



HAL
open science

Bayesian imaging using Plug & Play priors: when Langevin meets Tweedie

Rémi Laumont, Valentin de Bortoli, Andrés Almansa, Julie Delon, Alain Durmus, Marcelo Pereyra

► **To cite this version:**

Rémi Laumont, Valentin de Bortoli, Andrés Almansa, Julie Delon, Alain Durmus, et al.. Bayesian imaging using Plug & Play priors: when Langevin meets Tweedie. 2021. hal-03161400v1

HAL Id: hal-03161400

<https://hal.science/hal-03161400v1>

Preprint submitted on 6 Mar 2021 (v1), last revised 19 Jan 2022 (v3)

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.

Bayesian imaging using Plug & Play priors: when Langevin meets Tweedie *

Rémi Laumont ^{†‡} Valentin De Bortoli ^{†§} Andrés Almansa [‡]
Julie Delon [‡] Alain Durmus [¶] Marcelo Pereyra ^{||}

March 6, 2021

Abstract

Since the seminal work of Venkatakrishnan et al. [83] in 2013, *Plug & Play* (PnP) methods have become ubiquitous in Bayesian imaging. These methods derive Minimum Mean Square Error (MMSE) or Maximum A Posteriori (MAP) estimators for inverse problems in imaging by combining an explicit likelihood function with a prior that is implicitly defined by an image denoising algorithm. The PnP algorithms proposed in the literature mainly differ in the iterative schemes they use for optimisation or for sampling. In the case of optimisation schemes, some recent works guarantee the convergence to a fixed point, albeit not necessarily a MAP estimate. In the case of sampling schemes, to the best of our knowledge, there is no known proof of convergence. There also remain important open questions regarding whether the underlying Bayesian models and estimators are well defined, well-posed, and have the basic regularity properties required to support these numerical schemes. To address these limitations, this paper develops theory, methods, and provably convergent algorithms for performing Bayesian inference with PnP priors. We introduce two algorithms: 1) PnP-ULA (Plug & Play Unadjusted Langevin Algorithm) for Monte Carlo sampling and MMSE inference; and 2) PnP-SGD (Plug & Play Stochastic Gradient Descent) for MAP inference. Using recent results on the quantitative convergence of Markov chains, we establish detailed convergence guarantees for these two algorithms under realistic assumptions on the denoising operators used, with special attention to denoisers based on deep neural networks. We also show that these algorithms approximately target a decision-theoretically optimal Bayesian model that is well-posed. The proposed algorithms are demonstrated on several canonical problems such as image deblurring, inpainting, and denoising, where they are used for point estimation as well as for uncertainty visualisation and quantification.

1 Introduction

1.1 Bayesian inference in imaging inverse problems

Most inverse problems in imaging aim at reconstructing an unknown image $x \in \mathbb{R}^d$ from a degraded observation $y \in \mathcal{Y}$ (a given measurable vector space) under some assumptions on their mathematical relationship. For example, many works consider observation models of the form

*VDB was partially supported by EPSRC grant EP/R034710/1. RL was partially supported by grants from Région Ile-De-France. AD acknowledges support of the Lagrange Mathematical and Computing Research Center. MP was partially supported by EPSRC grant EP/T007346/1. Computer experiments for this work ran on a Titan Xp GPU donated by NVIDIA, as well as on HPC resources from GENCI-IDRIS (Grant 2020-AD011011641).

$y = \mathbf{A}(x) + n$, where $\mathbf{A} : \mathbb{R}^d \rightarrow \mathbf{Y}$ is a degradation operator modelling deterministic instrumental aspects of the observation process, and n is an unknown (stochastic) noise term taking values in \mathbf{Y} . The operator \mathbf{A} can be known or not, and is usually assumed to be linear (e.g., \mathbf{A} can represent blur, missing pixels, a projection, etc.).

Most of the time, the estimation of x from y is ill-posed or ill-conditioned¹ and additional assumptions on the unknown x are required in order to deliver meaningful estimates. The Bayesian statistical paradigm provides a natural framework to regularise such estimation problems so as to deliver accurate and well-posed solutions. Accordingly, the relationship between x and y is described by a statistical model with likelihood function $p(y|x)$, and the knowledge about x is encoded by the marginal or *prior* distribution for x , typically specified via a density function $p(x)$ or by its potential $U(x) = -\log p(x)$. Similarly, in some cases the likelihood $p(y|x)$ is specified via the potential $F(x, y) = -\log p(y|x)$. Unless explicitly stated otherwise, we henceforth assume that $\mathbf{Y} = \mathbb{R}^m$ for some $m \in \mathbb{N}$ and that all densities are defined w.r.t. to the appropriate Lebesgue measure.

The likelihood and prior define the joint distribution with density $p(x, y) = p(y|x)p(x)$, from which we derive the *posterior* distribution with density $p(x|y)$ where for any $x \in \mathbb{R}^d, y \in \mathbb{R}^m$

$$p(x|y) = p(y|x)p(x) / \int_{\mathbb{R}^d} p(y|\tilde{x})p(\tilde{x})d\tilde{x} , \quad (1)$$

which underpins all inference about x given the observation y . Most imaging methods derive Minimum Mean Square Error (MMSE) or Maximum A Posteriori (MAP) estimators

$$\hat{x}_{\text{MAP}} = \arg \max_{x \in \mathbb{R}^d} p(x|y) = \arg \min_{x \in \mathbb{R}^d} \{F(x, y) + U(x)\} , \quad (2)$$

$$\hat{x}_{\text{MMSE}} = \arg \min_{u \in \mathbb{R}^d} \mathbb{E}[\|x - u\|^2|y] = \mathbb{E}[x|y] = \int_{\mathbb{R}^d} \tilde{x}p(\tilde{x}|y)d\tilde{x} . \quad (3)$$

Bayesian models are conceptually tractable, modular, and involve terms with clearly defined roles. These terms can be simple to ease computations, or highly sophisticated in order to accurately describe the observation processes as well as the available prior knowledge [73, 4]. Bayesian methods can naturally address problems involving partially unknown models (e.g., blind, semi-blind and unsupervised problems), and provide a range of decision-theoretic strategies to calibrate models, diagnose model misspecification, and compare alternative models. Bayesian decision theory also allows quantifying the uncertainty in the delivered solutions in order to inform decision-making in a rigorous and principled manner [69, 10].

The quality of the inference about x given y depends on how accurately the specified prior represents the true marginal distribution for x . Most works in the Bayesian imaging literature consider relatively simple priors promoting sparsity in transformed domains or piece-wise regularity (e.g., involving the ℓ_1 norm or the total-variation pseudo-norm [73, 25, 59, 65]), Markov random fields [1], or learning-based priors like patch-based Gaussian or Gaussian mixture models [96, 92, 2, 82, 48]. Special attention is given in the literature to models that have specific factorisation structures or that are log-concave, as this enables the use of Bayesian computation algorithms that scale efficiently to high-dimensions and which have detailed convergence guarantees. For example, proximal splitting convex optimisation algorithms compute \hat{x}_{MAP} and basic forms of uncertainty quantification [69], whereas proximal Markov chain Monte Carlo stochastic algorithms draw samples from $p(x|y)$ in order to compute \hat{x}_{MMSE} and perform uncertainty quantification analyses as well as other advanced inferences such as automatic model calibration and model misspecification diagnosis [65, 38].

¹That is, either the estimation problem does not admit a unique solution, or there exists a unique solution but it is not Lipschitz continuous w.r.t. to perturbations in the data y .

1.2 Bayesian computation in imaging inverse problems

There is a vast literature on Bayesian computation methodology for models related to imaging sciences (see, e.g., [66]). Here, we briefly summarise efficient high-dimensional Bayesian computation strategies derived from the Langevin stochastic differential equation (SDE)

$$\begin{aligned} d\mathbf{X}_t &= \nabla \log p(\mathbf{X}_t|y) + \sqrt{2}d\mathbf{B}_t \\ &= \nabla \log p(y|\mathbf{X}_t) + \nabla \log p(\mathbf{X}_t) + \sqrt{2}d\mathbf{B}_t, \end{aligned} \quad (4)$$

where $(\mathbf{B}_t)_{t \geq 0}$ is a d -dimensional Brownian motion. When $p(x|y)$ is proper and smooth, with $x \mapsto \nabla \log p(x|y)$ Lipschitz continuous², then, for any initial condition $\mathbf{X}_0 \in \mathbb{R}^d$, the SDE (4) has a unique strong solution $(\mathbf{X}_t)_{t \geq 0}$ that admits the posterior of interest $p(x|y)$ as unique stationary density [71]. In addition, for any initial condition $\mathbf{X}_0 \in \mathbb{R}^d$ the distribution of \mathbf{X}_t converges towards the posterior distribution in total variation. Although solving (4) in continuous time is generally not possible, we can use discrete time approximations of (4) to generate samples that are approximately distributed according to $p(x|y)$. A natural choice is the Unadjusted Langevin algorithm (ULA) Markov chain $(X_k)_{k \geq 0}$ obtained from an Euler-Maruyama discretisation of (4), given by $X_0 \in \mathbb{R}^d$ and the following recursion for all $k \in \mathbb{N}$

$$X_{k+1} = X_k + \delta_k \nabla \log p(y|X_k) + \delta_k \nabla \log p(X_k) + \sqrt{2\delta_k} Z_{k+1}, \quad (5)$$

where $\{Z_k : k \in \mathbb{N}\}$ is a family of i.i.d Gaussian random variables with zero mean and identity covariance matrix and $(\delta_k)_{k \in \mathbb{N}}$ is a sequence of positive step-sizes which controls a trade-off between asymptotic accuracy and convergence speed [30, 37]. The approximation error involved in discretising (4) can be asymptotically removed at the expense of additional computation by combining (5) with a Metropolis-Hastings correction step, leading to the so-called Metropolis-adjusted Langevin Algorithm (MALA) [71]. There are approximations of (4) with better convergence properties than ULA (see, e.g., [67]), as well as other dynamics to construct efficient Monte Carlo algorithms [43, 28]. Nevertheless, ULA remains a highly useful Bayesian computation methodology that is straightforward to apply, robust to high dimensionality, and thoroughly theoretically understood.

Similarly, when $p(x|y)$ is proper and differentiable, with $x \mapsto \nabla \log p(x|y)$ Lipschitz continuous, it is possible to use first-order optimisation methods to compute maximisers of $p(x|y)$, *i.e.* MAP estimators. A natural choice in this context is the gradient algorithm given by $X_0 \in \mathbb{R}^d$ and the following recursion for all $k \in \mathbb{N}$

$$X_{k+1} = X_k + \delta_k \nabla \log p(y|X_k) + \delta_k \nabla \log p(X_k),$$

which converges to critical points of $p(x|y)$ under mild assumptions on the sequence of step-sizes $(\delta_k)_{k \in \mathbb{N}}$ [64]. Alternatively, the stochastic gradient descent (SGD) variant

$$X_{k+1} = X_k + \delta_k \nabla \log p(y|X_k) + \delta_k \nabla \log p(X_k) + \delta_k Z_{k+1},$$

where $\{Z_k : k \in \mathbb{N}\}$ is a family of i.i.d Gaussian random variables with zero mean and identity covariance matrix, is more robust to local minima and saddle points and hence more suitable when $x \mapsto p(x|y)$ is not log-concave on \mathbb{R}^d [21, 18]. Again, there are other optimisation schemes related to (4) with better convergence properties than SGD (see, e.g., [64] in the convex case), as well as other dynamics to construct efficient optimisation algorithms [53, 97]. Nevertheless, similarly to ULA, SGD is straightforward to apply, robust, and has a detailed convergence theory, making it a valuable algorithm in the imaging scientist's toolbox.

²That is, there exists $L \geq 0$ such that for any $x_1, x_2 \in \mathbb{R}^d$, $\|\nabla \log p(x_1|y) - \nabla \log p(x_2|y)\| \leq L\|x_1 - x_2\|$

1.3 Machine learning and Plug & Play approaches in imaging inverse problems

In an apparently different direction, machine learning approaches have recently gained a considerable importance in the field of imaging inverse problems, particularly strategies based on deep neural networks. Indeed, neural networks can be trained as regressors to learn the function $y \mapsto \hat{x}_{\text{MMSE}}$ empirically from a huge dataset of examples $\{x'_i, y'_i\}_{i=1}^N$, where $N \in \mathbb{N}$ is the size of the training dataset. Many recent works on the topic report unprecedented accuracy. This training can be agnostic [35, 93, 95, 41, 75, 40] or exploit the knowledge of \mathbf{A} in the network architecture via unrolled optimization techniques [45, 27, 34, 42]. However, solutions encoded by end-to-end neural networks are mostly problem specific and not easily adapted to reflect changes in the problem (e.g., in instrumental settings). There also exists concerns regarding the stability of such approaches for general reconstruction problem [6, 5].

A natural strategy to reconcile the strengths of the Bayesian paradigm and neural networks is provided by *Plug & Play* approaches. These data-driven regularisation approaches learn an implicit representation of the prior density $p(x)$ (or its potential $U(x) = -\log p(x)$) while keeping an explicit likelihood density, which is usually assumed to be known and calibrated. More precisely, using a denoising algorithm D_ε , most often implemented by a previously trained neural network, *Plug & Play* approaches seek to derive an approximation of the gradient ∇U (called the score) [11, 12] or of the proximal operator prox_U [60, 94, 26, 50, 74], which is then used within an iterative minimisation scheme to approximate \hat{x}_{MAP} . While these approaches have shown some remarkable empirical performance, they rely on hybrid algorithms that are poorly understood and that in some cases fail to converge. Indeed, their convergence properties remain an important open question, especially when D_ε is implemented as a neural network. This question has been recently studied in the literature but only partly answered [74, 90, 79]. Similar *Plug & Play* iterative schemes have been proposed to approximately draw samples from $p(x|y)$ in order to approximate \hat{x}_{MMSE} [3, 47, 49], although again with no theoretical guarantees on the delivered solutions.

1.4 Contributions summary

One of the aims of this paper is to gain significant insights into Bayesian inference with *Plug & Play* priors by formalising them in the Bayesian statistical framework and exploring some open questions on the topic. In addition to aspects regarding the convergence of these schemes, we raise crucial inference questions. For example, it is not clear what constraints should be imposed on D_ε so that it leads to a Bayesian model whose posterior $p_\varepsilon(x|y)$ is proper and admits a density w.r.t. to an appropriate Lebesgue measure (this is essential for the definition of \hat{x}_{MAP} and to study its properties). Similarly, we wonder under what conditions on D_ε the posterior $p_\varepsilon(x|y)$ is well-posed (i.e., Lipschitz continuous w.r.t y for a suitable metric) [55]. We also raise questions about the existence and the accurate estimation of \hat{x}_{MMSE} and other moments of $x|y$. There also remain questions regarding the frequentist properties of Bayesian models with *Plug & Play* priors, as well as about basic properties of the marginal likelihood $p(y) = \int p(y|\tilde{x})p(\tilde{x})d\tilde{x}$, which plays a central role in model calibration and model selection tasks [84]. Lastly, a mature literature on *Plug & Play* Bayesian methodology should also address the diagnoses of model misspecification directly from y . This arises, for example, when D_ε is learnt from training data that is not representative of x , leading to a posterior $p(x|y)$ with poor accuracy properties.

In this paper, we provide a novel and unified framework of theory, methods, and algorithms for Bayesian analysis and computation with models involving *Plug & Play* priors. We establish detailed convergence guarantees based on [16] for the following Bayesian computation algorithms:

- PnP-SGD, for approximately computing \hat{x}_{MAP} using SGD;
- PnP-ULA, for approximately sampling from the posterior $p(x|y)$ using ULA.

These two algorithms share a similar form and satisfy the following recursion: for any $k \in \mathbb{N}$

$$X_{k+1} = X_k - \delta_k \nabla F(X_k, y) - (\delta_k / \varepsilon)(X_k - D_\varepsilon(X_k)) + \begin{cases} \delta_k Z_{k+1} & \text{(PnP-SGD)} \\ \sqrt{2\delta_k} Z_{k+1} & \text{(PnP-ULA)} \end{cases},$$

where $\{Z_k : k \in \mathbb{N}\}$ is a family of independent Gaussian random variables with zero mean and identity covariance matrix, $D_\varepsilon : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a denoiser operator and $(\delta_k)_{k \in \mathbb{N}}$ is a sequence of step-sizes. Convergence guarantees are provided for both algorithms, under clear assumptions on the denoiser D_ε that are realistic for several well-known denoisers, in particular some neural network-based denoisers. Our results show that these algorithms approximately target a decision-theoretically optimal Bayesian model that is well defined, well-posed, and sufficiently regular to allow efficient Bayesian computation by optimisation and Langevin sampling. We explicitly characterise the estimation bias introduced by using tractable denoising operators (as opposed to the intractable oracle Bayesian model), as well as the bias arising from the use of discrete approximations and a finite number of iterations. In addition, we provide experimental implementations of PnP-ULA and PnP-SGD for a specific neural network denoiser [74] shown to satisfy our convergence guarantees. These are demonstrated on several canonical inverse problems such as denoising, deblurring or inpainting, and compared with the state-of-the-art method PnP-ADMM [74].

The paper is organized as follows. In Section 2, we summarise previous works on *Plug & Play* approaches for MAP estimation and MMSE estimation via posterior sampling. Section 3 then presents our framework for studying Bayesian inference methods with *Plug & Play* priors as well as the algorithms PnP-ULA and PnP-SGD. Section 4 demonstrates the proposed methodology on several canonical imaging inverse problems. This is then followed by a detailed theoretical analysis and convergence proofs in Section 5. In Section 6, we conclude and discuss perspectives for future work.

