theorem 6. Following the proof for Theorem 5 but without invoking **MD1**, we obtain the following inequality that is similar to (69):

$$\begin{aligned} \eta_{k+1} \mathbb{E} \|h(\theta_k)\|_{\theta_k}^2 &\leq V(\theta_k) - V(\theta_{k+1}) + \eta_{k+1} \mathbb{E} \left[\langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \middle| \mathcal{F}_k \right] \\ &\quad + \eta_{k+1} \mathbb{E} \left[\langle \text{grad } V(\theta_k), b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \rangle_{\theta_k} \middle| \mathcal{F}_k \right] \\ &\quad + (L\eta_{k+1}^2/2) \mathbb{E} \left[\left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \middle| \mathcal{F}_k \right]. \end{aligned} \tag{77}$$

First order retraction Under **R4**, the terms on right hand side in (77) can be upper bounded one-by-one as:

$$\begin{aligned} &\mathbb{E} \left[\langle \text{grad } V(\theta_k), b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \rangle_{\theta_k} \right] \\ &\quad \leq V_\infty b_\infty + 3\eta_{k+1} V_\infty \mathcal{L}_\infty^{(1)} \{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + e_\infty^2 \}, \\ &\mathbb{E} \left[\left\| H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\|_{\theta_k}^2 \right] \\ &\quad \leq 4 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + e_\infty^2 + 27(\mathcal{L}_\infty^{(1)})^2 (h_\infty^4 + b_\infty^4 + e_\infty^4) \eta_{k+1}^2 \right\}. \end{aligned}$$

As such, summing up the inequality (77) from $k = 0$ to $k = n$ and rearranging terms yield

$$\begin{aligned}
& \sum_{k=0}^n \eta_{k+1} \{ \underline{c} - (3V_\infty \mathcal{L}_\infty^{(1)} + 2L) \eta_{k+1} \} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\
& \leq \mathbb{E} [V(\theta_0) - V(\theta_{n+1})] + \mathbb{E} \left[\sum_{k=0}^n \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \right] + V_\infty b_\infty \Gamma_{n+1} \\
& \quad + \{ 3\mathcal{L}_\infty^{(1)} V_\infty + 2L \} (b_\infty^2 + e_\infty^2) \Gamma_{n+1}^{(2)} + 54L (\mathcal{L}_\infty^{(1)})^2 (h_\infty^4 + b_\infty^4 + e_\infty^4) \Gamma_{n+1}^{(4)}.
\end{aligned} \tag{78}$$

Therefore, combining (78) and (74), we obtain

$$\begin{aligned}
& \sum_{k=0}^n \eta_{k+1} \{ \underline{c} - (3V_\infty \mathcal{L}_\infty^{(1)} + 2L + D^{\text{Ret}}) \eta_{k+1} \} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\
& \leq \mathbb{E} [V(\theta_0) - V(\theta_{n+1})] + C^{\text{Ret}}(\eta_1) + C_{\hat{e}}^{\text{Ret}} \Gamma_{n+1}^{(2)} + E_{(1)}^{\text{Ret}} \Gamma_{n+1}^{(3)} + V_\infty b_\infty \Gamma_{n+1} \\
& \quad + \{ 3\mathcal{L}_\infty^{(1)} V_\infty + 2L \} (b_\infty^2 + e_\infty^2) \Gamma_{n+1}^{(2)} + 54L (\mathcal{L}_\infty^{(1)})^2 (h_\infty^4 + b_\infty^4 + e_\infty^4) \Gamma_{n+1}^{(4)}.
\end{aligned}$$

Since $\eta_{k+1} \leq \underline{c} / (2(3V_\infty \mathcal{L}_\infty^{(1)} + 2L + D^{\text{Ret}}))$, using $V(\theta) \geq 0$ gives

$$\begin{aligned}
& (\underline{c}/2) \sum_{k=0}^n \eta_{k+1} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\
& \leq \mathbb{E} [V(\theta_0)] + C^{\text{Ret}}(\eta_1) + V_\infty b_\infty \Gamma_{n+1} + \left\{ C_{\hat{e}}^{\text{Ret}} + (3\mathcal{L}_\infty^{(1)} V_\infty + 2L) (b_\infty^2 + e_\infty^2) \right\} \Gamma_{n+1}^{(2)} \\
& \quad + E_{(1)}^{\text{Ret}} \Gamma_{n+1}^{(3)} + 54L (\mathcal{L}_\infty^{(1)})^2 (h_\infty^4 + b_\infty^4 + e_\infty^4) \Gamma_{n+1}^{(4)}.
\end{aligned}$$

Finally, we have

$$\begin{aligned}
& \mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \\
& \leq 2V_\infty b_\infty / \underline{c} + 2(\underline{c} \Gamma_{n+1})^{-1} \left\{ \mathbb{E} [V(\theta_0)] + C^{\text{Ret}}(\eta_1) + \left\{ C_{\hat{e}}^{\text{Ret}} + (3\mathcal{L}_\infty^{(1)} V_\infty + 2L) (b_\infty^2 + e_\infty^2) \right\} \Gamma_{n+1}^{(2)} \right\} \\
& \quad + 2(\underline{c} \Gamma_{n+1})^{-1} \left\{ E_{(1)}^{\text{Ret}} \Gamma_{n+1}^{(3)} + 54L (\mathcal{L}_\infty^{(1)})^2 (h_\infty^4 + b_\infty^4 + e_\infty^4) \Gamma_{n+1}^{(4)} \right\}.
\end{aligned}$$

Collecting terms and computing the constants yield (15).