2 A survey of Plug & Play methods in imaging

In the context of imaging inverse problems, *Plug & Play* methods aim at using a carefully chosen denoiser $D_\varepsilon : \mathbb{R}^d \rightarrow \mathbb{R}^d$ to implicitly define an image prior. This is achieved by relating D_ε to a gradient or proximal operator associated with the prior density. In what follows, we describe how these approaches have been widely used to derive MAP estimators (Section 2.1), or less frequently for posterior sampling and computation of MMSE estimators (Section 2.2). In our discussion, we pay particular attention to questions related to algorithmic convergence, as this has been the main focus of the literature. However, as mentioned previously, there remain many other fundamental open questions. For example, questions related to the correct definition of the Bayesian models involved, the existence of the estimators under consideration and the impact of model misspecification in *Plug & Play* models and their diagnosis.

2.1 *Plug & Play* MAP estimators

Plug & Play approaches seeking to derive a MAP estimator for the problem (2) can be divided in two families, depending on whether D_ε is itself a MAP estimator or an MMSE estimator³ for

³Notice that although it is conceptually helpful to distinguish these two cases (in order to make a historical and practical survey of the subject), there are clear theoretical connections between the two approaches. Indeed,

an image denoising problem.

Plug & Play MAP estimators using proximal splitting If the denoiser D_ε is a MAP estimator, there exists a prior potential U such that for each image $x \in \mathbb{R}^d$, $D_\varepsilon(x) = \arg \min_{\tilde{x} \in \mathbb{R}^d} \{\frac{1}{2}\|x - \tilde{x}\|^2 + \varepsilon U(\tilde{x})\} = \text{prox}_{\varepsilon U}(x)$. Any optimization scheme making use of a proximal descent on the prior can then be used to solve (2) via D_ε . For instance, the half-quadratic splitting (HQS) algorithm solves a series of simpler problems

$$(\hat{x}_{\text{MAP},\varepsilon}, \hat{z}_{\text{MAP},\varepsilon}) = \arg \min_{x,z \in \mathbb{R}^d} E_\varepsilon(x,z), \quad \text{where } E_\varepsilon(x,z) = F(x,y) + \|x-z\|^2/(2\varepsilon) + U(z),$$

and each subproblem is solved by an alternate minimization scheme,

$$x_{k+1} = \arg \min_x E_\varepsilon(x, z_k) = \text{prox}_{\varepsilon F(\cdot, y)}(z_k), \quad (6)$$

$$z_{k+1} = \arg \min_z E_\varepsilon(x_{k+1}, z) = \text{prox}_{\varepsilon U}(x_{k+1}) = D_\varepsilon(x_{k+1}). \quad (7)$$

When U is convex, such splitting schemes and many variants (including HQS, primal-dual methods, FBS or ISTA, ADMM, etc.) are well understood and proved to converge to the global optimum [20]. They have also been successfully used for non-convex U like patch-based Gaussian mixture models (GMM) as pioneered for external learning by Zoran & Weiss in [96]. The use of splitting schemes with non-convex GMM priors was later refined with convergence guarantees for scene-adapted learning [81].

Following the seminal work [83], this kind of splitting schemes have become popular in cases where U is unknown, but a denoiser D_ε is available and assumed to be a good approximation of $\text{prox}_{\varepsilon U}$. As popular and useful these methods have become, their convergence properties are largely unknown: it remains difficult to show that they converge to the MAP estimator, or even a critical point of (2). Indeed, for most denoisers D_ε , there is no guarantee that there exists a potential U such that $D_\varepsilon = \text{prox}_{\varepsilon U}$. In [77], Sreehari *et al.* establish some sufficient conditions for this to happen: D_ε must be differentiable, and its Jacobian J_{D_ε} should be symmetric with eigenvalues within the $[0, 1]$ interval (which ensures non expansiveness). These assumptions hold for transform-domain thresholding denoisers and for variants of Non Local means [22] where symmetry is explicitly enforced [77]. However, the two assumptions are unfortunately false for most popular denoisers, including Non Local Means [22], BM3D [29], Non Local Bayes [56] and neural networks based denoisers like DnCNN [93], as observed in [68].

More recently Ryu *et al.* [74] proposed a convergence proof of the *Plug & Play* Alternating Direction Method of Multipliers (ADMM) and Forward Backward Splitting (FBS) schemes, based on an alternative assumption on the denoiser D_ε , trained to ensure that its residual operator $D_\varepsilon - \text{Id}$ is L-Lipschitz, with a Lipschitz constant L which depends both on the data fitting term and the denoiser. With this assumption they show convergence to a fixed point, but since there is still no guarantee that the denoiser D_ε is the proximal operator of a potential U , we cannot be sure whether this fixed point is a MAP estimator for some prior or not. In addition the proof requires F to be strongly convex (which excludes all cases where \mathbf{A} is not full rank and de facto excludes some of the applications considered in [74]) and it imposes quite restrictive assumptions on the regularization parameter. More recently, Sun et al [79] reformulated the *Plug & Play* ADMM algorithm with different convergence conditions, notably without the assumption that F is strongly convex, but with more restrictive conditions on the denoiser D_ε ⁴. As mentioned

under regularity conditions on the Bayesian model involved, MAP denoisers can be expressed as MMSE denoisers under an alternative (albeit often unknown) Bayesian model [46]. However this equivalence can not always be exploited in practice and has been mostly ignored in the literature on *Plug & Play* methods until very recently with the work of Xu *et al.* [90] to be presented later.

⁴In [79], the residual $\text{Id} - D_\varepsilon$ is assumed to be firmly non expansive, which is equivalent to say that D_ε is firmly non expansive, see [8, Proposition 4.4].

previously, the convergence of the algorithms is only the first of many important questions that need to be addressed.

Plug & Play MAP estimators using gradient descent Second, assume that D_ε is a MMSE estimator, *i.e.* $D_\varepsilon(x_\varepsilon) = \mathbb{E}[x|x_\varepsilon]$, where $x \sim p$ and $x_\varepsilon \sim \mathcal{N}(x, \varepsilon \text{Id})$ conditionally to x . In this case, Tweedie’s identity [39] establishes the following link between D_ε and $p_\varepsilon = p * G_\varepsilon$, a smoothed version of the prior by a Gaussian kernel G_ε with variance ε , for any $x \in \mathbb{R}^d$

$$\nabla U_\varepsilon(x) = -\nabla \log p_\varepsilon(x) = (x - D_\varepsilon(x))/\varepsilon, \quad (8)$$

where $U_\varepsilon = -\log(p_\varepsilon)$. This relation can be used in any gradient descent scheme involving ∇U_ε and it is at the core of the two algorithms PnP-SGD and PnP-ULA presented in this paper.

A similar relation is derived by Romano *et al.* in [72] where they present the Regularization by Denoising (RED) method. Instead of using Tweedie’s identity, the RED method solves equation (2) via gradient descent with explicit regularization $U_\varepsilon(x) = (1/2)\langle x, x - D_\varepsilon(x) \rangle$. As shown in [68], under the assumptions that D_ε is locally homogeneous and has symmetric Jacobian, this implies that for any $x \in \mathbb{R}^d$, $\nabla U_\varepsilon(x) = x - D_\varepsilon(x)$, which is (up to a scaling factor $1/\varepsilon$) the same expression as Tweedie’s identity in equation (8). Unfortunately, as pointed out before, these assumptions on D_ε are not satisfied by most commonly used denoisers, and the convergence of the RED algorithms for these denoisers remains unproven.

Following on from this, Xu *et al.* [90] very recently proposed a convergence study for a *Plug & Play* ISTA (Iterated Soft Thresholding Algorithm) scheme, which alternates proximal descents on F with gradient descents on U_ε . They derive convergence guarantees under very mild conditions on both the denoiser and the data fitting term. However, Xu *et al.* [90] assumes that D_ε is an exact MMSE denoiser, so their theoretical results do not carry to denoisers learned from training data and implemented by neural networks.

The PnP-SGD optimisation algorithm that will be presented in this paper is close to the ones presented in [72], and is shown to converge under much milder conditions than previously assumed, and in particular when D_ε is not an exact MAP or MMSE denoiser. Importantly, our convergence proof is valid for the neural network denoiser used in [74] (a variant of DnCNN [93] with a contractive residual) and for the native Non Local Means [23].

2.2 Plug & Play posterior sampling and MMSE estimators

The idea of using *Plug & Play* priors to draw samples from the posterior distribution $p(x|y)$ has also been explored in the imaging literature [47, 49, 52]. Both of these approaches start from a Langevin algorithm to draw samples from the posterior distribution. Drawing on several ideas developed in [3], the authors of [47] use a Taylor expansion of D_ε (chosen as a denoising autoencoder) and approximate $\nabla U(x)$ by $(x - D_\varepsilon(x))/\varepsilon$ for small values of ε in the Langevin scheme, while the more recent [49] uses Tweedie’s identity more directly (8) and also within a Langevin scheme. The sampling algorithms studied therein are illustrated with several linear inverse problems, but without any convergence or accuracy guarantees.

The availability of clear convergence guarantees for sampling algorithms is arguably even more important than for optimisation schemes, as the former can suffer from a range of convergence issues that are not always easy to detect. For example, Langevin algorithms can admit the correct target measure but fail to converge with exponential rate [71] or deliver solutions with an unacceptably large bias. Clear convergence results are also essential for setting algorithm parameters, as well as for using the algorithms within more complex numerical schemes (e.g. in the context of hierarchical or empirical Bayesian models [84, 31]).

To the best of our knowledge, the PnP-ULA scheme presented in this paper is the first provably convergent Monte Carlo sampling algorithm for implicit *Plug & Play* priors in imaging. This algorithm and the associated convergence results set the grounds for more advanced sampling methods with superior convergence properties (e.g., [67]).

3 Bayesian inference with Plug & Play priors: theory methods and algorithms

3.1 Bayesian modelling and analysis with *Plug & Play* priors

In this section we formalise a Bayesian framework for *Plug & Play* priors and analyse some of their main properties. As explained previously, we are interested in the estimation of the unknown image x from an observation y when the problem is ill-conditioned or ill-posed, resulting in significant uncertainty about the value of x . Adopting a Bayesian approach, we seek to use prior knowledge about x to reduce the uncertainty and render the estimation problem well-posed. This prior knowledge is specified by the marginal distribution for x , which is combined with the likelihood function to derive the posterior distribution for $x|y$.

In this paper, we analyse *Plug & Play* Bayesian models for $x|y$ through the prism of *M-complete* Bayesian modelling [9]. Accordingly, there exists a true -albeit unknown and intractable- marginal distribution for x and posterior distribution for $x|y$. Basing inferences on these true marginal and posterior distributions is theoretically optimal both in terms of point estimation and in term of delivering Bayesian probabilities that are valid from a frequentist viewpoint. We henceforth use μ to denote this optimal prior distribution for x on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$, and when μ admits a density w.r.t. the Lebesgue measure on \mathbb{R}^d , we denote this density by p^* . In the latter case, the posterior distribution for $x|y$ associated with the marginal μ also admits a density that is given for any $x \in \mathbb{R}^d$ and $y \in \mathbb{R}^m$ by

$$p^*(x|y) = p(y|x)p^*(x) / \int_{\mathbb{R}^d} p(y|\tilde{x})p^*(\tilde{x})d\tilde{x}. \quad (9)$$

Unlike most Bayesian imaging approaches that operate implicitly in an *M-closed* manner and treat their postulated Bayesian models as true models (see [9] for more details), we regard p^* (or more precisely μ) as a fundamental property of the unknown x , and other models as operational approximations specified by the practitioner (either analytically, algorithmically, or from data). This distinction will be useful for using the oracle (9) as a reference for tractable Bayesian models and algorithms.

In this conceptual construction, the marginal μ naturally depends on the imaging application considered and could be the distribution of natural images of the size and resolution of x , or that of a class of images related to a specific application. Also, in problems where there is training data $\{x'_i\}_{i=1}^N$ available, we regard $\{x'_i\}_{i=1}^N$ as samples from μ . Lastly, we note that the posterior for $x|y$ remains well defined when μ does not admit a density; this is important to provide robustness to situations where p^* is nearly degenerate or improper. For clarity, our presentation assumes that p^* exists, although this is not strictly required ⁵.

Notice that because μ is unknown, $p^*(x|y)$ might not verify the basic desiderata for efficient Bayesian computation: i.e., $p^*(x|y)$ need not be proper and differentiable, with $\nabla \log p^*(x|y)$ Lipschitz continuous. In order to guarantee that gradient based algorithms such as ULA and SGD are well posed, we introduce a regularised approximation μ_ε obtained via the convolution of μ with a Gaussian smoothing kernel with bandwidth $\varepsilon > 0$. By construction, μ_ε has a smooth

⁵Operating without densities requires measure disintegration concepts that are technical [76].

proper density p_ε given for any $x \in \mathbb{R}^d$ and $\varepsilon > 0$ by

$$p_\varepsilon^*(x) = (2\pi\varepsilon)^{-d/2} \int_{\mathbb{R}^d} \exp[-\|x - \tilde{x}\|_2^2/(2\varepsilon)] p^*(\tilde{x}) d\tilde{x} .$$

Henceforth, we also assume that the likelihood satisfies the following standard conditions:

H1. For any $y \in \mathbb{R}^m$, $\sup_{x \in \mathbb{R}^d} p(y|x) < +\infty$, $p(y|\cdot) \in C^1(\mathbb{R}^d, (0, +\infty))$ and there exists $L_y > 0$ such that $\nabla \log(p(y|\cdot))$ is L_y Lipschitz continuous.

Equipped with the approximation p_ε^* and **H1**, we use Bayes' theorem to involve the likelihood $p(y|x)$ and derive the posterior density $p_\varepsilon^*(x|y)$, given for any $\varepsilon > 0$ and $x \in \mathbb{R}^d$ by

$$p_\varepsilon^*(x|y) = p(y|x)p_\varepsilon^*(x) / \int_{\mathbb{R}^d} p(y|\tilde{x})p_\varepsilon^*(\tilde{x})d\tilde{x} . \quad (10)$$

This regularised oracle model will enable defining well-posed gradient-based algorithms, which will serve as basis or reference for *Plug & Play* based ULA and SGD algorithms.

Proposition 1 below establishes that the regularised prior $p_\varepsilon^*(x)$ and posterior $p_\varepsilon^*(x|y)$ are proper and smooth, and that they can be made arbitrarily close to the unregularised oracle models $p^*(x)$ and $p^*(x|y)$ by reducing ε , with the approximation error vanishing as $\varepsilon \rightarrow 0$. For any $f : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable such that $\int_{\mathbb{R}^d} |f(\tilde{x})| d\tilde{x} < +\infty$ we denote $\|f\|_1 = \int_{\mathbb{R}^d} |f(\tilde{x})| d\tilde{x}$.

Proposition 1. Assume **H1**. Then, for any $\varepsilon > 0$ and $y \in \mathbb{R}^m$, the following hold:

- (a) p_ε^* and $p_\varepsilon^*(\cdot|y)$ are proper.
- (b) For any $k \in \mathbb{N}$, $p_\varepsilon^* \in C^k(\mathbb{R}^d)$. In addition, if $p(y|\cdot) \in C^k(\mathbb{R}^d)$ then $p_\varepsilon^*(\cdot|y) \in C^k(\mathbb{R}^d, \mathbb{R})$.
- (c) Let $k \in \mathbb{N}$. If $\int_{\mathbb{R}^d} \|\tilde{x}\|^k p^*(x) d\tilde{x} < +\infty$ then $\int_{\mathbb{R}^d} \|\tilde{x}\|^k p_\varepsilon^*(\tilde{x}|y) d\tilde{x} < +\infty$.
- (d) $\lim_{\varepsilon \rightarrow 0} \|p_\varepsilon^*(\cdot|y) - p^*(\cdot|y)\|_1 = 0$.
- (e) In addition, if there exist $\kappa, \beta \geq 0$ such that for any $x \in \mathbb{R}^d$, $\|p^* - p^*(\cdot - x)\|_1 \leq \|x\|^\beta$, then there exists $C \geq 0$ such that $\|p_\varepsilon^*(\cdot|y) - p^*(\cdot|y)\|_1 \leq C\varepsilon^{\beta/2}$.

Proof. The proof of Proposition 1-(a)-(b)-(c) are straightforward using the dominated convergence theorem. Finally, Proposition 1-(d)-(e) is an application of Proposition 20. \square

With the smoothness conditions required for gradient algorithms in mind, we note that under **H1** and $p(y|\cdot) \in C^1(\mathbb{R}^d)$, $x \mapsto \nabla \log p_\varepsilon^*(x|y)$ is well-defined and continuous. However, $\nabla \log p_\varepsilon^*(x|y)$ might not be Lipschitz continuous, as this requires an additional regularity condition on μ . To analyse this important Lipschitz condition, as well as to set the grounds for *Plug & Play* methods that define priors implicitly through a denoising algorithms, we introduce the oracle MMSE denoiser D_ε^* defined for any $x \in \mathbb{R}^d$ and $\varepsilon > 0$ by

$$D_\varepsilon^*(x) = (2\pi\varepsilon)^{-d/2} \int_{\mathbb{R}^d} \tilde{x} \exp[-\|x - \tilde{x}\|_2^2/(2\varepsilon)] p^*(\tilde{x}) d\tilde{x} .$$

Under the assumption that the expected mean square error (MSE) is finite, D_ε^* is the MMSE estimator to recover an image $x \sim \mu$ from a noisy observation $x_\varepsilon \sim \mathcal{N}(x, \varepsilon \text{Id})$ [70]. From Tweedie's identity [39], the gradient $x \mapsto \nabla \log p_\varepsilon^*(x)$ is related to D_ε^* for any $x \in \mathbb{R}^d$ by

$$\varepsilon \nabla \log p_\varepsilon^*(x) = D_\varepsilon^*(x) - x ,$$

and hence $x \mapsto \nabla \log p_\varepsilon^*(x|y)$ is Lipschitz continuous if and only if D_ε^* has this property. We argue that this is a natural assumption on D_ε^* , as it is essentially equivalent to assuming that

the denoising problem underpinning D_ε^\star is well-posed (recall that an inverse problem is said to be well posed if its solution is unique and Lipschitz continuous w.r.t to the observation [78]).