Second order retraction Under **R5**, the terms on right hand side in (77) can be upper bounded one-by-one as:

$$\begin{aligned}
& \mathbb{E} \left[\left\langle \text{grad } V(\theta_k), b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1}) \right\rangle_{\theta_k} \right] \leq V_\infty b_\infty + 9\eta_{k+1}^2 V_\infty \mathcal{L}_\infty^{(2)} \{ h_\infty^3 + b_\infty^3 + e_\infty^3 \}, \\
& \mathbb{E} \left[\|H_{\theta_k}(X_{k+1}) + b_{\theta_k}(X_{k+1}) + \Delta_{\theta_k, \eta_{k+1}}(X_{k+1})\|_{\theta_k}^2 \right] \\
& \quad \leq 4 \left\{ \|h(\theta_k)\|_{\theta_k}^2 + b_\infty^2 + e_\infty^2 + 243 (\mathcal{L}_\infty^{(2)})^2 (h_\infty^6 + b_\infty^6 + e_\infty^6) \eta_{k+1}^4 \right\}.
\end{aligned}$$

As such, summing up the inequality (77) from $k = 0$ to $k = n$ and rearranging terms yield

$$\begin{aligned} & \sum_{k=0}^n \eta_{k+1} \{\underline{c} - 2L\eta_{k+1}\} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\ & \leq \mathbb{E} [V(\theta_0) - V(\theta_{n+1})] + \mathbb{E} \left[\sum_{k=0}^n \eta_{k+1} \langle \text{grad } V(\theta_k), e_{\theta_k}(X_{k+1}) \rangle_{\theta_k} \right] + V_\infty b_\infty \Gamma_{n+1} \\ & \quad + 2L(b_\infty^2 + e_\infty^2) \Gamma_{n+1}^{(2)} + 9\mathcal{L}_\infty^{(2)} V_\infty (h_\infty^3 + b_\infty^3 + e_\infty^3) \Gamma_{n+1}^{(3)} + 486L(\mathcal{L}_\infty^{(2)})^2 (h_\infty^6 + b_\infty^6 + e_\infty^6) \Gamma_{n+1}^{(6)}. \end{aligned}$$

Combining with (75) leads to

$$\begin{aligned} & \sum_{k=0}^n \eta_{k+1} \{\underline{c} - (2L + D^{\text{Ret}})\eta_{k+1}\} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\ & \leq \mathbb{E} [V(\theta_0) - V(\theta_{n+1})] + C^{\text{Ret}}(\eta_1) + C_{\hat{e}}^{\text{Ret}} \Gamma_{n+1}^{(2)} + E_{(2)}^{\text{Ret}} \Gamma_{n+1}^{(3)} + V_\infty b_\infty \Gamma_{n+1} \\ & \quad + 2L(b_\infty^2 + e_\infty^2) \Gamma_{n+1}^{(2)} + 9\mathcal{L}_\infty^{(2)} V_\infty (h_\infty^3 + b_\infty^3 + e_\infty^3) \Gamma_{n+1}^{(3)} + 486L(\mathcal{L}_\infty^{(2)})^2 (h_\infty^6 + b_\infty^6 + e_\infty^6) \Gamma_{n+1}^{(6)}. \end{aligned}$$

Since we have set $(2L + D^{\text{Ret}})\eta_{k+1} \leq \underline{c}/2$, using $V(\theta) \geq 0$ gives

$$\begin{aligned} & (\underline{c}/2) \sum_{k=0}^n \eta_{k+1} \mathbb{E} \left[\|h(\theta_k)\|_{\theta_k}^2 \right] \\ & \leq \mathbb{E} [V(\theta_0)] + C^{\text{Ret}}(\eta_1) + (C_{\hat{e}}^{\text{Ret}} + 2L(b_\infty^2 + e_\infty^2)) \Gamma_{n+1}^{(2)} + V_\infty b_\infty \Gamma_{n+1} \\ & \quad + \{E_{(2)}^{\text{Ret}} + 9\mathcal{L}_\infty^{(2)} V_\infty (h_\infty^3 + b_\infty^3 + e_\infty^3)\} \Gamma_{n+1}^{(3)} + 486L(\mathcal{L}_\infty^{(2)})^2 (h_\infty^6 + b_\infty^6 + e_\infty^6) \Gamma_{n+1}^{(6)}. \end{aligned}$$

Finally, it leads to

$$\begin{aligned} & \mathbb{E} \left[\|h(\theta_{I_n})\|_{\theta_{I_n}}^2 \right] \\ & \leq 2V_\infty b_\infty / \underline{c} + 2(\underline{c} \Gamma_{n+1})^{-1} \left\{ \mathbb{E} [V(\theta_0)] + C^{\text{Ret}}(\eta_1) + (C_{\hat{e}}^{\text{Ret}} + 2L(b_\infty^2 + e_\infty^2)) \Gamma_{n+1}^{(2)} \right\} \\ & \quad + 2(\underline{c} \Gamma_{n+1})^{-1} \left\{ \{E_{(2)}^{\text{Ret}} + 9\mathcal{L}_\infty^{(2)} V_\infty (h_\infty^3 + b_\infty^3 + e_\infty^3)\} \Gamma_{n+1}^{(3)} + 486L(\mathcal{L}_\infty^{(2)})^2 (h_\infty^6 + b_\infty^6 + e_\infty^6) \Gamma_{n+1}^{(6)} \right\}. \end{aligned}$$

Collecting terms and identifying the constants lead to (16). \square

D.4 Proof of Proposition 7

First, for any $\theta \in \mathbb{S}^d$, the Riemannian exponential map is given for any $u \in T_\theta \mathbb{S}^d$ by (see [18, Proposition 5.27 and its proof])

$$\text{Exp}_\theta(u) = \cos(\|u\|) \theta + \sin(\|u\|) (u / \|u\|). \quad (79)$$

In addition, the retraction Ret given by (18) can be written as

$$\text{Ret}_\theta(u) = \text{Exp}_\theta \left(\arctan(\|u\|) \frac{u}{\|u\|} \right) \quad (80)$$

This can be proven by replacing the identities,

$$\cos(\arctan(\|u\|)) = \frac{1}{\sqrt{1 + \|u\|^2}} \quad \sin(\arctan(\|u\|)) = \frac{\|u\|}{\sqrt{1 + \|u\|^2}}$$

into (79). Indeed, this yields,

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

To see that this is equal to $\text{Ret}_\theta(u)$, note that $1 + \|u\|^2 = \|\theta + u\|^2$, because $\|\theta\| = 1$ and u is orthogonal to θ (since $u \in T_\theta\Theta$). Then, (80) follows from (18).

The following are now proven.

- Condition **R1** is satisfied: this condition is just the definition of a retraction, as given in [1].
- Condition **R2** is satisfied: the cut locus of a point θ on the sphere S^d is $\text{Cut}(\theta) = \{-\theta\}$ [18] (Page 308). The Riemannian (that is, spherical) distance between θ and $-\theta$ is $\rho_\Theta(\theta, -\theta) = \pi$. On the other hand, from (80), $\rho_\Theta(\theta, \text{Ret}_\theta(u)) < \frac{\pi}{2}$ because $\arctan(\|u\|) < \pi/2$ for all $u \in T_\theta\Theta$. It is then clear that $\text{Ret}_\theta(u) \neq \{-\theta\}$ for any $u \in T_\theta\Theta$.
- Condition **R3** is satisfied: the isometry group of $\Theta = S^d$ is $G = O(d)$, the group of $d \times d$ orthogonal matrices. The action of G on Θ is given by matrix-vector multiplication, $g \cdot \theta = g\theta$ and $g \cdot u = gu$. From (18),

$$g \cdot \text{Ret}_\theta(u) = \frac{g \cdot (\theta + u)}{\|g \cdot (\theta + u)\|} \quad (81)$$

However, since g is an orthogonal matrix, g preserves Euclidean norms, so $\|g \cdot (\theta + u)\| = \|\theta + u\|$. Replacing into (81),

$$g \cdot \text{Ret}_\theta(u) = \frac{g \cdot (\theta + u)}{\|g \cdot (\theta + u)\|} = \frac{g \cdot \theta + g \cdot u}{\|g \cdot \theta + g \cdot u\|} \quad (82)$$

where the second equality follows since the action of g is linear. Finally, the right-hand side of (82) is $\text{Ret}_{g \cdot \theta}(g \cdot u)$.

- Condition **R4** is satisfied: from (80) and **R2**,

$$\Phi_\theta(u) = \arctan(\|u\|) \frac{u}{\|u\|} \quad (83)$$

The required second derivative can now be computed, thanks to the identity,

$$D^2\Phi_\theta(tu)[u, u] = \frac{d^2}{dt^2}\Phi_\theta(tu) \quad (84)$$

Indeed, using (83) and (84),

$$D^2\Phi_\theta(tu)[u, u] = \frac{d^2}{dt^2} \arctan(t\|u\|) \frac{u}{\|u\|} = \|u\|^2 \left(f_2(t\|u\|) \frac{u}{\|u\|} \right)$$

where f_2 is the second derivative of the arctan function, so $|f_2(x)| \leq 1$ for real x . Now, since $\Theta = S^d$, here $\|u\|_\theta = \|u\|$. Thus, Condition **R4** is satisfied with $\mathcal{L}^{(1)}(\theta) = 1$.

- Condition **R5-(i)** is satisfied: recall (59) from the proof of Lemma 13. This states,

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

From (84), it then follows,

$$D_t\dot{\gamma}(0) = \left. \frac{d^2}{dt^2} \right|_{t=0} \Phi_\theta(tu)$$

Since Φ_θ is given by (83), an elementary calculation shows the right-hand side is here equal to zero.

- Condition **R5-(ii)** is satisfied: the proof is similar to the above one for **R4**. Here, instead of (84), it is enough to use

$$D^3\Phi_\theta(tu)[u, u, u] = \frac{d^3}{dt^3} \Phi_\theta(tu)$$

Using (83), this shows that **R5-(i)** is satisfied with $\mathcal{L}^{(2)}(\theta) = 2$.

- Φ_θ has bounded first derivative: from (83), by differentiating,

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

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

D.5 Proof of Proposition 8

First, by [?, Equation 2.32], the Riemannian exponential map at θ is given for $u \in T_\theta\text{Gr}_r(\mathbb{R}^d)$, with $u = B_\perp C$, $C \in \mathbb{R}^{(d-r) \times r}$:

$$\text{Exp}_\theta(u) = \left[(B, B_\perp) \exp \left(\begin{pmatrix} 0 & -C^\top \\ C & 0 \end{pmatrix} \right) \begin{pmatrix} I_r \\ 0_{d-r \times r} \end{pmatrix} \right] \quad (85)$$

where \exp is the matrix exponential. In addition, we show below that the retraction Ret defined by (21) can be written on the form

$$\text{Ret}_\theta(u) = \text{Exp}_\theta(\Phi_\theta(u)) \quad \Phi_\theta(u) = B_\perp V \arctan(a) U^\top \quad (86)$$

for $u \in T_\theta\Theta$ with $u = B_\perp C$ as in (19), where C has singular value decomposition $C = V a U^\top$. Here, V is $(d-r) \times (d-r)$ orthogonal and U is $r \times r$ orthogonal. Moreover, $\arctan(a)$ is obtained by taking the arctangent of each element of the matrix a . Accepting (86), it is possible to show that.