Proposition 2 states that the Lipschitz continuity of $x \mapsto \nabla \log p_\varepsilon^\star(x)$ is equivalent to assuming that the expected MSE involved in using D_ε^\star to estimate x from $x_\varepsilon \sim \mathcal{N}(x, \varepsilon \text{Id})$, where x has marginal μ , is finite and uniformly upper bounded for all $x_\varepsilon \in \mathbb{R}^d$.

Proposition 2. *Assume H1. Let $\varepsilon > 0$. $\nabla \log p_\varepsilon^\star$ is Lipschitz continuous if and only if there exists $C \geq 0$ such that for any $x_\varepsilon \in \mathbb{R}^d$*

$$\int_{\mathbb{R}^d} \|x - D_\varepsilon^\star(x_\varepsilon)\|^2 g_\varepsilon(x|x_\varepsilon) dx \leq C, \quad (11)$$

where $g_\varepsilon(\cdot|x_\varepsilon)$ is the density of the conditional distribution of the unknown image $x \in \mathbb{R}^d$ with marginal distribution μ , given a noisy observation $x_\varepsilon \sim \mathcal{N}(x, \varepsilon \text{Id})$.

Proof. The proof is postponed to Lemma 27. □

These results can be generalised to hold under the weaker assumption that the expected MSE for D_ε^\star is finite but not uniformly bounded, as in this case $x \mapsto \nabla \log p_\varepsilon^\star(x|y)$ is locally instead of globally Lipschitz continuous (we postpone this technical extension to future work). The pathological case where D_ε^\star does not have a finite MSE arises when μ is such that the denoising problem does not admit a Bayesian estimator w.r.t. to the MSE loss. In this case, Tweedie's identity does not hold either. It is not clear at this point what kind of theoretical guarantees *Plug & Play* strategies can offer for such severely ill-posed problems.

Before concluding this section, we study whether the oracle $p^\star(x|y)$ is well posed in the sense of Hadamard, i.e., if $p^\star(x|y)$ changes continuously w.r.t. y under a suitable probability metric (see [55]). We answer positively to this question in Proposition 3 which states that, under mild assumptions on the likelihood, $p^\star(x|y)$ is locally Lipschitz continuous w.r.t. y for an appropriate metric. This stability result implies, for example, that the MMSE estimator derived from $p^\star(x|y)$ is locally Lipschitz continuous w.r.t. y , and hence stable w.r.t. small perturbations of y . Note that a similar property holds for the regularised posterior $p_\varepsilon^\star(x|y)$. In particular, Proposition 3 holds for Gaussian likelihoods (see Section 5 for details).

Proposition 3. *Assume that there exist $\Phi_1 : \mathbb{R}^d \rightarrow [0, +\infty)$ and $\Phi_2 : \mathbb{R}^m \rightarrow [0, +\infty)$ such that for any $x \in \mathbb{R}^d$ and $y_1, y_2 \in \mathbb{R}^m$*

$$\left| \log(p(y_1|x)) - \log(p(y_2|x)) \right| \leq (\Phi_1(x) + \Phi_2(y_1) + \Phi_2(y_2)) \|y_1 - y_2\|, \quad (12)$$

and for any $c > 0$, $\int_{\mathbb{R}^d} (1 + \Phi_1(\tilde{x})) \exp[c\Phi_1(\tilde{x})] p^\star(x) d\tilde{x} < +\infty$. Then $y \mapsto p^\star(\cdot|y)$ is locally Lipschitz w.r.t $\|\cdot\|_1$, i.e., for any compact set \mathcal{K} there exists $C_{\mathcal{K}} \geq 0$ such that for any $y_1, y_2 \in \mathcal{K}$, $\|p^\star(\cdot|y_1) - p^\star(\cdot|y_2)\|_1 \leq C_{\mathcal{K}} \|y_1 - y_2\|$.

Proof. The proof is a straightforward application of Proposition 9. □

To conclude, starting from the decision-theoretically optimal model $p^\star(x|y)$, we have constructed a regularised approximation $p_\varepsilon^\star(x|y)$ that is proper and smooth by construction. Under mild assumptions on $p(y|x)$, the approximation $p_\varepsilon^\star(x|y)$ is well posed in the sense of Hadamard and can be made arbitrarily close to the oracle $p^\star(x|y)$ by controlling ε . Moreover, we established that $x \mapsto \nabla \log p_\varepsilon^\star(x)$ is Lipschitz continuous when the problem of Gaussian image denoising for μ under the MSE loss is well posed. This allows analysing gradient-based algorithms for performing Bayesian computation for $p_\varepsilon^\star(x|y)$, which will play a central role in our subsequent analysis of *Plug & Play* schemes that mimic these idealised algorithms.

3.2 Bayesian computation with *Plug & Play* priors: oracle algorithms

We are now ready to analyse ULA and SGD for the regularised oracle $p_\varepsilon^*(x|y)$ and set the grounds for deriving *Plug & Play* schemes based on a generic denoiser D_ε . With the guarantees of Proposition 1 and Proposition 2 in mind, we use Tweedie's identity to express SGD for computing maximisers of $p_\varepsilon^*(x|y)$ as the following recursion: given $X_0 \in \mathbb{R}^d$ and for any $k \in \mathbb{N}$,

$$X_{k+1} = (1 - \delta_k/\varepsilon)X_k + \delta_k \nabla \log p(y|X_k) + (\delta_k/\varepsilon)D_\varepsilon^*(X_k) + 2\delta_k Z_{k+1}. \quad (13)$$

where we recall that $\{Z_k : k \in \mathbb{N}\}$ are i.i.d standard Gaussian random variables on \mathbb{R}^d . Proposition 4 below states that all the sequences generated by (13) that stay within a compact set converge to a critical point of $p_\varepsilon^*(x|y)$ almost surely (the compactness conditions is standard to establish convergence results for stochastic optimisation algorithms [33, 32, 61]).

Proposition 4. *Assume H1 and that $x \mapsto \nabla \log p_\varepsilon^*(x|y)$ is locally Lipschitz continuous. In addition, assume that $\lim_{k \rightarrow +\infty} \delta_k = 0$, $\sum_{k \in \mathbb{N}} \delta_k = +\infty$ and $\sum_{k \in \mathbb{N}} \delta_k^2 < +\infty$. Let*

$A_\varepsilon = \{\omega \in \Omega : \text{there exists } k_0 \in \mathbb{N}, K \subset \mathbb{R}^d \text{ compact such that for any } k \geq k_0, X_k(\omega) \in K\}$, where $(\Omega, \mathcal{F}, \mathbb{P})$ is the probability space on which $(X_k)_{k \in \mathbb{N}}$ is defined. Then, for any $\omega \in A_\varepsilon$

$$\limsup_{k \rightarrow +\infty} d(X_k(\omega), S_\varepsilon) = 0, \quad \text{with } S_\varepsilon = \{x \in \mathbb{R}^d : \nabla \log p_\varepsilon^*(x|y) = 0\}.$$

Proof. The proof is similar to the one of Proposition 22 upon setting $D_\varepsilon \leftarrow D_\varepsilon^*$ and using [80, Theorem 2.1-(a)] instead of [80, Theorem 2.1-(c)]. \square

Similarly, ULA applied to $p_\varepsilon^*(x|y)$ is given by the following recursion: for any $k \in \mathbb{N}$

$$X_{k+1} = (1 - \delta_k/\varepsilon)X_k + \delta_k \nabla \log p(y|X_k) + (\delta_k/\varepsilon)D_\varepsilon^*(X_k) + \sqrt{2\delta_k}Z_{k+1}, \quad (14)$$

Under standard assumptions on the sequence of step-sizes $(\delta_k)_{k \in \mathbb{N}}$, the sequence generated by (14) is a Markov chain which admits an invariant probability distribution whose density is provably close to $p_\varepsilon^*(x|y)$. However, to guarantee geometrically fast convergence to $p_\varepsilon^*(x|y)$ -a basic property for efficient Bayesian computation- it is also necessary to control the behaviour of the tails of $p_\varepsilon^*(x|y)$. We consider two alternatives to guarantee geometric convergence with markedly different bias-variance trade-offs: one with excellent accuracy guarantees but that requires using small step-sizes $(\delta_k)_{k \in \mathbb{N}}$ and hence has a higher computational cost, and another one that allows taking larger step-sizes $(\delta_k)_{k \in \mathbb{N}}$ to improve convergence speed at the expense of weaker guarantees in terms of estimation bias.

For the first method, in the spirit of the Moreau-Yosida regularised ULA, we introduce a minor modification to (14) so that the conditions for geometrically fast convergence are automatically verified, at the expense of some small estimation bias [30, 37]. More precisely, we target the modified posterior density $p_{\varepsilon, C}^*$ where for any $x \in \mathbb{R}^d$ we have

$$p_{\varepsilon, C}^*(x|y) \propto p_\varepsilon(x|y) \exp[-d^2(x, C)/\lambda],$$

where $C \subset \mathbb{R}^d$ is some large compact convex set that controls the far tail behaviour of $p_{\varepsilon, C}^*(x|y)$, $d(\cdot, C)$ is the Euclidean distance to C , and $\lambda > 0$ is a tail regularisation parameter that we set so that $p_{\varepsilon, C}^*(x|y)$ satisfies a curvature condition as $\|x\| \rightarrow \infty$ (see Section 5 for details). This leads to the following Markov chain to sample approximately from $p_\varepsilon^*(x|y)$ (and from $p^*(x|y)$ when the latter exists):

$$\begin{aligned} X_{k+1} = & (1 - \delta_k/\varepsilon)X_k + \delta_k \nabla \log p(y|X_k) + (\delta_k/\varepsilon)D_\varepsilon^*(X_k) \\ & + (\delta_k/\lambda)(\Pi_C(X_k) - X_k) + \sqrt{2\delta_k}Z_{k+1}. \end{aligned} \quad (15)$$

Proposition 5 below provides detailed non-asymptotic convergence bounds on the estimation error involved in using (15) for Bayesian computation w.r.t. $p_\varepsilon^*(x|y)$ and $p^*(x|y)$. The bounds for computation w.r.t. $p_\varepsilon^*(x|y)$ include three terms related to the size of \mathbf{C} , the sequence of step-sizes $(\delta_k)_{k \in \mathbb{N}}$ (assumed here to be constant $\delta_k = \delta$ for any $k \in \mathbb{N}$ for presentation clarity), and the number of iterations n . The bound w.r.t. $p^*(x|y)$ includes an additional term related to the smoothing effect of $\varepsilon > 0$.

Proposition 5. *Assume H1, that $\int_{\mathbb{R}^d} \|\tilde{x}\|^4 p^*(\tilde{x}) d\tilde{x} < +\infty$ and that (11) is satisfied. Then the following hold:*

(a) *There exists $\varepsilon_0 > 0$ such that for any $\varepsilon \in (0, \varepsilon_0]$ there exist $\lambda_0, \bar{\delta}, C_1 > 0$ such that for any \mathbf{C} convex compact with $\bar{\mathbf{B}}(0, R_{\mathbf{C}}) \subset \mathbf{C}$, $R_{\mathbf{C}} > 0$ there exists $C_2 > 0$ such that for any $h : \mathbb{R}^d \rightarrow \mathbb{R}$ with $\sup_{x \in \mathbb{R}^d} \{|h(x)| (1 + \|x\|^2)^{-1}\} \leq 1$, $n \in \mathbb{N}^*$, $\lambda \in (0, \lambda_0]$ and $\delta \in (0, \bar{\delta}]$ we have*

$$\left| n^{-1} \sum_{k=1}^n \mathbb{E}[h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) p_\varepsilon^*(\tilde{x}|y) d\tilde{x} \right| \leq \left\{ C_1 R_{\mathbf{C}}^{-1} + C_2 (\delta^{1/2} + (n\delta)^{-1}) \right\} (1 + \mathbb{E}[\|X_0\|^4]).$$

(b) *In addition, if the conditions of Proposition 1-(e) are satisfied, then there exist $\varepsilon_0, C_0 > 0$ such that for any $\varepsilon \in (0, \varepsilon_0]$ there exist $\lambda_0, \bar{\delta}, C_1 > 0$ such that for any \mathbf{C} convex compact with $\bar{\mathbf{B}}(0, R_{\mathbf{C}}) \subset \mathbf{C}$, $R_{\mathbf{C}} > 0$ there exists $C_2 > 0$ such that for any $h : \mathbb{R}^d \rightarrow \mathbb{R}$ with $\sup_{x \in \mathbb{R}^d} \{|h(x)| (1 + \|x\|^2)^{-1}\} \leq 1$, $n \in \mathbb{N}^*$, $\lambda \in (0, \lambda_0]$ and $\delta \in (0, \bar{\delta}]$ we have*

$$\left| n^{-1} \sum_{k=1}^n \mathbb{E}[h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) p^*(\tilde{x}|y) d\tilde{x} \right| \leq \left\{ C_0 \varepsilon^{\beta/4} + C_1 R_{\mathbf{C}}^{-1} + C_2 (\delta^{1/2} + (n\delta)^{-1}) \right\} (1 + \mathbb{E}[\|X_0\|^4]).$$

Proof. The proof follows from Proposition 16 and Proposition 21 with $D_\varepsilon \leftarrow D_\varepsilon^*$. \square

An alternative strategy (which we Projected PnP-ULA, *i.e.* PPnP-ULA) for ensuring numerical stability and fast convergence to the invariant measure is to modify ULA to include a hard projection onto \mathbf{C} , *i.e.* $(X_k)_{k \in \mathbb{N}}$ is defined by $X_0 \in \mathbf{C}$ and the following recursion for any $k \in \mathbb{N}$

$$X_{k+1} = \Pi_{\mathbf{C}} \left[(1 - \delta_k/\varepsilon) X_k + \delta_k \nabla \log p(y|X_k) + (\delta_k/\varepsilon) D_\varepsilon^*(X_k) + \sqrt{2\delta_k} Z_{k+1} \right], \quad (16)$$

where we notice that, by construction, the chain cannot exit \mathbf{C} because of the action of the projection operator $\Pi_{\mathbf{C}}$. This guarantees geometrically fast convergence to the chain's invariant measure without imposing any restrictions on the step-sizes $(\delta_k)_{k \in \mathbb{N}}$, at the expense of bias from truncating the support of $p_\varepsilon^*(x|y)$. However, we cannot explicitly control the bias associated with using (16), unless we impose strict restrictions on $(\delta_k)_{k \in \mathbb{N}}$, in which case Proposition 18 regarding the bias of (16) holds. Despite weaker accuracy guarantees, (16) is interesting because it allows taking larger step-sizes to improve convergence speed without risking numerical instability. Conversely, (15) can become explosive if the step-sizes are too large.

To summarise, the ULA and SGD algorithms defined by (13), (15) and (16) are well defined and have clear theoretical convergence guarantees, but they cannot be implemented in practice as they require evaluating $D_\varepsilon^*(x)$. Nevertheless, they provide a powerful and natural strategy for analysing *Plug & Play* SGD and ULA schemes that can be viewed as approximations of (13), (15) and (16).

3.3 Bayesian computation with *Plug & Play* priors: tractable algorithms

We are now ready to study the PnP-ULA, Projected PnP-ULA (PPnP-ULA), and PnP-SGD *Plug & Play* schemes described in Algorithms 1, 2 and 3, obtained by replacing D_ε^* in (15), (16) and (13) with a generic denoiser D_ε . We pay special attention to state-of-the-art image denoisers implemented by neural networks, which are usually explicitly trained to achieve optimal MSE performance and hence to approximate D_ε^* . For generality, we allow the use of a regularisation parameter $\alpha > 0$ to control amount of regularity enforced by the prior (in our experiments we set $\alpha = 1$). For presentation clarity, in this section we present simplified convergence results for 1, 2 and 3 by assuming $\alpha = 1$ and constant a step-size $\delta_k = \delta$ for any $k \in \mathbb{N}$. These results are presented in full generality in Section 5.

Algorithm 1 PnP-ULA

Require: $n \in \mathbb{N}$, $y \in \mathbb{R}^m$, $\varepsilon, \lambda, \alpha, \delta > 0$, $C \subset \mathbb{R}^d$ convex and compact

Ensure: $2\lambda(2L_y + \alpha L/\varepsilon) \leq 1$ and $\delta < (1/3)(L_y + 1/\lambda + \alpha L/\varepsilon)^{-1}$

Initialization: Set $X_0 = y$ and $k = 0$.

for $k = 0 : N$ **do**

$Z_{k+1} \sim \mathcal{N}(0, \text{Id})$

$X_{k+1} = X_k + \delta \nabla \log(p(y|X_k)) + (\alpha\delta/\varepsilon)(D_\varepsilon(X_k) - X_k) - (\delta/\lambda)(\Pi_C(X_k) - X_k) + \sqrt{2\delta}Z_{k+1}$

end for

return $\{X_k : k \in \{0, \dots, N+1\}\}$

Algorithm 2 PPnP-ULA

Require: $n \in \mathbb{N}$, $y \in \mathbb{R}^m$, $\varepsilon, \lambda, \alpha, \delta > 0$, $C \subset \mathbb{R}^d$ convex and compact

Initialization: Set $X_0 = y$ and $k = 0$.

for $k = 0 : N$ **do**

$Z_{k+1} \sim \mathcal{N}(0, \text{Id})$

$X_{k+1} = \Pi_C \left(X_k + \delta \nabla \log(p(y|X_k)) + (\alpha\delta/\varepsilon)(D_\varepsilon(X_k) - X_k) + \sqrt{2\delta}Z_{k+1} \right)$

end for

return $\{X_k : k \in \{0, \dots, N+1\}\}$

Proposition 6 provides convergence guarantees for PnP-ULA in Algorithm 1 under the assumption that D_ε is Lipschitz continuous, and explicitly characterises the additional bias that using D_ε instead of D_ε^* introduces in the resulting Monte Carlo estimates. Similarly Proposition 7 provides convergence guarantees for PPnP-ULA in Algorithm 2. Finally, Proposition 8 provides convergence guarantees for PnP-SGD in Algorithm 3. The assumption that D_ε is Lipschitz continuous is verified, for example, by the neural network denoiser [74] that controls the Lipschitz constant of the residual $(\text{Id} - D_\varepsilon)$ during the training process, and which we use in the experiments of Section 4.