- Condition **R1** is satisfied: this condition is just the definition of a retraction, as given in [1].
- Condition **R2** is satisfied: when $\Theta = \text{Gr}_r(\mathbb{R}^d)$, the cut locus of each $\theta \in \Theta$ is given by [?],

$$\text{Cut}(\theta) = \left\{ \text{Exp}_\theta(B_\perp C) \mid C = V a U^\top; \|a\|_\infty = \frac{\pi}{2} \right\}$$

where $\|a\|_\infty = \max_{ij} |a_{ij}|$. From (86), for any $u \in T_\theta\Theta$, one has

$$\text{Ret}_\theta(u) = \text{Exp}_\theta(B_\perp C(u)) \text{ where } C(u) = V \arctan(a) U^\top$$

Since $\|\arctan(a)\|_\infty < \pi/2$, it follows that $\text{Ret}_\theta(u) \notin \text{Cut}(\theta)$.

- Condition **R3** is satisfied: the isometry group of $\Theta = \text{Gr}_r(\mathbb{R}^d)$ is $G = O(d)$, the group of $d \times d$ orthogonal matrices. The action of G on Θ is given by $g \cdot \theta = g(\theta)$ (the image of the subspace θ of \mathbb{R}^d by the orthogonal transformation g). If $u \in T_\theta\Theta$ is given by (19), then $g \cdot u = gu$ is a matrix product.

Note that, if $\theta = [B]$ for some $B \in \text{St}_r(\mathbb{R}^d)$, then $g \cdot \theta = [gB]$. Applying this property in (21),

$$g \cdot \text{Ret}_\theta(u) = g([B + u]) = [gB + gu] \quad (87)$$

But, since $g \cdot \theta = [gB]$ and $gu = g \cdot u$, (87) implies

$$g \cdot \text{Ret}_\theta(u) = \text{Ret}_{g \cdot \theta}(g \cdot u)$$

which is (13).

- Condition **R4** is satisfied: the required second derivative is computed using the identity (this is a repetition of (84)),

$$D^2\Phi_\theta(tu)[u, u] = \frac{d^2}{dt^2}\Phi_\theta(tu) \quad (88)$$

Using (86) and (88),

$$D^2\Phi_\theta(tu)[u, u] = B_\perp V (a \odot a \odot f_2(ta)) U^\top \quad (89)$$

where \odot denotes the Kronecker product, and f_2 is the second derivative of the arctan function (again, this is applied to each element of the matrix (ta)). From [?] (Page 314)

$$\left\| D^2\Phi_\theta(tu)[u, u] \right\|_\theta^2 = \text{tr} \left((a \odot a \odot f_2(ta))(a \odot a \odot f_2(ta))^\top \right)$$

where tr denotes the trace. Then, using the fact that $|f_2(x)| \leq 1$ for real x , the right-hand side is less than $\text{tr}(aa^\top)$, which is equal to $\|u\|_\theta^2$. Thus, Condition **R4** is satisfied with $\mathcal{L}^{(1)}(\theta) = 1$.

- Condition **R5-(i)** is satisfied: recall (59) from the proof of Lemma 13. This states,

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

Setting $t = 0$ in (89), it then follows

$$D_t \dot{\gamma}(0) = B_\perp V (a \odot a \odot f_2(0)) U^\top$$

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

- Condition **R5-(ii)** is satisfied: the proof is similar to the above one for **R4**. Here, instead of (88), it is enough to use

$$D^3 \Phi_\theta(tu)[u, u, u] = \frac{d^3}{dt^3} \Phi_\theta(tu)$$

by computing the derivative, as in (89), it can be shown that **R5** is satisfied with $\mathcal{L}^{(2)}(\theta) = 2$.

- Φ_θ does not have bounded first derivative: assume $r > 1$. Recall that $\Phi_\theta(u)$ is given by (86), which can be written

$$\Phi_\theta(u) = B_\perp \varphi(\psi_\theta(u)) \quad (90)$$

where, $\psi_\theta : \mathbb{R}^{d \times r} \rightarrow \mathbb{R}^{(d-r) \times r}$ and $\varphi : \mathbb{R}^{(d-r) \times r} \rightarrow \mathbb{R}^{(d-r) \times r}$ are given by

$$\psi_\theta(u) = B_\perp^\top u \quad \varphi(C) = V \arctan(a) U^\top \quad (91)$$

whenever C has singular value decomposition $C = VaU^\top$. Indeed, if $u = B_\perp C$ as in (19), then $\psi_\theta(u) = C$, so (90) is equivalent to (86). From (90) and (91), by an application of the chain rule

$$D\Phi_\theta(u)[v] = B_\perp D\varphi(C)[B_\perp^\top v]$$

for $v \in T_\theta \Theta$, where $C = \psi_\theta(u)$. Now, to show that $D\Phi_\theta(u)$ is not bounded, it is enough to show that $D\varphi(C)$ is not bounded. However,

$$D\varphi(C)[w] = DV[w] (\arctan(a)) U^\top + VD (\arctan(a))[w] U^\top + V (\arctan(a)) DU[w]$$

for $w \in \mathbb{R}^{(d-r) \times r}$, where DV , $D(\arctan(a))$ and DU denote the derivatives of V , $\arctan(a)$ and U , as functions of C , by an abuse of notation. To simplify the proof, assume, without loss of generality, that C is a square matrix (for example, if $d - r \geq r$, it is enough to add zero columns to C). With this assumption, the following formulae hold [?],

$$DV[w] = V \left[F \odot \left(V^\top w U a + a U^\top w^\top V \right) \right]$$

$$\begin{aligned} D(\arctan(a))[w] &= \mathbf{I}_{(d-r)} \left[V^\top w U \right] \\ DU[w] &= U \left[F \odot \left(a V^\top w U + U^\top w^\top V a \right) \right] \end{aligned}$$

where F is the matrix with entries $F_{ij} = (a_j^2 - a_i^2)^{-1}$ for $i \neq j$ and $F_{ii} = 0$. However, taking $w = V\omega a U^\top$ where ω is a $(d-r) \times (d-r)$ antisymmetric matrix, yields $D(\arctan(a))[w] = 0$ and $DU[w] = 0$, while

$$DV[w] = V [G \odot \omega]$$

where G has matrix elements $G_{ij} = a_i^2 a_j^2 / (a_j^2 - a_i^2)$ for $i \neq j$ and $G_{ii} = 0$. Clearly, these do not remain bounded as $a_i - a_j \rightarrow 0$.

Proof of (86) here, let $s = d - r$ and assume, without any loss of generality, that $s \geq r$.

Recall that, in (86), $u = B_\perp C$ where C has singular value decomposition $C = V a U^\top$. Here, V and U are orthogonal, and $a = (\alpha, 0_{r \times s-r})^\top$, with $r \times r$ diagonal matrix α . Write $\Phi_\theta(u)$ under the form

$$\Phi_\theta(u) = B_\perp C(u) \text{ where } C(u) = V \arctan(a) U^\top \quad (92)$$

Using (85), it follows

$$\text{Exp}_\theta(\Phi_\theta(u)) = \left[Q \exp \begin{pmatrix} 0 & -C(u)^\top \\ C(u) & 0 \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ 0_{s \times r} \end{pmatrix} \right] \quad (93)$$

where $Q = (B, B_\perp)$. The aim is to show this is equal to $\text{Ret}_\theta(u)$, given by (21). Using the expression of $C(u)$ in (92), and performing the matrix multiplication, it is possible to check that

$$\begin{pmatrix} 0 & -C(u)^\top \\ C(u) & 0 \end{pmatrix} = \begin{pmatrix} U & \\ & V \end{pmatrix} \begin{pmatrix} 0 & -\arctan(a)^\top \\ \arctan(a) & 0 \end{pmatrix} \begin{pmatrix} U^\top & \\ & V^\top \end{pmatrix} \quad (94)$$

Recall $\exp(AXA^{-1}) = A \exp(X) A^{-1}$ for any square matrices A and X , where A is invertible. It follows from (94),

$$\exp \begin{pmatrix} 0 & -C(u)^\top \\ C(u) & 0 \end{pmatrix} = \begin{pmatrix} U & \\ & V \end{pmatrix} \exp \begin{pmatrix} 0 & -\arctan(a)^\top \\ \arctan(a) & 0 \end{pmatrix} \begin{pmatrix} U^\top & \\ & V^\top \end{pmatrix}$$

By plugging this into (93), and noticing that

$$\left[\begin{pmatrix} U^\top & \\ & V^\top \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ 0_{s \times r} \end{pmatrix} \right] = \left[\begin{pmatrix} \mathbf{I}_r \\ 0_{s \times r} \end{pmatrix} \right]$$

it follows

$$\text{Exp}_\theta(\Phi_\theta(u)) = \left[Q \begin{pmatrix} U & \\ & V \end{pmatrix} \exp \begin{pmatrix} 0 & -\arctan(a)^\top \\ \arctan(a) & 0 \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ 0_{s \times r} \end{pmatrix} \right] \quad (95)$$

If $f = (\phi, 0_{r \times s-r})^\top$ where ϕ is $r \times r$ diagonal, then, under the assumption that $s \geq r$, (this is proven in detail, at the end of the present proof),

$$\exp \begin{pmatrix} 0 & -f^\top \\ f & 0 \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ 0_{s \times r} \end{pmatrix} = \begin{pmatrix} C(\phi) \\ S(f) \end{pmatrix} \quad (96)$$

where $C(\phi) = \cos(\phi)$ and $S(f) = \sin(f)$, with the functions \cos and \sin applied to each matrix element of ϕ and f , respectively. The identity (96) can be used to evaluate the matrix exponential in (95), since $a = (\alpha, 0_{r \times s-r})^\top$. This yields,

$$\text{Exp}_\theta(\Phi_\theta(u)) = \left[Q \begin{pmatrix} U & \\ & V \end{pmatrix} \begin{pmatrix} \cos(\arctan(\alpha)) \\ \sin(\arctan(a)) \end{pmatrix} \right]$$