Proposition 6. *Assume the same conditions as Proposition 5. In addition, assume that there exist $\varepsilon_0 > 0$, $L \geq 0$, and a function $M : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ such that for any $\varepsilon \in (0, \varepsilon_0]$, $x_1, x_2 \in \mathbb{R}^d$ and $x \in \bar{B}(0, R)$ we have*

$$\|D_\varepsilon(x_1) - D_\varepsilon(x_2)\| \leq L \|x_1 - x_2\|, \quad \|D_\varepsilon(x) - D_\varepsilon^*(x)\| \leq M(R).$$

Then the following hold:

Algorithm 3 PnP-SGD

Require: $n, n_{\text{burnin}} \in \mathbb{N}$, $y \in \mathbb{R}^m$, $\varepsilon, \lambda, \alpha, \delta > 0$, $\mathbf{C} \subset \mathbb{R}^d$ convex and compact

Initialization: Set $X_0 = \tilde{x}$ and $k = 0$.

for $k = 0 : N$ **do**

$Z_{k+1} \sim \mathcal{N}(0, \text{Id})$

if $k \leq n_{\text{burnin}}$ **then**

$X_{k+1} = X_k + \delta_0 \nabla \log(p(y|X_k)) + (\delta_0 \alpha / \varepsilon)(D_\varepsilon(X_k) - X_k) + \delta_0 Z_{k+1}$

end if

if $k \geq n_{\text{burnin}}$ **then**

$X_{k+1} = X_k + \delta_k \nabla \log(p(y|X_k)) + (\delta_k \alpha / \varepsilon)(D_\varepsilon(X_k) - X_k) + \delta_k Z_{k+1}$

$\delta_{k+1} = \delta_0 (k + 1 - n_{\text{burnin}})^{-0.8}$

end if

end for

return $\{X_k : k \in \{0, \dots, N + 1\}\}$

(a) There exists $\varepsilon_0 > 0$ such that for any $\varepsilon \in (0, \varepsilon_0]$ there exist $\lambda_0, \bar{\delta}, C_1 > 0$ such that for any \mathbf{C} convex compact with $\bar{\mathbf{B}}(0, R_{\mathbf{C}}) \subset \mathbf{C}$, $R_{\mathbf{C}} > 0$ there exists $C_2 > 0$ such that for any $h : \mathbb{R}^d \rightarrow \mathbb{R}$ with $\sup_{x \in \mathbb{R}^d} \{ |h(x)| (1 + \|x\|^2)^{-1} \} \leq 1$, $n \in \mathbb{N}^*$, $\lambda \in (0, \lambda_0]$ and $\delta \in (0, \bar{\delta}]$ we have

$$\left| n^{-1} \sum_{k=1}^n \mathbb{E}[h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) p_\varepsilon^*(\tilde{x}|y) d\tilde{x} \right| \leq \left\{ C_1 R_{\mathbf{C}}^{-1} + C_2 (\delta^{1/2} + \inf_{R>0} \{ \mathbf{M}(R) + \exp[-R] \} + (n\delta)^{-1}) \right\} (1 + \mathbb{E}[\|X_0\|^4]).$$

(b) In addition, if the conditions of Proposition 1-(e) are satisfied, then there exist $\varepsilon_0, C_0 > 0$ such that for any $\varepsilon \in (0, \varepsilon_0]$ there exist $\lambda_0, \bar{\delta}, C_1 > 0$ such that for any \mathbf{C} convex compact with $\bar{\mathbf{B}}(0, R_{\mathbf{C}}) \subset \mathbf{C}$, $R_{\mathbf{C}} > 0$ there exists $C_2 > 0$ such that for any $h : \mathbb{R}^d \rightarrow \mathbb{R}$ with $\sup_{x \in \mathbb{R}^d} \{ |h(x)| (1 + \|x\|^2)^{-1} \} \leq 1$, $n \in \mathbb{N}^*$, $\lambda \in (0, \lambda_0]$ and $\delta \in (0, \bar{\delta}]$ we have

$$\left| n^{-1} \sum_{k=1}^n \mathbb{E}[h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) p^*(\tilde{x}|y) d\tilde{x} \right| \leq \left\{ C_0 \varepsilon^{\beta/4} + C_1 R_{\mathbf{C}}^{-1} + C_2 (\delta^{1/2} + \inf_{R>0} \{ \mathbf{M}(R) + \exp[-R] \} + (n\delta)^{-1}) \right\} (1 + \mathbb{E}[\|X_0\|^4]).$$

Proof. The result is an application of Proposition 16 and Proposition 21. Explicit bounds on the parameters $\bar{\delta}$ and λ_0 are provided in Proposition 21, see also Algorithm 1. \square

Proposition 7. Assume the same conditions as Proposition 6. Then, for any $\varepsilon \in (0, \varepsilon_0]$, $\delta > 0$ and \mathbf{C} convex and compact with $0 \in \mathbf{C}$, there exists a probability measure $\pi_{\varepsilon, \delta}^{\mathbf{C}}$, $\tilde{A}_{\mathbf{C}} \geq 0$ and $\tilde{\rho}_{\mathbf{C}} \in [0, 1)$ such that for any $n \in \mathbb{N}$

$$\|\mathcal{L}(X_n) - \pi_{\varepsilon, \delta}^{\mathbf{C}}\|_{\text{TV}} \leq \tilde{A}_{\mathbf{C}} \rho_{\mathbf{C}}^{n\delta},$$

where $\mathcal{L}(X_n)$ is the distribution of X_n given by PnP-UULA and $\|\cdot\|_{\text{TV}}$ is the total variation norm.

Proof. The proof is a straightforward application of Proposition 17. \square

Proposition 8. *Assume the same conditions as Proposition 6. In addition, assume that $x \mapsto \log(p(y|x))$ is real-analytic and that $\lim_{k \rightarrow +\infty} \delta_k = 0$, $\sum_{k \in \mathbb{N}} \delta_k = +\infty$ and $\sum_{k \in \mathbb{N}} \delta_k^2 < +\infty$. Let $R > 0$, $\mathsf{K} \subset \overline{\mathsf{B}}(0, R)$ be a compact set, $X_0 \in \mathbb{R}^d$ and $\mathsf{A}_{\varepsilon, \mathsf{K}} \in \mathcal{F}$ given by*

$$\mathsf{A}_{\varepsilon, \mathsf{K}} = \{\omega \in \Omega : \text{there exists } k_0 \in \mathbb{N} \text{ such that for any } k \geq k_0, X_k(\omega) \in \mathsf{K}\}.$$

Then there exist $C_{\varepsilon, \mathsf{K}} \geq 0$ and $r_{\varepsilon, \mathsf{K}} \in (0, 1)$ such that $\limsup_{k \rightarrow +\infty} d(X_k(\omega), \mathsf{S}_{\varepsilon, \mathsf{K}}) \leq C_{\varepsilon, \mathsf{K}} M_R^{r_{\varepsilon, \mathsf{K}}}$ for any $\omega \in \mathsf{A}_{\varepsilon, \mathsf{K}}$, with $\mathsf{S}_{\varepsilon, \mathsf{K}} = \{x \in \mathsf{K} : \nabla \log p_\varepsilon^(x|y) = 0\}$.*

Proof. The result is a straightforward application of Proposition 22. \square

Lastly, it is worth mentioning that Algorithm 1, Algorithm 2 and Algorithm 3 can be straightforwardly modified to incorporate additional regularisation terms. More precisely, one could consider a prior defined as the (normalised) product of a *Plug & Play* term and an explicit analytical term. In that case, one should simply modify the recursion defining the Markov chain by adding the gradient associated with the analytical term. In a manner akin to [38], analytical terms that are not smooth are involved via their proximal operator.

The following section demonstrates the proposed PnP-ULA, PPnPPULA and PnP-SGD methodology on three canonical imaging inverse problems, namely image deblurring, inpainting, and denoising. Before concluding this section, it is worth emphasising that, in addition to being important in their own right, Algorithm 1, Algorithm 2 and Algorithm 3 and the associated theoretical results set the grounds for analysing more advanced stochastic simulation and optimisation schemes for performing Bayesian inference with *Plug & Play* priors, in particular accelerated optimisation and sampling algorithms [67]. This is an important perspective for future work.

4 Experimental study

This section illustrates the behavior of PnP-ULA, PPnPP-ULA and PnP-SGD for several inverse problems. We start with two classical imaging problems: deblurring and inpainting. For these two problems, we first compare the pointwise estimators (MMSE and MAP) provided by these algorithms, along with the MAP estimator provided by the recent PnP-ADMM method [74]. Following on from this, we show how to analyse the convergence of the PnP-ULA algorithm, and conclude the experiments with an uncertainty visualisation analysis. This is then followed by an illustration of uncertainty quantification in a medical imaging problem.

In each of these experiments the observation model takes the form

$$y = \mathbf{A}x + n, \tag{17}$$

where $x \in \mathbb{R}^d$ is the unknown original image, $y \in \mathbb{R}^d$ the observed image, n is a realization of a Gaussian i.i.d. centered noise with variance $\sigma^2 \text{Id}$, and \mathbf{A} is a linear operator corresponding to the considered degradation. The log-likelihood for this case writes $\log p(y|x) = -\|\mathbf{A}x - y\|^2 / (2\sigma^2)$. For all algorithms, the denoising operator D_ε is chosen as the pretrained denoising neural network introduced in [74], since it satisfies that $(D_\varepsilon - \text{Id})$ is L-Lipschitz with $L < 1$.

4.1 Implementation guidelines and parameter setting

In the following, we provide some simple and robust rules in order to set the parameters of the different algorithms, in particular the discretization step-size δ and the tail regularization parameter λ for PnP-ULA.

Step-size δ The parameter δ controls the asymptotic accuracy of PnP-ULA and PPnP-ULA, as well as the speed of convergence to stationarity. This leads to the following bias-variance trade-off. For large values of δ , the Markov chain has low auto-correlation and converges quickly to its stationary regime. Consequently, the Monte Carlo estimates computed from the chain exhibit low asymptotic variance, at the expense of some asymptotic bias. On the contrary, small values of delta δ produce a Markov chain that explores the parameter space less efficiently, but more accurately. As a result, the asymptotic bias is smaller, but the variance is larger. In the context of inverse problems that are high-dimensional and ill-posed, properly exploring the solution space can take a large number of iterations. For this reason, we recommend using large values of δ , at the expense of some bias. In addition, in PnP-ULA, δ is also subject to a numerical stability constraint related to the inverse of the Lipschitz constant of $\nabla \log p_\varepsilon(x|y)$; namely, we require $\delta < (1/3)(\alpha L/\varepsilon + L_y + 1/\lambda)^{-1}$ where L and L_y are respectively the Lipschitz constant of the denoiser residual $(D_\varepsilon - \text{Id})$ and the Lipschitz constant of the log-likelihood gradient. In our experiments, $L = 1$ and $L_y = \|\mathbf{A}^* \mathbf{A}\|/\sigma^2$, so we choose δ just below the upper bound $1/3(\alpha/\varepsilon + \|\mathbf{A}^* \mathbf{A}\|/\sigma^2 + 1/\lambda)^{-1}$, where \mathbf{A}^* is the adjoint of \mathbf{A} . For PPnP-ULA, we recommend setting $\delta < (\alpha L/\varepsilon + L_y)^{-1}$ to prevent excessive bias.

Parameter λ The parameter λ controls the tail behaviour of the target density. As previously explained, it must be set so that the tails of the target density decay sufficient fast to ensure convergence at a geometric rate, a key property for guaranteeing that the Monte Carlo estimates computed from the chain are consistent and subject to a Central Limit Theorem with the standard $\mathcal{O}(\sqrt{k})$ rate. More precisely, we require $\lambda \in (0, 1/2(\alpha L/\varepsilon + 2L_y))$. Within this admissible range, if λ is too small this limits the maximal δ and leads to a slow Markov chain. For this reason, we recommend setting λ as large as possible below $(2\alpha L/\varepsilon + 4L_y)^{-1}$. By choosing λ and δ in this manner, both Proposition 6 and Proposition 8 hold.

Other parameters The compact set \mathbf{C} is defined as $\mathbf{C} = [-1, 2]^d$, even if in practice no samples went outside of \mathbf{C} in all our experiments, which suggests that the tail decay conditions hold without explicitly enforcing them. In all our experiments, we set the regularisation parameter α to its default value 1, and set the noise level of the denoiser D_ε to $\varepsilon = (5/255)^2$. The only exception is the denoiser used in PnP-ADMM for the inpainting experiments, trained at a noise level of $(40/255)^2$ to improve the performance. All algorithms are implemented using Python and the PyTorch library. Our experiments are run on an Intel Xeon CPU E5-2609 server with a Nvidia Titan XP graphic card.

4.2 Deblurring and inpainting: pointwise estimation

In this section, we address two inverse problems of the form (17): deblurring, where \mathbf{A} is a 9×9 box blur operator, and inpainting, where the matrix \mathbf{A} is diagonal, composed of 0 (for missing values) and 1. For the inpainting experiments, we focus on a case where 80% of the pixels are hidden. In all experiments, we add a Gaussian noise with variance $\sigma^2 = (1/255)^2$. Figure 1 shows the three original 256×256 images used in these experiments. Among them, the *Traffic* image is of particular interest since it consists of a mix of piecewise-constant areas and complex textures. Figure 2 shows the corresponding blurry images and Figure 3 shows the images to inpaint.

Because we will provide later in this section a complete and thorough convergence study of PnP-ULA, we run the algorithm for 25×10^6 iterations (65 hours) and 2.5×10^6 burn-in iterations, although the simple MMSE estimator can be obtained much faster in practice, as we will see later. After this burn-in period, we consider a thinned version of the Markov chain by

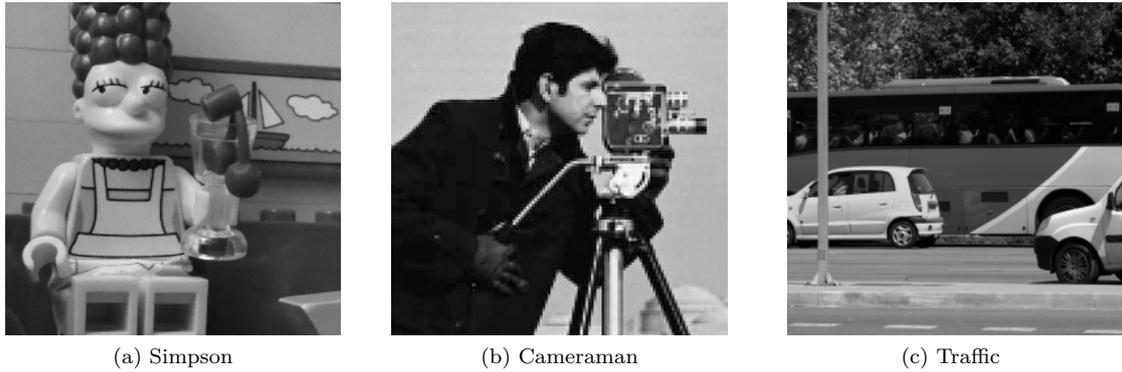


Figure 1: Original images used for the deblurring and inpainting experiments.

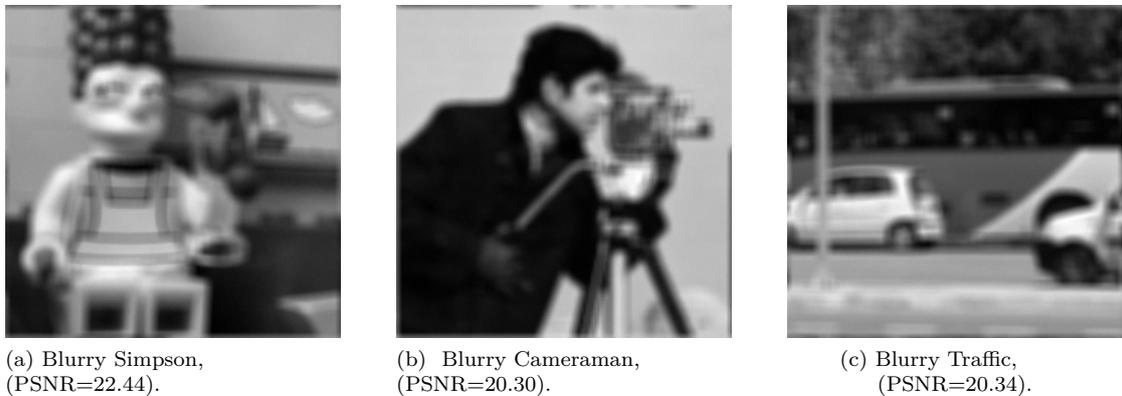


Figure 2: Images of Figure 1, blurred using a 9×9 -box-filter operator and corrupted by an additive Gaussian white noise with standard deviation $\sigma = 1/255$.

storing a sample every 2500 iterations, and the MMSE is obtained here as the average of these samples, both for the deblurring and inpainting experiments. For simplicity, the algorithm is always initialized with the degraded observation y . The PnP-SGD algorithm is initialized with the solution provided by PnP-ADMM and run for 10^4 iterations with 5×10^3 burn-in iterations (2 minutes). We emphasize the fact that PnP-ADMM does not minimize the same objective function as PnP-SGD. As a result, there is no reason that the fixed point provided by PnP-ADMM and the PnP-SGD stationary point are the same. We compare these results with the solution provided by the PnP-ADMM algorithm, initialized with the degraded observation y , running during 10^3 iterations for deblurring and 10^4 iterations for inpainting.