However, since $\cos(\arctan(x)) = 1/(1+x^2)^{1/2}$ and $\sin(\arctan(x)) = x/(1+x^2)^{1/2}$, this becomes

$$\text{Exp}_\theta(\Phi_\theta(u)) = \left[Q \begin{pmatrix} U & \\ & V \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ a \end{pmatrix} (\mathbf{I}_r + \alpha)^{-1/2} \right] = \left[Q \begin{pmatrix} U & \\ & V \end{pmatrix} \begin{pmatrix} \mathbf{I}_r \\ a \end{pmatrix} \right] \quad (97)$$

where the second equality holds because $(\mathbf{I}_r + \alpha)$ is invertible (the diagonal elements of α are the singular values of C , and are therefore positive). It follows from (97) that

$$\text{Exp}_\theta(\Phi_\theta(u)) = \left[Q \begin{pmatrix} U \\ Va \end{pmatrix} \right] = \left[Q \begin{pmatrix} U \\ Va \end{pmatrix} U^\top \right] \quad (98)$$

where the second equality holds because U^\top is an invertible $r \times r$ matrix (which therefore does not change the span of the columns of the overall matrix product). Performing the matrix product in (98), and noting $UU^\top = \mathbf{I}_r$ and $C = VaU^\top$, it finally follows that

$$\text{Exp}_\theta(\Phi_\theta(u)) = \left[Q \begin{pmatrix} \mathbf{I}_r \\ C \end{pmatrix} \right] = \text{Exp}_\theta(\Phi_\theta(u)) = \left[(B, B_\perp) \begin{pmatrix} \mathbf{I}_r \\ C \end{pmatrix} \right]$$

From $u = B_\perp C$, this immediately implies

$$\text{Exp}_\theta(\Phi_\theta(u)) = [B + u]$$

which means $\text{Exp}_\theta(\Phi_\theta(u)) = \text{Ret}_\theta(u)$, as required in (86).

Proof of (96) this follows from

$$\exp \begin{pmatrix} 0 & -f^\top \\ f & 0 \end{pmatrix} = \begin{pmatrix} C(\phi) & -S(f)^\top \\ S(f) & C(\phi) \end{pmatrix} \quad (99)$$

where $C(\phi)$ and $S(f)$ are as in (96) and where (recall it is assumed that $s \geq r$),

$$C(\phi) = \begin{pmatrix} \cos(\phi) & 0_{r \times s-r} \\ 0_{s-r \times r} & 0_{s-r \times s-r} \end{pmatrix}$$

To prove (99), write

$$\begin{pmatrix} 0 & -f^\top \\ f & 0 \end{pmatrix} = \sum_{i=1}^r \phi_i \mathbf{b}_{r+i} \quad (100)$$

where $f = (\phi, 0_{r \times s-r})^\top$ with diagonal ϕ , and where

$$\mathbf{b}_{r+i} = e_{r+i,i} - e_{i,r+i}$$

with $e_{j,k}$ a matrix all of whose elements are zero, except the one at row j and column k , which is equal to 1. One easily checks the matrices \mathbf{b}_{r+i} commute with each other. Therefore, (100) implies

$$\exp \begin{pmatrix} 0 & -f^\top \\ f & 0 \end{pmatrix} = \prod_{i=1}^r \exp(\phi_i \mathbf{b}_{r+i}) \quad (101)$$

However, it is elementary that

$$\exp(\phi_i \mathbf{b}_{r+i}) = \mathbf{I}_d + (\cos(\phi_i) - 1) \mathbf{a}_{r+i} + \sin(\phi_i) \mathbf{b}_{r+i} \quad (102)$$

where \mathbf{I}_d is the $d \times d$ identity matrix and

$$\mathbf{a}_{r+i} = e_{i,i} + e_{r+i,r+i}$$

Finally, (99) follows from (101) and (102), after noting the matrix products, for $i \neq k$,

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

which can be checked immediately.

E Proofs of Section 5.3

We present lemmas to verify **A2–A3**, **MD1** for the robust barycenter problem. In addition, we show that the function $V(\theta)$ of (26) 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 (28) finds a stationary point of (26), the strict g-convexity of $V(\theta)$ guarantees that such stationary point is globally optimal and is unique.

Lemma 16. *Let Θ be a Hadamard manifold with sectional curvature bounded below by $-\kappa^2$, $x \in \Theta$ and $\delta > 0$. Define for any $\theta \in \Theta$,*

$$V_2(\theta) = \rho_{\Theta}^2(x, \theta) \quad \text{and} \quad V_1(\theta) = \delta^2 \left[V_2(\theta)/\delta^2 + 1 \right]^{1/2} - \delta^2 .$$

Then, V_1 is a smooth function and its Riemannian gradient is given for any $\theta \in \Theta$ by

$$\text{grad } V_1(\theta) = - \text{Exp}_{\theta}^{-1}(x) \Big/ \left[V_2(\theta)/\delta^2 + 1 \right]^{1/2} . \quad (103)$$

Moreover, for any $\theta \in \Theta$, $v \in T_{\theta}\Theta \setminus \{0\}$, its Hessian satisfies

$$0 < \text{Hess } V_1(\theta)(v, v) \leq (1 + \delta\kappa) \|v\|_{\theta}^2 .$$

Proof. The proof relies heavily on the computation of $\text{grad } V_2$ and the Hessian comparison of V_2 done in [? , Theorem 5.6.1]. Indeed, [? , Theorem 5.6.1] shows V_2 is smooth and that for any $\theta \in \Theta$,

$$\text{grad } V_2(\theta) = -2\text{Exp}_x^{-1}(\theta) .$$

Hence, (103) 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 $v \in T_v\Theta$, 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 Appendix A.8. The product rule for the covariant derivative [? , 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

$$\begin{aligned} \text{Hess } V_1(\theta) = & - \text{Exp}_{\theta}^{-1}(x) \otimes \text{Exp}_{\theta}^{-1}(x) \Big/ \left[\delta^2 \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{3/2} \right] \\ & + \text{Hess } V_2(\theta) \Big/ \left[2 \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{1/2} \right] . \end{aligned} \quad (104)$$

Let $\theta \in \Theta$ and $v \in T_{\theta}\Theta \setminus \{0\}$. 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,

$$0 \leq \text{Exp}_{\theta}^{-1}(x) \otimes \text{Exp}_{\theta}^{-1}(x)(v, v) \leq \left\| \text{Exp}_{\theta}^{-1}(x) \right\|_{\theta}^2 \|v\|_{\theta}^2 = \rho_{\Theta}^2(\theta, x) \|v\|_{\theta}^2 , \quad (105)$$

since $\|\text{Exp}_{\theta}^{-1}(x)\|_{\theta} = \rho_{\Theta}(\theta, x)$. On the other hand, [? , Theorem 5.6.1] implies that

$$2 \|v\|_{\theta}^2 \leq \text{Hess } V_2(\theta)(v, v) \leq 2\kappa\rho_{\Theta}(\theta, x) \coth[\kappa\rho_{\Theta}(\theta, x)] \|v\|_{\theta}^2 . \quad (106)$$

Now, combining (105)-(106) in (104), and using $t \coth(t) \leq 1 + t$ for $t \geq 0$, it follows that

$$\begin{aligned} \text{Hess } V_1(\theta)(v, v) &\leq [1 + \kappa \rho_\Theta(\theta, x)] \left/ \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{1/2} \right. \|v\|_\theta^2 \\ &\leq (1 + \kappa\delta) \|v\|_\theta^2, \end{aligned}$$

where we have used $1 + V_2(\theta)/\delta^2 \geq \max(1, \rho_\Theta^2(\theta, x)/\delta^2)$. Similarly, we obtain

$$\begin{aligned} \text{Hess } V_1(\theta)(v, v) &\geq - \left(\rho_\Theta^2(\theta, x) \|v\|_\theta^2 \right) \left/ \left(\delta^2 \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{3/2} \right) \right. + \|v\|_\theta^2 \left/ \left\{ 1 + V_2(\theta)/\delta^2 \right\}^{1/2} \right. \\ &> 0, \end{aligned}$$

which concludes the proof. \square

Lemma 17. *Let Θ be a Hadamard manifold with sectional curvature bounded below by $-\kappa^2$. Furthermore, consider the Lyapunov function given by (26) and stochastic approximation scheme (28). Then, A 2 and A 3 are satisfied, with $\underline{c} = \bar{c} = 1$, $L = 1 + \delta\kappa$, and $b_\infty = 0$. Moreover, MD 1 is satisfied, with $\sigma_0 = \delta^2$ and $\sigma_1 = 0$.*

Proof. To show that A 2 holds with the stated values of \underline{c} , \bar{c} and L , note that the scheme (28) can be written

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

where the stochastic update $H_{\theta_n}(X_{n+1})$ is given by

$$H_\theta(x) = -\text{Exp}_\theta^{-1}(x) \left/ \left[1 + \{\rho_\Theta(\theta, x)/\delta\}^2 \right]^{1/2} \right. . \quad (107)$$

Then, both $\text{grad } V$ and $\text{Hess } V$ can be computed differentiating under the integral in (26). Using Lemma 16, 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 $\|\text{Exp}_\theta^{-1}(x)\|_\theta = \rho_\Theta(\theta, x)$, we have that for any $x, \theta \in \Theta$, $\|\text{grad } V_1(\theta)\|_\theta \leq 1$ and, for any $x, \theta \in \Theta, v \in \mathbb{T}_\theta\Theta$,

$$|\text{Hess } V_1(\theta)(v, v)| \leq (1 + \kappa\delta) \|v\|_\theta^2.$$

Under these domination conditions, using Lebesgue's dominated convergence theorem, we have for any $\theta \in \Theta$,

$$\text{grad } V(\theta) = - \int_\Theta \text{Exp}_\theta^{-1}(x) \left/ \left[1 + \{\rho_\Theta(\theta, x)/\delta\}^2 \right]^{1/2} \right. \pi(\text{d}x),$$

and for any $v \in \mathbb{T}_\theta\Theta$,

$$\text{Hess } V(\theta)(v, v) = \int_\Theta \text{Hess } V_1(\theta)(v, v) \pi(\text{d}x).$$

Thus, for any $v \in \mathbb{T}_\theta \Theta \setminus \{0\}$,

$$0 < \text{Hess } V(\theta)(v, v) \leq (1 + \kappa\delta) \|v\|_\theta^2 . \quad (108)$$

This last inequality proves that the operator norm of $\text{Hess } V$ is upper bounded by $1 + \delta\kappa$. Therefore, using Lemma 10, it follows that **A2-(b)** holds with $L = 1 + \delta\kappa$.