Results Figure 4 and Figure 5 show the deblurring and inpainting results obtained by three algorithms: the MMSE for PnP-ULA and the MAP for PnP-SGD and PnP-ADMM. We also provide on the same figures the *Peak Signal-To Noise Ratio* (PSNR) and the *Structural Similarity Index* (SSIM) [87, 88] for all these experiments. From a quantitative viewpoint, the MMSE computed by PnP-ULA achieves the best SSIM results for all images and the best PSNR for five out of six experiments, while the MAP estimator computed by PnP-SGD usually shows

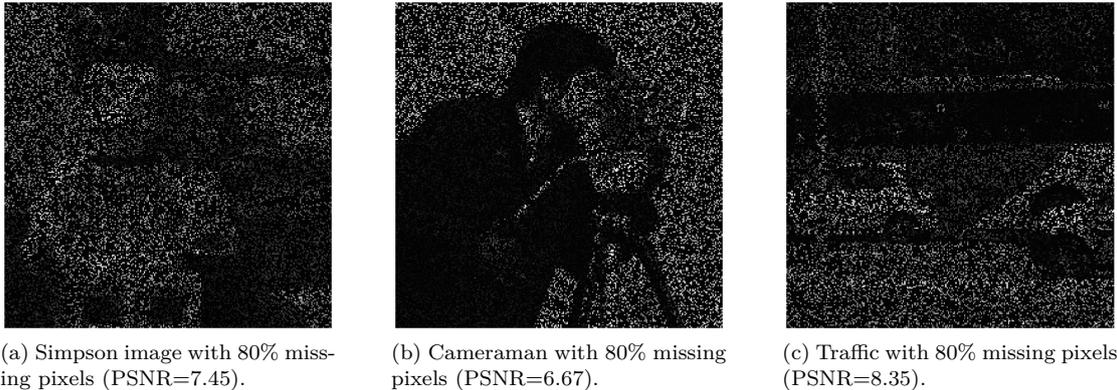


Figure 3: Images of Figure 1, with 80% missing pixels and corrupted by an additive Gaussian white noise with standard deviation $\sigma = 1/255$.

the lowest performance of the three algorithms. Interestingly, a careful visual inspection of the deblurring and inpainting results shows that the estimators computed by the three algorithms are quite different (see also Figure 6 and Figure 7). For the deblurring experiment, the results of PnP-SGD are not as sharp as the ones of the two other estimators, while the PnP-ADMM results seem to contain more details but exhibit a grid pattern artefact on all images. The MMSE computed by PnP-ULA seems to provide the best visual results, both sharp and without important visual defects. For the inpainting experiments, PnP-ULA and PnP-ADMM provide quite close visual results, although PnP-ADMM struggles to restore some specific regions (like the license plate in the Traffic image).

Observe that the inferior results of PnP-SGD and PnP-ADMM may be due to several factors and, on their own, do not imply that the MMSE estimator is more accurate than the MAP estimator in these cases. Unlike PnP-ULA, PnP-SGD and PnP-ADMM are actually more prone to algorithmic failures: *(i)* they may get stuck in a sub-optimal critical point of the posterior (PnP-SGD) or at a fixed point that is not even a critical point (PnP-ADMM); *(ii)* convergence of PnP-ADMM is not guaranteed since \mathbf{A} is not full-rank, so the data fitting term is not strongly convex; *(iii)* for the MAP estimator, the optimal regularization parameter α often differs from the canonical choice $\alpha = 1$ used here.

Faster MMSE estimation As suggested above, computing MMSE estimators does not necessitate millions of iterations of PnP-ULA. For the deblurring experiment, it is possible to use PPnP-ULA to obtain similar MMSE results in a only a few minutes. In this case, we remove the burn-in phase, use no thinning, and set δ to a value that is 24 times larger than the one specified for PnP-ULA in the previous section. To be more specific, for the *Simpsons* image, 15 000 iterations of PPnP-ULA (2 minutes) are enough to attain a SSIM of 0.94 and PSNR of 34.06. The convergence of the algorithm is ensured by the projection step. Convergence is slower for the inpainting experiments because in this case the likelihood has severe identifiability issues. For that problem, MMSE results as the one presented before for PnP-ULA necessitate approximately 10^5 PPnP-ULA iterations (15 minutes).

We emphasise at this point that we recommend to always use PPnP-ULA to warm-start PnP-ULA. Here we used PnP-ULA with $X_0 = y$ and a long number of burn-in iterations only to illustrate the robustness of PnP-ULA to a poor initialisation, as well as to simplify the pre-

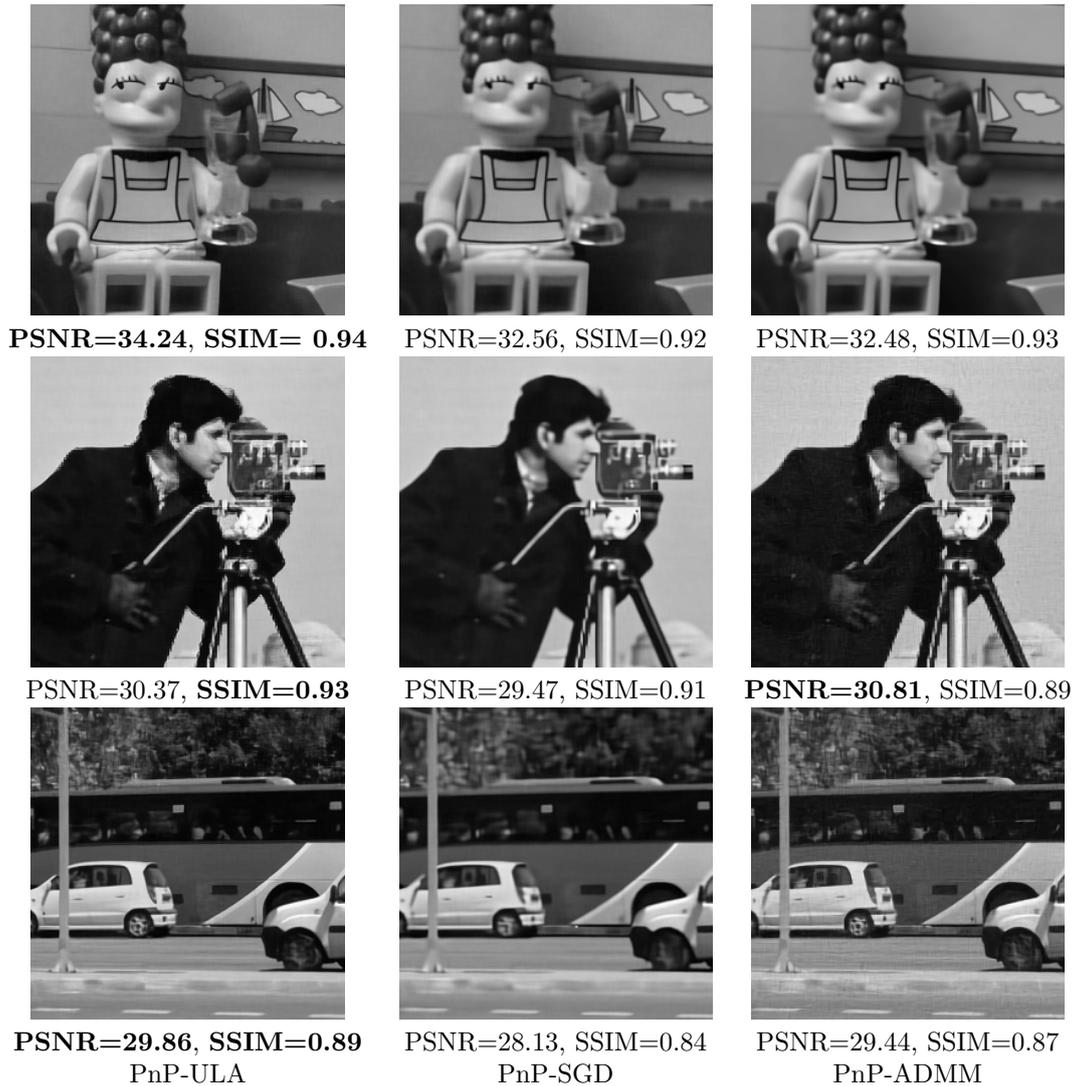


Figure 4: Deblurring of the images presented in Figure 2, using (left to right) PnP-ULA, PnP-SGD with $\alpha = 1$ and PnP-ADMM with the default parameters.

sensation of the experiments.

4.3 Deblurring and inpainting: convergence analysis for PnP-ULA

When using a sampling algorithm such as PnP-ULA, it is essential to check that the state space is correctly explored. We provide in Figure 8 and Figure 9 two convergence diagnostics for the deblurring and inpainting experiments of the previous section. These results are obtained for the *Simpson* image, but similar results were obtained for the two other images.

For each figure, on the left, we show the evolution of the Euclidean distance between the final MMSE estimate and the current stored sample (after the burn-in period, and hence in

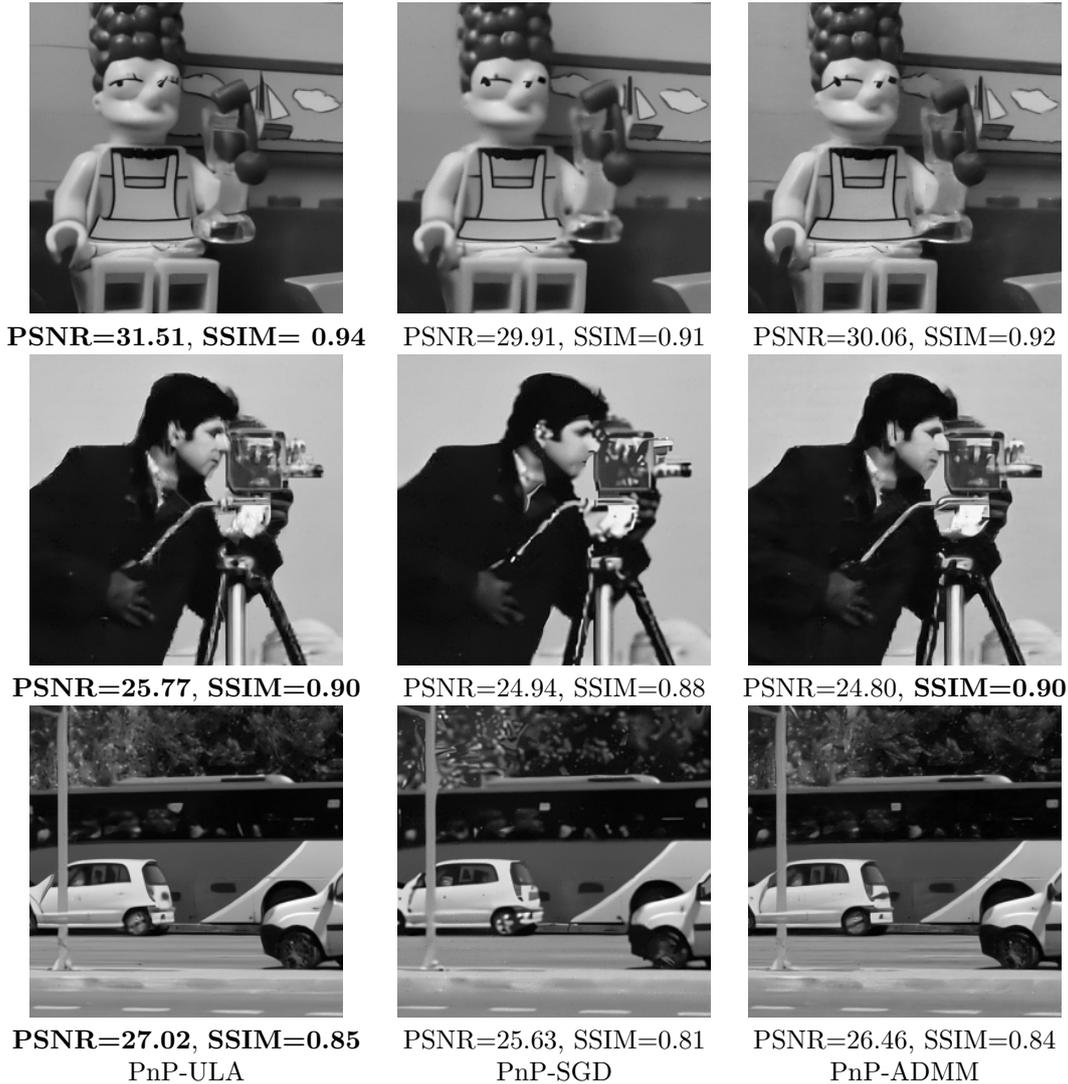


Figure 5: Inpainting of the images presented in Figure 3 using (left to right) PnP-ULA, PnP-SGD with $\alpha = 1$ and PnP-ADMM with the default parameters. PnP-ULA achieves better results when compared with the PnP-SGD and PnP-ADMM methods.

what considered to be a stationary regime). The fluctuations around the posterior mean and the absence of structure in these plots are a first indication that the chain has a low auto-correlation and adequately explores the solution space. On the right, we show the sample auto-correlation functions (ACFs) for the chains generated by PnP-ULA. These ACF plots measure how fast samples become uncorrelated, once the chain is in stationary regime. A fast decay of the auto-correlation is associated with a short mixing-time of the Markov chain, which in turn implies accurate Monte Carlo estimates. On the contrary a slow decay of the auto-correlation indicates that the Markov chain is moving slowly, which leads to Monte Carlo estimates with high variance. Because computing and visualising a multivariate ACF is difficult, here we show the ACF of the

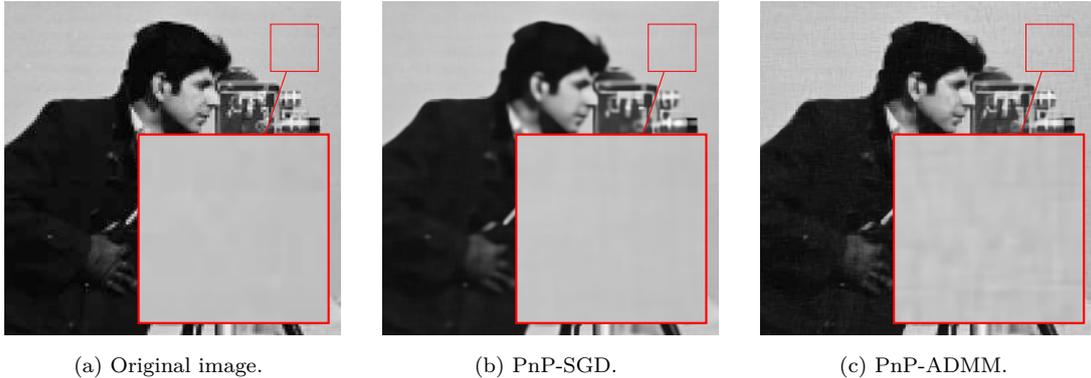


Figure 6: Comparison of the PnP-SGD and the PnP-ADMM solutions for the deblurring problem on the *Cameraman* image in Figure 1b.

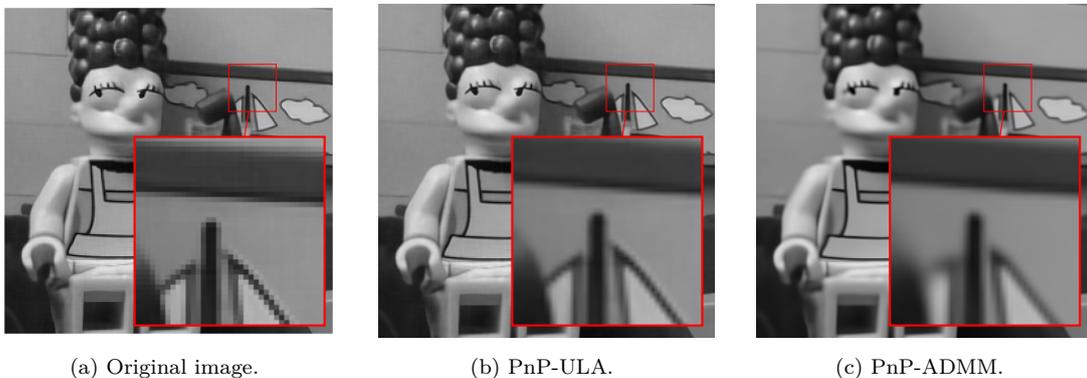


Figure 7: Comparison of the PnP-ULA and the PnP-ADMM solutions for the deblurring problem on the *Simpson* image in Figure 1a. Edges are sharper in the image obtained with PnP-ULA compared to the one obtained with PnP-ADMM.

chain along the slowest and the fastest one-dimensional sub-spaces in the Fourier domain⁶. In the deblurring case, these plots suggest that to achieve independence, samples need to be separated by approximately 7.5×10^4 iterations. As for the inpainting experiment, in the slowest direction, this requires in the order of 2.5×10^5 iterations. This is due to the fact that the likelihood function for the inpainting problem has severe identifiability issues, which in turn leads to a posterior distribution that is highly anisotropic and more difficult to explore as a result. To guarantee reliable Monte Carlo estimates, the total number of iterations should be significantly larger than the number of iterations required to produce two independent samples.

4.4 Deblurring and inpainting: uncertainty visualisation study

One of the benefits of sampling from the posterior distribution with PnP-ULA is that we can probe the uncertainty in the delivered solutions. In the following, we present an uncertainty

⁶The slowest direction corresponds to the Fourier coefficient with the highest variance along the samples.

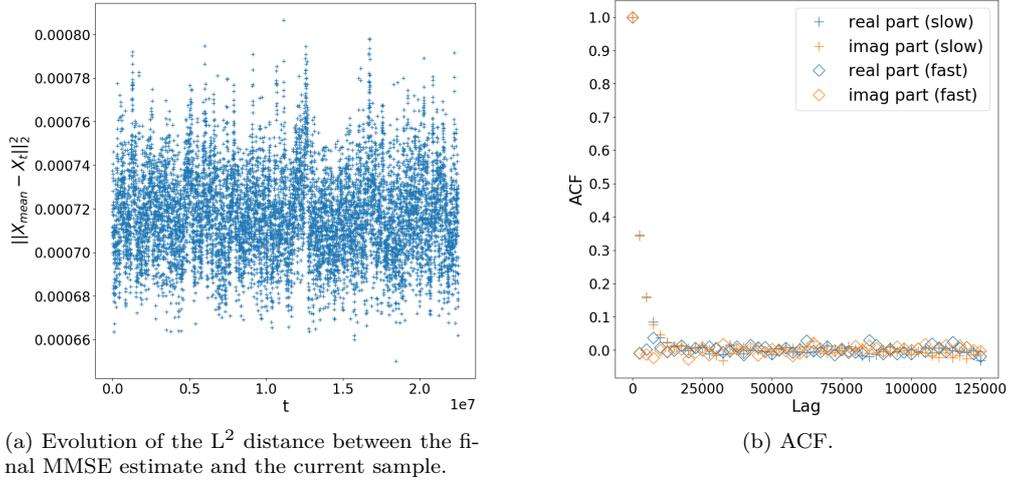


Figure 8: Quantitative analysis of the convergence of the PnP-ULA algorithm for the deblurring of the *Simpson* image (Figure 1a). After the burn-in iterations, the mean-square distance between the samples of the PnP-ULA Markov chains and the posterior mean fluctuates without exhibiting a pattern. From the ACF plots (Figure 8b), we derive that the Markov chain iterates become uncorrelated for a time-shift of 7.5×10^4 iterations.

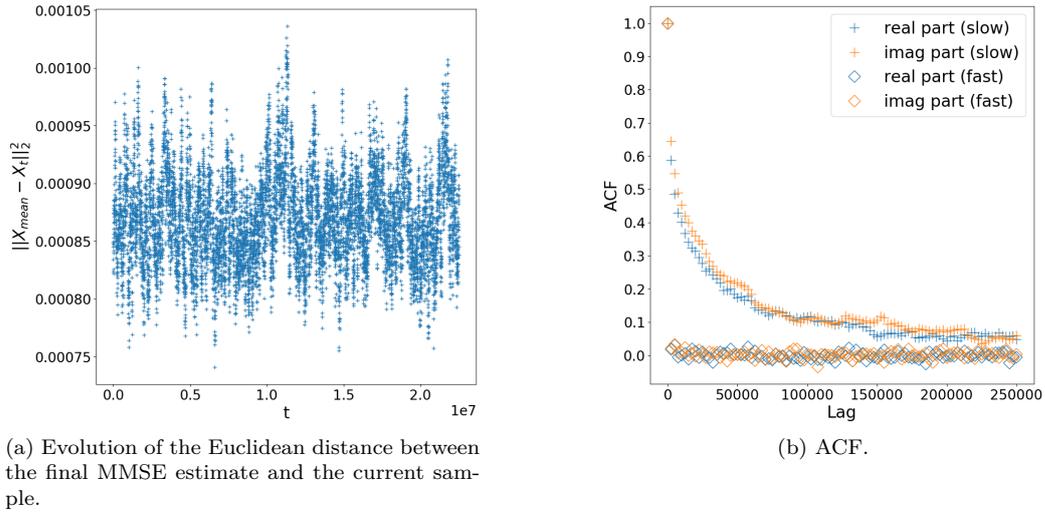


Figure 9: Quantitative analysis of the convergence of the PnP-ULA algorithm for the inpainting problem and *Simpson* image. On the left, the absence of pattern in the Euclidean distance between the final MMSE estimate and the current sample suggests that we properly explore the space. On the right, the ACF plot tells us that the Markov chain samples are still not independent after 2.5×10^5 iterations.

visualisation analysis that is useful for displaying the uncertainty related to image structures of different sizes and located in different regions of the image (see [24] for more details). The

analysis proceeds as follows. First, Figure 10 and Figure 11 show the marginal posterior standard deviation associated with each image pixel, as computed by PnP-ULA for three images and for the deblurring and inpainting problems, respectively. As could be expected, we observe that highly uncertain pixels are concentrated around the edges of the reconstructed images, but also on textured areas, although the uncertainty remains low on observed pixels for the inpainting problem. Notice that the dynamic range of the pixel standard deviation is larger for the inpainting problem than for the deblurring problem, which might suggest (again) that the problem has a higher level of intrinsic uncertainty.

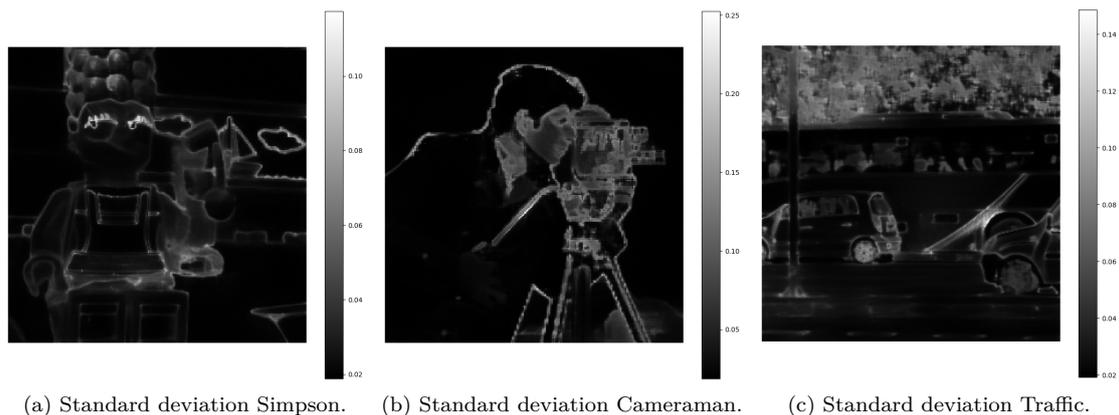


Figure 10: Standard deviation for the deblurring problem. On simple images such as *Simpson* (see Figure 1a), most of the uncertainty is located around the edges.

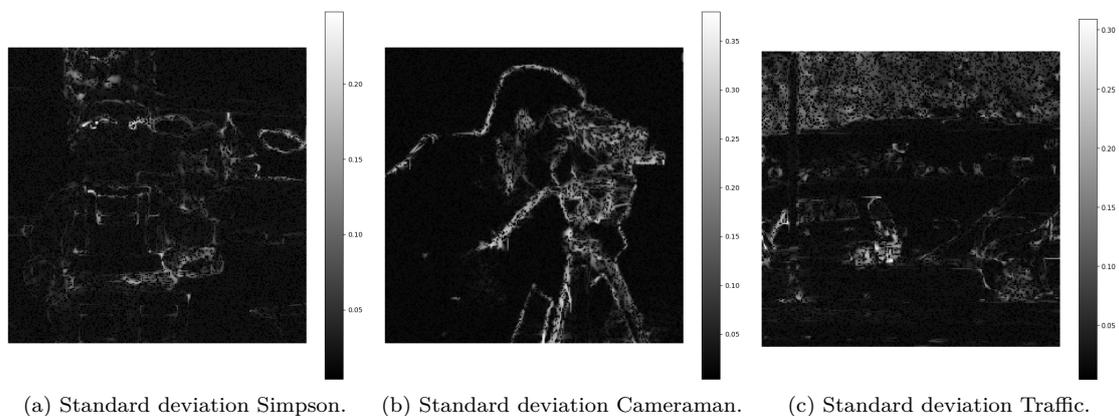


Figure 11: Standard deviation for the inpainting problem.

Following on from this, to explore the uncertainty for structures that are larger than one pixel, Figure 12 and Figure 13 report the marginal standard deviation associated with higher scales. More precisely, for different values of the scale i , we downsample the stored samples by a factor 2^i before computing the standard deviation. This downsampling step permits to quantify the uncertainty of larger or lower-frequency structures, such as the bottom of the glass in the

deblurring experiment in Figure 4. At each scale, we see that the uncertainty of the estimate is much more localized for the inpainting problem (resulting in higher uncertainty values in some specific regions) and more spread out for deblurring, certainly because of the different nature of the linear degradation \mathbf{A} .

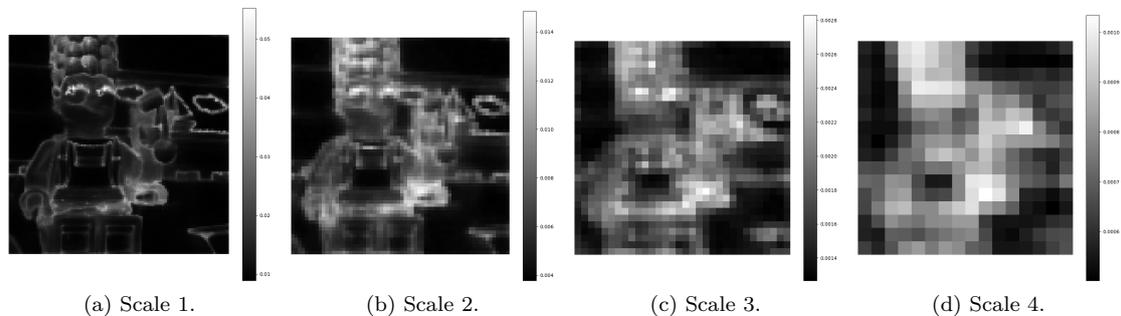


Figure 12: Marginal posterior standard deviation of the *Simpson* image for the deblurring problem at different scales. The scale i corresponds to a downsampling by a factor $2i$ of the original image size.

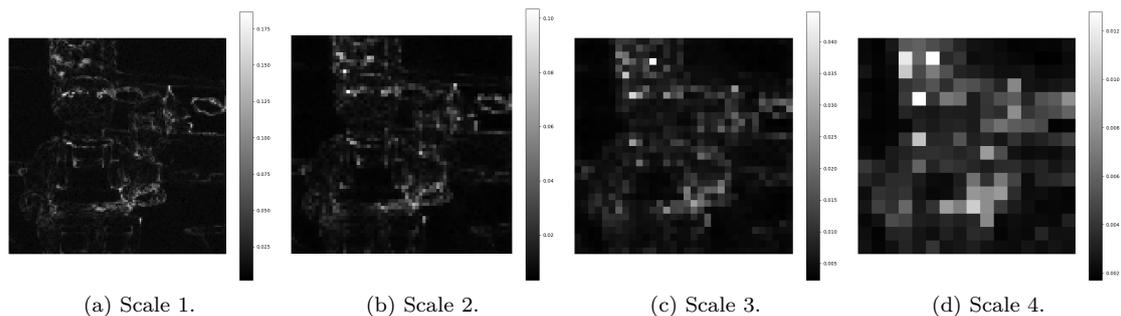


Figure 13: Marginal posterior standard deviation of the *Simpson* image for the inpainting problem at different scales. The scale i corresponds to a downsampling by a factor $2i$ of the original sample size.

4.5 Uncertainty quantification in medical imaging

We conclude our numerical experiments with an illustration of uncertainty quantification for a medical imaging problem. For the sake of simplicity, we focus on a simple denoising problem ($\mathbf{A} = \text{Id}$), with a Gaussian i.i.d centered noise of standard deviation $\sigma = 25/255$. Figure 14 shows the original image x on the left and the corrupted version y on the right. The image x is an axial computed tomography (CT) slice from the *Deep Lesion* dataset [91]. This CT slice is particularly interesting, because it contains an annotated lesion. Our goal is twofold: we want to perform pointwise denoising estimation, as well as to quantify the uncertainty on the estimation of the lesion's size.

For this experiment, PnP-ULA is initialized with the noisy observation y and run during 5×10^6 iterations with 5×10^4 burn-in iterations. Samples of the burn-in period are discarded

and we only consider a thinned Markov chain of samples stored every 500 iterations. The regularisation parameter α is set to its default value 1.

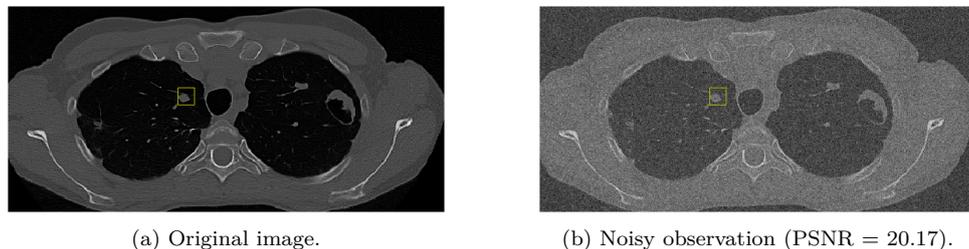


Figure 14: Original CT slice and its noisy observation obtained by adding a Gaussian white noise with standard deviation $\sigma = 25/255$. The yellow square corresponds to the lesion annotation provided by [91].

Figure 15a shows the MMSE solution computed by PnP-ULA. Visually, this restoration seems accurate, and this is confirmed by the PSNR score. Figure 15b shows the marginal posterior standard deviation for the image pixels. Edges are the main source of uncertainty. As we can see, the outline of the lesion is a region of high uncertainty, which motivates the following simple uncertainty study on the lesion’s size. To this aim, we define as a region of interest the yellow square annotated in the slice. Then, for each sample, we perform a simple segmentation of the lesion by using a thresholding technique, and estimate the lesion’s size by counting the number of pixels in the segmented region (of course, more sophisticated segmentation techniques could be used too). Figure 16 shows the histogram of the estimated lesion sizes, as well as the true size and lower and upper limits of a 95% credible interval. We observe that the uncertainty is of the order of 10%. This type of analysis would be useful, for example, in situations where it is necessary to use a series of images acquired on different dates to monitor the evolution of the size of the lesion.

5 Theoretical analysis

In this section, we provide a theoretical study of the long-time behavior of PnP-ULA Algorithm 1 and PnP-SGD Algorithm 3, as well as quantitative bounds on the approximation of the posterior (10) introduced in Section 3. Let $q_y : \mathbb{R}^d \rightarrow (0, +\infty)$ be the likelihood function, *i.e.* $p(\cdot|x) \propto q_y$,

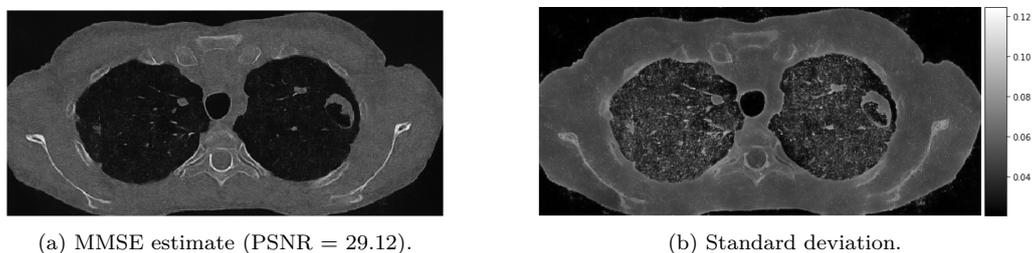


Figure 15: MMSE and standard deviation estimates for the denoising problem.

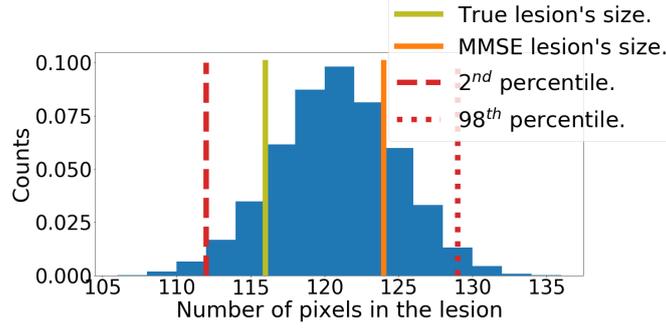


Figure 16: Histogram of the estimated lesion's size in the samples generated by PnP-ULA. The red lines correspond to the 2nd and 98th percentiles, the green line to the number of pixels in the lesion associated with the true image x and the orange line to the number of pixels constituting the lesion in the MMSE estimate.

and μ the prior distribution with density p with respect to the Lebesgue measure. For any $\varepsilon > 0$ we recall that p_ε is given by the Gaussian smoothing of p with level ε , for any $x \in \mathbb{R}^d$ by

$$p_\varepsilon(x) = (2\pi\varepsilon)^{-d/2} \int_{\mathbb{R}^d} \exp[-\|x - \tilde{x}\|^2 / (2\varepsilon)] p(\tilde{x}) d\tilde{x} .$$

Note that even in the case where μ does not admit a density with respect to the Lebesgue measure, p_ε is well-defined upon setting for any $x \in \mathbb{R}^d$

$$p_\varepsilon(x) = (2\pi\varepsilon)^{-d/2} \int_{\mathbb{R}^d} \exp[-\|x - \tilde{x}\|^2 / (2\varepsilon)] d\mu(\tilde{x}) .$$

One typical example of likelihood function that we consider in our numerical illustration, see Section 4, is $q_y(x) = \exp[-\|\mathbf{A}x - y\|^2 / (2\sigma^2)]$ for any $x \in \mathbb{R}^d$ with $\sigma > 0$ and $\mathbf{A} \in \mathbb{R}^{m \times d}$. Before turning to the analysis of the convergence of the introduced algorithms we state the following proposition which ensures the regularity of the posterior model w.r.t to the observation y .