It remains to prove that Assumptions **MD1** is satisfied with $\sigma_0 = \delta^2$ and $\sigma_1 = 0$. To do so, note first from (107), that

$$\|H_{\theta_n}(X_{n+1})\|_{\theta_n}^2 = \left\| \text{Exp}_{\theta_n}^{-1}(X_{n+1}) \right\|_{\theta_n}^2 / \left(1 + \{\rho_\Theta(\theta, X_{n+1})/\delta\}^2 \right) .$$

Since $\|\text{Exp}_{\theta_n}^{-1}(X_{n+1})\|_{\theta_n}^2 = \rho_\Theta^2(\theta_n, X_{n+1})$, it follows that $\|H_{\theta_n}(X_{n+1})\|_{\theta_n}^2 \leq \delta^2$. Thus, in the notation of (5),

$$\mathbb{E} \left[\|e_{\theta_n}(X_{n+1})\|_{\theta_n}^2 \mid \mathcal{F}_n \right] = \mathbb{E} \left[\|H_{\theta_n}(X_{n+1}) - h(\theta_n)\|_{\theta_n}^2 \mid \mathcal{F}_n \right] \leq \delta^2 \quad (109)$$

since the conditional variance is bounded by the mean square. Now, the required values of σ_0 and σ_1 can be read from (109). \square

Proposition 18. *Let Θ be a Hadamard manifold with sectional curvature bounded below, and let π be a probability distribution on Θ . Assume there exists some $\tau \in \Theta$ such that*

$$\int_\Theta \rho_\Theta(\tau, x) \pi(dx) < +\infty . \quad (110)$$

Then, the function $V : \Theta \rightarrow \mathbb{R}_+$ given by (26) is geodesically strictly convex, but not strongly convex, in general. Moreover, V has a unique global minimizer $\theta^ \in \Theta$. In other words, π has a unique robust barycenter θ^* .*

Proof. Under condition (110), the function V takes finite values, $V(\theta) < +\infty$ for any $\theta \in \Theta$. 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 (27) and (26), this inequality implies

$$V(\theta) \leq \delta \int_\Theta \rho_\Theta(\theta, x) \pi(dx) \quad \text{for } \theta \in \Theta$$

and, furthermore, by the triangle inequality,

$$V(\theta) \leq \delta \int_\Theta (\rho_\Theta(\theta, \tau) + \rho_\Theta(\tau, x)) \pi(dx) = \rho_\Theta(\theta, \tau) + \int_\Theta \rho_\Theta(\tau, x) \pi(dx) \quad (111)$$

Then, it follows from (110) and (111) that $V(\theta) < +\infty$ for $\theta \in \Theta$.

Further, from (108) in the proof of Lemma 17, Hess $V(\theta) \succ 0$, so V has strictly positive-definite Riemannian Hessian, and is therefore geodesically strictly convex. To see that $V(\theta)$ may fail to be strongly convex, consider the case where $\pi = \delta_\tau$ (here, δ_τ is the Dirac distribution, concentrated at $\tau \in \Theta$). By (26), it then follows

$$V(\theta) = \tilde{\rho}(\theta, \tau) \quad (112)$$

Then, let $\gamma(t)$ be a geodesic through τ , given by $\gamma(t) = \text{Exp}_\tau(tu)$ for $t \in \mathbb{R}$, where $u \in T_\tau\Theta$ has $\|u\|_\tau = 1$. If $V(\theta)$ is given by (112), then it follows from (27)

$$(V \circ \gamma)(t) = \delta^2 \left[1 + \{t/\delta\}^2 \right]^{1/2} - \delta^2$$

but this is not a strongly convex function of $t \in \mathbb{R}$. Therefore, V is not geodesically strongly convex on Θ [?, p. 187].

It remains to show that V has a unique global minimizer θ^* . Since V is geodesically strictly convex, and bounded below (indeed, V is positive), it is enough to show that V is coercive, in the sense that $V(\theta) \rightarrow +\infty$ when $\rho_\Theta(\theta, \tau) \rightarrow +\infty$. To do so, note the following inequality holds for all real x ,

$$\delta^2 \left[1 + \{x/\delta\}^2 \right]^{1/2} - \delta^2 \geq \frac{\delta}{\sqrt{2}}(x - \delta)$$

Then, using (27) and (26), this inequality implies

$$V(\theta) \geq \frac{\delta}{\sqrt{2}} \int_\Theta (\rho_\Theta(\theta, x) - \delta) \pi(dx)$$

or, by the triangle inequality,

$$V(\theta) \geq \frac{\delta}{\sqrt{2}} \int_\Theta (|\rho_\Theta(\theta, \tau) - \rho_\Theta(\tau, x)| - \delta) \pi(dx) \quad \text{for } \theta \in \Theta$$

However, this directly yields

$$V(\theta) \geq \frac{\delta}{\sqrt{2}} \rho_\Theta(\theta, \tau) - \frac{\delta}{\sqrt{2}} \int_\Theta (\rho_\Theta(\tau, x) + \delta) \pi(dx)$$

Clearly, the right-hand side increases to $+\infty$ when $\rho_\Theta(\theta, \tau) \rightarrow +\infty$. Thus, V is indeed coercive.

Finally, to show that V has a unique global minimizer θ^* , let $V_* = \inf\{V(\theta); \theta \in \Theta\}$ and $(\theta_n; n = 1, 2, \dots)$ a sequence of points in Θ such that $\lim V(\theta_n) = V_*$. since V takes finite values, $V_* < +\infty$. Therefore, there exists some $R > 0$ such that $\rho_\Theta(\theta_n, \tau) < R$ for all n . This is because, otherwise, $\rho_\Theta(\theta_n, \tau) \rightarrow +\infty$, and thus $V(\theta_n) \rightarrow +\infty$, since V is coercive. Because Θ is a complete Riemannian manifold, the metric ball $B(\tau, R)$ has compact closure (a consequence of the Hopf-Rinow theorem [18]). 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 V , it is clear that $V(\theta^*) = \lim V(\theta_n) = V_*$, so θ^* is indeed a global minimizer of V . This global minimizer is unique because V is geodesically strictly convex. \square