Proposition 9. *Assume that there exist $\Phi_1 : \mathbb{R}^d \rightarrow [0, +\infty)$ and $\Phi_2 : \mathbb{R}^m \rightarrow [0, +\infty)$ such that for any $x \in \mathbb{R}^d$ and $y_1, y_2 \in \mathbb{R}^m$*

$$\|\log(q_{y_1}(x)) - \log(q_{y_2}(x))\| \leq (\Phi_1(x) + \Phi_2(y_1) + \Phi_2(y_2)) \|y_1 - y_2\| , \quad (18)$$

and for any $c > 0$, $\int_{\mathbb{R}^d} (1 + \Phi_1(\tilde{x})) \exp[c\Phi_1(\tilde{x})] p(\tilde{x}) d\tilde{x} < +\infty$. Then $y \mapsto \pi_y$ is locally Lipschitz w.r.t the total variation $\|\cdot\|_{TV}$, where for any $x \in \mathbb{R}^d, y \in \mathbb{R}^m$ we have

$$(d\pi_y / d\text{Leb})(x) = q_y(x)p(x) \Big/ \int_{\mathbb{R}^d} q_y(\tilde{x})p(\tilde{x})d\tilde{x} . \quad (19)$$

Proof. The proof is postponed to Appendix B. □

We highlight that Proposition 9 also holds if we replace p in (19) by p_ε for any $\varepsilon > 0$. Note that (18) is satisfied in the case where $q_y(x) = \exp[-\|\mathbf{A}x - y\|^2 / (2\sigma^2)]$ with $\Phi_1(x) = 2\|\mathbf{A}x\|/\sigma^2$ and $\Phi_2(y) = \|y\|/\sigma^2$.

We consider the following assumption on q_y and the prior p for some hyperparameter $\alpha > 0$ and an observation $y \in \mathbb{R}^m$.

H2. *The following hold:*

- (a) $\int_{\mathbb{R}^d} q_y(\tilde{x})p^\alpha(\tilde{x})d\tilde{x} < +\infty$ and for any $\varepsilon > 0$, $\int_{\mathbb{R}^d} q_y(\tilde{x})p_\varepsilon^\alpha(\tilde{x})d\tilde{x} < +\infty$.
(b) $\int_{\mathbb{R}^d} \|\tilde{x}\|^2 p(x)dx < +\infty$.

Note that if $\alpha = 1$, **H2-(a)** hold under **H1**, see Proposition 1. Under **H2-(a)**, define π the target probability distribution for any $x \in \mathbb{R}^d$ by

$$(d\pi/d\text{Leb})(x) = q_y(x)p^\alpha(x) \Big/ \int_{\mathbb{R}^d} q_y(\tilde{x})p^\alpha(\tilde{x})d\tilde{x} . \quad (20)$$

Note that for ease of notation, we do not explicitly highlight the dependency of the posterior distribution π with respect to the hyperparameter $\alpha > 0$, since it is fixed in the rest of this section. We also consider the family of probability distributions $\{\pi_\varepsilon : \varepsilon > 0\}$ given for any $\varepsilon > 0$ and $x \in \mathbb{R}^d$ by

$$(d\pi_\varepsilon/d\text{Leb})(x) = q_y(x)p_\varepsilon^\alpha(x) \Big/ \int_{\mathbb{R}^d} q_y(\tilde{x})p_\varepsilon^\alpha(\tilde{x})d\tilde{x} , \quad (21)$$

We divide our study into three parts. We first establish non-asymptotic bounds between the iterates of PnP-ULA and π_ε with respect to the total variation distance for any $\varepsilon > 0$, in Section 5.2. Second, we show the quantitative convergence of $(\pi_\varepsilon)_{\varepsilon>0}$ towards π with respect to the same metric as $\varepsilon \rightarrow 0$ in Section 5.3. Finally, in Section 5.4 the case of PnP-SGD is considered under the same set of assumptions as PnP-ULA. Before stating our results, additional notation is in order.

5.1 Notation

Denote by $\mathcal{B}(\mathbb{R}^d)$ the Borel σ -field of \mathbb{R}^d , and for $f : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable, $\|f\|_\infty = \sup_{x \in \mathbb{R}^d} |f(x)|$. For μ a probability measure on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ and f a μ -integrable function, denote by $\mu(f)$ the integral of f w.r.t. μ . For $f : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable, the V -norm of f is given by $\|f\|_V = \sup_{x \in \mathbb{R}^d} |f(x)|/V(x)$. Let ξ be a finite signed measure on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. The V -total variation distance of ξ is defined as

$$\|\xi\|_V = \sup_{\|f\|_V \leq 1} \left| \int_{\mathbb{R}^d} f(x)d\xi(x) \right| .$$

If $V = 1$, then $\|\cdot\|_V$ is the total variation denoted by $\|\cdot\|_{\text{TV}}$. Let \mathbf{U} be an open set of \mathbb{R}^d . For any pair of measurable spaces $(\mathbf{X}, \mathcal{X})$ and $(\mathbf{Y}, \mathcal{Y})$, measurable function $f : (\mathbf{X}, \mathcal{X}) \rightarrow (\mathbf{Y}, \mathcal{Y})$ and measure μ on $(\mathbf{X}, \mathcal{X})$ we denote by $f_{\#}\mu$ the pushforward measure of μ on $(\mathbf{Y}, \mathcal{Y})$ given for any $\mathbf{A} \in \mathcal{Y}$ by $f_{\#}\mu(\mathbf{A}) = \mu(f^{-1}(\mathbf{A}))$. We denote $\mathcal{P}(\mathbb{R}^d)$ the set of probability measures over $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ and for any $m \in \mathbb{N}$, $\mathcal{P}_m(\mathbb{R}^d) = \{\nu \in \mathcal{P}(\mathbb{R}^d) : \int_{\mathbb{R}^d} \|\tilde{x}\|^m d\tilde{x} < +\infty\}$.

We denote by $C^k(\mathbf{U}, \mathbb{R}^m)$ the set of \mathbb{R}^m -valued k -differentiable functions, respectively the set of compactly supported \mathbb{R}^m -valued and k -differentiable functions. Let $f : \mathbf{U} \rightarrow \mathbb{R}$, we denote by ∇f , the gradient of f if it exists. f is said to be m -convex with $m \geq 0$ if for all $x, y \in \mathbb{R}^d$ and $t \in [0, 1]$,

$$f(tx + (1-t)y) \leq tf(x) + (1-t)f(y) - mt(1-t)\|x-y\|^2/2 .$$

For any $a \in \mathbb{R}^d$ and $R > 0$, denote $B(a, R)$ the open ball centered at a with radius R . Let $(\mathbf{X}, \mathcal{X})$ and $(\mathbf{Y}, \mathcal{Y})$ be two measurable spaces. A Markov kernel \mathbf{P} is a mapping $\mathbf{K} : \mathbf{X} \times \mathcal{Y} \rightarrow [0, 1]$ such that for any $x \in \mathbf{X}$, $\mathbf{P}(x, \cdot)$ is a probability measure and for any $\mathbf{A} \in \mathcal{Y}$, $\mathbf{P}(\cdot, \mathbf{A})$ is measurable. For any probability measure μ on $(\mathbf{X}, \mathcal{X})$ and measurable function $f : \mathbf{Y} \rightarrow \mathbb{R}_+$ we denote $\mu\mathbf{P} = \int_{\mathbf{X}} \mathbf{P}(x, \cdot) d\mu(x)$ and $\mathbf{P}f = \int_{\mathbf{Y}} f(y)\mathbf{P}(\cdot, dy)$. In what follows the Dirac mass at $x \in \mathbb{R}^d$ is

denoted by δ_x . For any $x \in \mathbb{R}^d$, we denote $\tau_x : \mathbb{R}^d \rightarrow \mathbb{R}^d$ the translation operator given for any $\tilde{x} \in \mathbb{R}^d$ by $\tau_x(\tilde{x}) = \tilde{x} - x$. The complement of a set $A \subset \mathbb{R}^d$, is denoted by A^c . All densities are w.r.t. the Lebesgue measure unless stated otherwise.

Finally, for any compact set $C \subset \mathbb{R}^d$ we consider $\mathcal{K}_C = \{K : K \text{ is compact and } K \subset C\}$. (\mathcal{K}_C, d_C) is a metric space where the Hausdorff distance d_C is given for any $K_1, K_2 \in \mathcal{K}_C$ by

$$d_C(K_1, K_2) = \inf\{\varepsilon > 0 : K_1 \subset K_2 + \overline{B}(0, \varepsilon), K_2 \subset K_1 + \overline{B}(0, \varepsilon)\}, \quad (22)$$

where for any pair of sets $A, B \subset \mathbb{R}^d$ we have $A + B = \{x + y : x \in A, y \in B\}$.

5.2 Convergence of PnP-ULA

In this section, we fix $\varepsilon > 0$ and derive quantitative bounds between the iterates of PnP-ULA and π_ε with respect to the total variation. To address this issue, we first show that PnP-ULA is geometrically ergodic and establish non-asymptotic bounds between the corresponding Markov kernel and its invariant distribution. Second, we analyse the distance between this stationary distribution and π_ε for the same metric. In Section 5.2.1, we derive our convergence results in the case where $\log(q_y)$ is differentiable with Lipschitz gradient and strongly concave. In Section 5.2.2, we extend our results to the case where $\log(q_y)$ has only Lipschitz gradient and satisfies a one-sided Lipschitz condition. Finally in Section 5.2.3 we deal with the case of PPnP-ULA.

For any $\varepsilon > 0$ we define $g_\varepsilon : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, +\infty)$ for any $x_1, x_2 \in \mathbb{R}^d$ by

$$g_\varepsilon(x_1|x_2) = p(x_1) \exp[-\|x_2 - x_1\|^2 / (2\varepsilon)] \bigg/ \int_{\mathbb{R}^d} p(\tilde{x}) \exp[-\|x_2 - \tilde{x}\|^2 / (2\varepsilon)] d\tilde{x} . \quad (23)$$

Note that $g(\cdot|X_\varepsilon)$ is the density with respect to the Lebesgue measure of the distribution of X given X_ε , where X is sampled according to the prior distribution μ and $X_\varepsilon = X + \varepsilon^{1/2}Z$ where Z is a Gaussian random variable with zero mean and identity covariance matrix.

Throughout, this section, we consider the following assumption on the family of denoising operators $\{D_\varepsilon : \varepsilon > 0\}$ which will ensure that PnP-ULA approximately targets π_ε . The same assumption will be considered when studying PnP-SGD to approximate the maximum a posteriori estimator associated with π_ε .

H3 (R). *There exist $\varepsilon_0 > 0$, $M_R \geq 0$ and $L \geq 0$ such that for any $\varepsilon \in (0, \varepsilon_0]$, $x_1, x_2 \in \mathbb{R}^d$ and $x \in \overline{B}(0, R)$ we have*

$$\|(\text{Id} - D_\varepsilon)(x_1) - (\text{Id} - D_\varepsilon)(x_2)\| \leq L \|x_1 - x_2\| , \quad \|D_\varepsilon(x) - D_\varepsilon^*(x)\| \leq M_R , \quad (24)$$

where

$$D_\varepsilon^*(x_1) = \int_{\mathbb{R}^d} \tilde{x} g_\varepsilon(\tilde{x}|x_1) d\tilde{x} . \quad (25)$$

The first condition in (24) regarding the smoothness property of the denoiser can be explicitly verified for a certain class of neural networks by adding a spectral regularization term for each layer of the neural network, see [74, 63].

Regarding the second condition in (24), small values of $M_R \geq 0$ for any $R > 0$ can be empirically enforced as follows. Consider a neural network $f_w : \mathbb{R}^d \rightarrow \mathbb{R}^d$, parameterized by its weights and bias gathered in $w \in \mathcal{W}$ where \mathcal{W} is some measurable space. For any $\varepsilon > 0$, let $\ell_\varepsilon : \mathcal{W} \rightarrow [0, +\infty)$ given for any $w \in \mathcal{W}$ by $\ell_\varepsilon(w) = \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - f_w(x_\varepsilon)\|^2 p_\varepsilon(x_\varepsilon) g_\varepsilon(x|x_\varepsilon) dx_\varepsilon dx$. Note that such a loss is considered in the Noise2Noise network introduced in [57].

We have that for any $w \in \mathcal{W}$, $\ell_\varepsilon(w) \geq \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - D_\varepsilon^*(x_\varepsilon)\|^2 p_\varepsilon(x_\varepsilon) g_\varepsilon(x|x_\varepsilon) dx dx_\varepsilon = \ell_\varepsilon^*$, since for any $x_\varepsilon \in \mathbb{R}^d$, $D_\varepsilon^*(x_\varepsilon) = \int_{\mathbb{R}^d} \tilde{x} g_\varepsilon(\tilde{x}|x_\varepsilon) d\tilde{x}$, see (25). Consider $w^* \in \mathcal{W}$ obtained after numerically minimizing ℓ_ε and satisfying $\ell_\varepsilon(w^*) \leq \ell_\varepsilon^* + \eta$ with $\eta > 0$. In this case, the following result ensures that (24) is satisfied with M_R of order $\eta^{1/(2d+2)}$ for any $R > 0$ and letting $D_\varepsilon = f_{w^*}$.

Proposition 10. Assume that for any $w \in \mathcal{W}$

$$\int_{\mathbb{R}^d} (\|x\|^2 + \|f_w(x_\varepsilon)\|^2) p_\varepsilon(x_\varepsilon) g_\varepsilon(x|x_\varepsilon) dx dx_\varepsilon < +\infty. \quad (26)$$

Let $R, \eta > 0$ and $w^* \in \mathcal{W}$ such that $\ell_\varepsilon(w^*) \leq \ell_\varepsilon^* + \eta$. In addition, assume that

$$\sup_{x_1, x_2 \in \bar{\mathbb{B}}(0, 2R)} \left\{ \|x_2 - x_1\|^{-1} (\|f_{w^*}(x_2) - f_{w^*}(x_1)\| + \|D_\varepsilon^*(x_2) - D_\varepsilon^*(x_1)\|) \right\} < +\infty,$$

where D_ε^* is given in (25). Then there exists $C_R, \bar{\eta}_R \geq 0$ such that if $\eta \in (0, \bar{\eta}_R]$ then for any $\tilde{x} \in \bar{\mathbb{B}}(0, R)$, $\|f_{w^*}(\tilde{x}) - D_\varepsilon^*(\tilde{x})\| \leq C_R \eta^{1/(2d+2)}$.

Proof. The proof is postponed to Appendix C.1. \square

Note that (26) is satisfied if for any $w \in \mathcal{W}$, $\sup_{x \in \mathbb{R}^d} \|f_w(x)\| (1 + \|x\|)^{-1} < +\infty$ and **H2** hold.

The second condition in (24) can also be enforced using universal properties of neural network. Indeed, if there exists $\tilde{M}_R \geq 0$ such that we have $\inf_{w \in \mathcal{W}} \sup_{x \in \bar{\mathbb{B}}(0, R)} \tilde{M}_R^{-1} \|f_w(x) - D_\varepsilon^*(x)\| \leq 1$ then the second condition in (24) holds upon letting $D_\varepsilon = f_{w^*}$ for an appropriate choice of weights $w^* \in \mathcal{W}$. This last inequality can be established using universal approximation theorems such as the one given [7, Section 4.7].

We recall that PnP-ULA is given by the following recursion: for any $k \in \mathbb{N}$

$$\begin{aligned} X_{k+1} &= X_k + \delta b_\varepsilon(X_k) + \sqrt{2\delta} Z_{k+1}, \\ b_\varepsilon(x) &= \nabla \log(q_y(x)) + \alpha P_\varepsilon(x) + \text{prox}_\lambda(\iota_C)(x), \quad P_\varepsilon(x) = (D_\varepsilon(x) - x)/\varepsilon, \end{aligned} \quad (27)$$

where $\delta > 0$ is a step-size, $\alpha, \varepsilon, \lambda > 0$ are hyperparameters of the algorithm, $C \subset \mathbb{R}^d$ is a closed convex set, $\{Z_k : k \in \mathbb{N}\}$ a family of i.i.d. Gaussian random variables with zero mean and identity covariance matrix and $\text{prox}_\lambda(\iota_C)$ the proximal operator of ι_C with step-size λ , see [8, Definition 12.23], where ι_C is the convex indicator of C defined for $x \in \mathbb{R}^d$ by $\iota_C = +\infty$ if $x \notin C$ and 0 if $x \in C$. Note that for any $x \in \mathbb{R}^d$ we have $\text{prox}_\lambda(\iota_C)(x) = (x - \Pi_C(x))/\lambda$, where Π_C is the projection onto C .

In what follows, for any $\delta > 0$ and $C \subset \mathbb{R}^d$ closed and convex, we denote by $R_{\varepsilon, \delta} : \mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d) \rightarrow [0, 1]$ the Markov kernel associated with the recursion (39) and given for any $x \in \mathbb{R}^d$ and $A \in \mathcal{B}(\mathbb{R}^d)$ by

$$R_{\varepsilon, \delta}(x, A) = (2\pi)^{-d/2} \int_{\mathbb{R}^d} \mathbf{1}_A(x + \delta b_\varepsilon(x) + \sqrt{2\delta}z) \exp[-\|z\|^2/2] dz. \quad (28)$$

Note that for ease of notation, we do not explicitly highlight the dependency of $R_{\varepsilon, \delta}$ and b_ε with respect to the hyperparameter $\alpha, \lambda > 0$ and C .

5.2.1 Strongly log-concave case

First, we consider the case where $x \mapsto \log(q_y)$ is strongly log-concave. In this setting, we set $C = \mathbb{R}^d$ in (39) and (28). We recall that in our image processing applications, we have that for any $x \in \mathbb{R}^d$, $q_y(x) = \exp[-\|\mathbf{A}x - y\|^2/(2\sigma^2)]$ and that q_y is strongly log-concave if and only if \mathbf{A} is invertible. This is the case for denoising tasks where $\mathbf{A} = \text{Id}$ and for deblurring tasks with convolution kernels which have full Fourier support.

We start with the following result which ensures that the Markov chain (39) is geometrically ergodic under **H3** for the Wasserstein metric \mathbf{W}_1 and in V -norm for $V : \mathbb{R}^d \rightarrow [1, +\infty)$ given for any $x \in \mathbb{R}^d$ by

$$V(x) = 1 + \|x\|^2. \quad (29)$$

Proposition 11. Assume **H1**, **H2** and **H3**(R) for some $R > 0$. Let $\alpha > 0$ and $\varepsilon \in (0, \varepsilon_0]$. If there exists $\mathfrak{m} > 0$ such that $\log(q_y)$ is \mathfrak{m} -concave with $\mathfrak{m} \geq 2\alpha L/\varepsilon$ then there exist $A_1 \geq 0$ and $\rho_1 \in [0, 1)$ such that for any $\delta \in (0, \bar{\delta}]$, $x_1, x_2 \in \mathbb{R}^d$ and $k \in \mathbb{N}$ we have

$$\begin{aligned} \|\delta_{x_1} R_{\varepsilon, \delta}^k - \delta_{x_2} R_{\varepsilon, \delta}^k\|_V &\leq A_1 \rho_1^{k\delta} (V^2(x_1) + V^2(x_2)) , \\ \mathbf{W}_1(\delta_{x_1} R_{\varepsilon, \delta}^k, \delta_{x_2} R_{\varepsilon, \delta}^k) &\leq A_1 \rho_1^{k\delta} \|x_1 - x_2\| , \end{aligned}$$

where V is given in (29) and $\bar{\delta} = \mathfrak{m}(L_y + \alpha L/\varepsilon)^{-2}/2$.

Proof. The proof is postponed to Appendix C.2. \square

The constants A_1 and ρ_1 do not depend on the dimension d but only on the parameters $\mathfrak{m}, L, L_y, \alpha, \varepsilon$. Note that a similar result can be established for \mathbf{W}_p for any $p \in \mathbb{N}^*$ instead of \mathbf{W}_1 . Under the conditions of Proposition 11 we have for any $\nu_1, \nu_2 \in \mathcal{P}_1(\mathbb{R}^d)$

$$\begin{aligned} \|\nu_1 R_{\varepsilon, \delta}^k - \nu_2 R_{\varepsilon, \delta}^k\|_V &\leq A_1 \rho_1^{k\delta} \left(\int_{\mathbb{R}^d} V^2(\tilde{x}) d\nu_1(\tilde{x}) + \int_{\mathbb{R}^d} V^2(\tilde{x}) d\nu_2(\tilde{x}) \right) , \\ \mathbf{W}_1(\nu_1 R_{\varepsilon, \delta}^k, \nu_2 R_{\varepsilon, \delta}^k) &\leq A_1 \rho_1^{k\delta} \left(\int_{\mathbb{R}^d} \|\tilde{x}\| d\nu_1(\tilde{x}) + \int_{\mathbb{R}^d} \|\tilde{x}\| d\nu_2(\tilde{x}) \right) . \end{aligned} \quad (30)$$

First, $(\mathcal{P}_1(\mathbb{R}^d), \mathbf{W}_1)$ is a complete metric space [85, Theorem 6.18]. Second, for any $\delta \in (0, \bar{\delta}]$, there exists $\mathfrak{m} \in \mathbb{N}^*$ such that $f^{\mathfrak{m}}$ is contractive with $f : \mathcal{P}_1(\mathbb{R}^d) \rightarrow \mathcal{P}_1(\mathbb{R}^d)$ given for any $\nu \in \mathcal{P}_1(\mathbb{R}^d)$ by $f(\nu) = \nu R_{\varepsilon, \delta}$ using Proposition 11. Therefore we can apply the Picard fixed point theorem and we obtain that $R_{\varepsilon, \delta}$ admits an invariant probability measure $\pi_{\varepsilon, \delta} \in \mathcal{P}_1(\mathbb{R}^d)$.

Therefore, since $\pi_{\varepsilon, \delta}$ is an invariant probability measure for $R_{\varepsilon, \delta}$ and $\pi_{\varepsilon, \delta} \in \mathcal{P}_1(\mathbb{R}^d)$, using (30), we have for any $\nu \in \mathcal{P}_1(\mathbb{R}^d)$

$$\begin{aligned} \|\nu R_{\varepsilon, \delta}^k - \pi_{\varepsilon, \delta}\|_V &\leq A_1 \rho_1^{k\delta} \left(\int_{\mathbb{R}^d} V^2(\tilde{x}) d\nu(\tilde{x}) + \int_{\mathbb{R}^d} V^2(\tilde{x}) d\pi_{\varepsilon, \delta}(\tilde{x}) \right) , \\ \mathbf{W}_1(\nu R_{\varepsilon, \delta}^k, \pi_{\varepsilon, \delta}) &\leq A_1 \rho_1^{k\delta} \left(\int_{\mathbb{R}^d} \|\tilde{x}\| d\nu(\tilde{x}) + \int_{\mathbb{R}^d} \|\tilde{x}\| d\pi_{\varepsilon, \delta}(\tilde{x}) \right) . \end{aligned}$$

Combining this result with the fact that for any $t \geq 0$, $(1 - e^{-t})^{-1} \leq 1 + t^{-1}$, we get that for any $n \in \mathbb{N}^*$ and $h : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable such that $\sup_{x \in \mathbb{R}^d} \{(1 + \|x\|^2)^{-1} |h(x)|\} < +\infty$

$$\left| n^{-1} \sum_{k=1}^n \mathbb{E}[h(X_k)] - \pi_{\varepsilon, \delta}[h] \right| \leq A_1 (\bar{\delta} + \log^{-1}(1/\rho_1)) \left(V^2(x) + \int_{\mathbb{R}^d} V^2(\tilde{x}) d\pi_{\varepsilon, \delta}(\tilde{x}) \right) / (n\delta) ,$$

where $(X_k)_{k \in \mathbb{N}}$ is the Markov chain given by (39) with starting point $X_0 = x$.

The following assumption is necessary to ensure that $x \mapsto \log(p_\varepsilon(x))$ has Lipschitz gradients, see Proposition 2.

H4. For any $\varepsilon > 0$, there exists $K_\varepsilon \geq 0$ such that for any $x \in \mathbb{R}^d$,

$$\int_{\mathbb{R}^d} \left\| \tilde{x} - \int_{\mathbb{R}^d} \tilde{x}' g_\varepsilon(\tilde{x}'|x) d\tilde{x}' \right\|^2 g_\varepsilon(\tilde{x}|x) d\tilde{x} \leq K_\varepsilon ,$$

with g_ε given in (23).

We emphasize that **H4** is not needed to establish the convergence of the Markov chain. However, we impose it in order to compare the stationary distribution of PnP-ULA with the target distribution π_ε . Depending on the prior distribution density p , **H4** may be checked by hand. Finally, note that **H4** can be extended to cover the case where the prior distribution μ does not admit a density with respect to the Lebesgue measure.

In the following proposition, we show that we can control the distance between $\pi_{\varepsilon,\delta}$ and π_ε based on the previous observations.

Proposition 12. *Assume **H1**, **H2**, **H3**(R) for some $R > 0$ and **H4**. Moreover, let $\alpha > 0$, $\varepsilon \in (0, \varepsilon_0]$ and assume that $\int_{\mathbb{R}^d} (1 + \|\tilde{x}\|^4) p_\varepsilon^\alpha(\tilde{x}) d\tilde{x} < +\infty$. In addition, if there exists $\mathfrak{m} > 0$ such that $\log(q_y)$ is \mathfrak{m} -concave with $\mathfrak{m} \geq (2\alpha/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon)$ and $\bar{\delta} = \mathfrak{m}(L_y + \alpha L/\varepsilon)^{-2}/2$, then for any $\delta \in (0, \bar{\delta}]$, $R_{\varepsilon,\delta}$ admits an invariant probability measure $\pi_{\varepsilon,\delta}$ and there exists $B_1 \geq 0$ such that for any $\delta \in (0, \bar{\delta}]$*

$$\|\pi_{\varepsilon,\delta} - \pi_\varepsilon\|_V \leq B_1(\delta^{1/2} + M_R + \exp[-R]), \quad (31)$$

where V is given in (29) and B_1 does not depend on R .

Proof. The proof is postponed to Appendix C.3. \square

The bound appearing in (31) depends on an extra hyperparameter $R > 0$ which may be optimized if **H3**(R) holds for any $R > 0$ and $\{M_R : R > 0\}$ can be expressed in a closed form. In particular if there exists $M \in (0, 1)$ such that for any $R > 0$, $M_R = M \times R$ then there exists $B_1 \geq 0$ such that for any $\delta \in (0, \bar{\delta}]$ and $R > 0$

$$\|\pi_{\varepsilon,\delta} - \pi_\varepsilon\|_V \leq B_1(\delta^{1/2} + M \log(1/M)),$$

by setting $R = \log(1/M)$. Similarly if there exists $M > 0$ such that for any $R > 0$, $M_R = M$ then there exists $B_1 \geq 0$ such that for any $\delta \in (0, \bar{\delta}]$ and $R > 0$

$$\|\pi_{\varepsilon,\delta} - \pi_\varepsilon\|_V \leq B_1(\delta^{1/2} + M),$$

by letting $R \rightarrow +\infty$.

We now combine Proposition 11 and Proposition 12 in order to control the bias of the Monte Carlo estimator obtained using PnP-ULA.

Proposition 13. *Assume **H1**, **H2**, **H3**(R) for some $R > 0$ and **H4**. Moreover, let $\alpha > 0$, $\varepsilon \in (0, \varepsilon_0]$ and assume that $\int_{\mathbb{R}^d} (1 + \|\tilde{x}\|^4) p_\varepsilon^\alpha(\tilde{x}) d\tilde{x} < +\infty$. In addition, if there exists $\mathfrak{m} > 0$ such that $\log(q_y)$ is \mathfrak{m} -concave with $\mathfrak{m} \geq (2\alpha/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon)$ and $\bar{\delta} = \mathfrak{m}(L_y + \alpha L/\varepsilon)^{-2}/2$, then there exists $C_{1,\varepsilon} \geq 0$ such that for any $h : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable with $\sup_{x \in \mathbb{R}^d} \{|h(x)| (1 + \|x\|^2)^{-1}\} \leq 1$, $n \in \mathbb{N}^*$, $\delta \in (0, \bar{\delta}]$ we have*

$$\left| n^{-1} \sum_{k=1}^n \mathbb{E}[h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) d\pi_\varepsilon(\tilde{x}) \right| \leq C_{1,\varepsilon}(\delta^{1/2} + M_R + \exp[-R] + (n\delta)^{-1})(1 + \|x\|^4).$$

Proof. The proof is straightforward upon combining Proposition 11 and Proposition 12. \square

In particular, applying Proposition 13 to the family $\{h_i\}_{i=1}^d$ where for any $i \in \{1, \dots, d\}$, $h_i(x) = x_i$ we get that

$$\left\| n^{-1} \sum_{k=1}^n \mathbb{E}[X_k] - \int_{\mathbb{R}^d} \tilde{x} d\pi_\varepsilon(\tilde{x}) \right\| \leq C_{1,\varepsilon}(\delta^{1/2} + M_R + \exp[-R] + (n\delta)^{-1})(1 + \|x\|^4),$$

and $n^{-1} \sum_{k=1}^n X_k$ is an approximation of the MMSE given by $\int_{\mathbb{R}^d} \tilde{x} d\pi_\varepsilon(\tilde{x})$.

5.2.2 General case

We now consider the general case where $\log(q_y)$ is no longer \mathfrak{m} -concave but satisfies a more general one-sided Lipschitz condition, *i.e.* we consider the following condition.

H5. *There exists $\mathfrak{m} \in \mathbb{R}$ such that for any $x_1, x_2 \in \mathbb{R}^d$ we have*

$$\langle \nabla \log(q_y)(x_2) - \nabla \log(q_y)(x_1), x_2 - x_1 \rangle \leq -\mathfrak{m} \|x_2 - x_1\|^2 .$$

Note that if **H5** is satisfied with $\mathfrak{m} > 0$ then $\log(q_y)$ is \mathfrak{m} -concave. Assume **H1** then **H5** holds for $\mathfrak{m} = -L_y$. However, it is possible that $\mathfrak{m} > -L_y$ which leads to better convergence rates for PnP-ULA. As a result even when **H1** holds we still consider **H5**. In order to deal with **H5** we no longer consider $\mathbb{C} = \mathbb{R}^d$ in (39) and (28), but instead $\mathbb{C} \subset \mathbb{R}^d$ is now some convex compact set fixed by the user.

In our image processing applications, we recall that for any $x \in \mathbb{R}^d$ we have, $q_y(x) = \exp[-\|\mathbf{A}x - y\|^2 / (2\sigma^2)]$. If \mathbf{A} is not invertible then $\log(q_y)$ is not \mathfrak{m} -concave with $\mathfrak{m} > 0$. This is the case, in our deblurring experiment if the convolution kernel does not have full support in the Fourier domain. In our inpainting application $\mathbf{A} = \text{diag}(a)$, where $a \in \{0, 1\}^d$, and for any $i \in \{1, \dots, d\}$, $a(i) = 0$ if and only if the i -th pixel is not visible (note that here we have considered a flattened version of the image). As a result, \mathbf{A} is not invertible even if only one pixel is hidden. Note that in these cases, **H5** holds with $\mathfrak{m} = 0$ since q_y is log-concave.

In what follows, we prove similar results to Section 5.2.1 in this general setting. We start by stating the counterpart of Proposition 11.

Proposition 14. *Assume **H1**, **H2**, **H3**(R) for some $R > 0$ and **H5**. Let $\alpha, \lambda > 0$, $\varepsilon \in (0, \varepsilon_0]$ such that $2\lambda(L_y + \alpha L/\varepsilon - \min(\mathfrak{m}, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$. Then for any $\mathbb{C} \subset \mathbb{R}^d$ convex and compact with $0 \in \mathbb{C}$, there exist $A_{1,\mathbb{C}} \geq 0$ and $\rho_{1,\mathbb{C}} \in [0, 1)$ such that for any $\delta \in (0, \bar{\delta}]$, $x_1, x_2 \in \mathbb{R}^d$ and $k \in \mathbb{N}$ we have*

$$\begin{aligned} \|\delta_{x_1} R_{\varepsilon,\delta}^k - \delta_{x_2} R_{\varepsilon,\delta}^k\|_V &\leq A_{1,\mathbb{C}} \rho_{1,\mathbb{C}}^{k\delta} (V^2(x_1) + V^2(x_2)) , \\ \mathbf{W}_1(\delta_{x_1} R_{\varepsilon,\delta}^k, \delta_{x_2} R_{\varepsilon,\delta}^k) &\leq A_{1,\mathbb{C}} \rho_{1,\mathbb{C}}^{k\delta} \|x_1 - x_2\| , \end{aligned}$$

where V is given in (29).

Proof. The proof is postponed to Appendix C.2. □

The discussion following Proposition 11 is still valid in this case. In particular, Proposition 14 implies that $R_{\varepsilon,\delta}$ admits an invariant probability measure $\pi_{\varepsilon,\delta}$. In the following proposition, we control the distance in V -norm between $\pi_{\varepsilon,\delta}$ and π_ε . This result is the counterpart of Proposition 12. Note that contrary to Proposition 12, we must take into account another source of error stemming from the fact that we have to incorporate the additional term $x \mapsto \text{prox}_\lambda(\iota_{\mathbb{C}})$ into (39) which ensures the geometric ergodicity of PnP-ULA in this general setting.

Proposition 15. *Assume **H1**, **H2**, **H3**(R) for some $R > 0$, **H4** and **H5**. Moreover, let $\alpha > 0$, $\varepsilon \in (0, \varepsilon_0]$ and assume that $\int_{\mathbb{R}^d} (1 + \|\tilde{x}\|^4) p_\varepsilon^\alpha(\tilde{x}) d\tilde{x} < +\infty$. Let $\lambda > 0$ such that $2\lambda(L_y + (\alpha/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon) - \min(\mathfrak{m}, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$. Then for any $\delta \in (0, \bar{\delta}]$ and \mathbb{C} convex and compact with $0 \in \mathbb{C}$, $R_{\varepsilon,\delta}$ admits an invariant probability measure $\pi_{\varepsilon,\delta}$. In addition, there exists $B_0 \geq 0$ such that for any \mathbb{C} convex compact with $\bar{\mathbb{B}}(0, R_{\mathbb{C}}) \subset \mathbb{C}$ and $R_{\mathbb{C}} > 0$, there exists $B_{1,\mathbb{C}} \geq 0$ such that for any $\delta \in (0, \bar{\delta}]$*

$$\|\pi_{\varepsilon,\delta} - \pi_\varepsilon\|_V \leq B_0 R_{\mathbb{C}}^{-1} + B_{1,\mathbb{C}} (\delta^{1/2} + M_R + \exp[-R]) ,$$

where V is given in (29).

Proof. The proof is postponed to Appendix C.3. \square

We now combine Proposition 14 and Proposition 15 in order to control the bias of the Monte Carlo estimator obtained using PnP-ULA. Note that the following proposition is the counterpart of Proposition 13.

Proposition 16. *Assume H2, H3(R) for some $R > 0$, H4 and H5. Moreover, let $\alpha > 0$, $\varepsilon \in (0, \varepsilon_0]$ and assume that $\int_{\mathbb{R}^d} (1 + \|\tilde{x}\|^4) p_\varepsilon^\alpha(\tilde{x}) d\tilde{x} < +\infty$. Let $\lambda > 0$ such that $2\lambda(L_y + (\alpha/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon) - \min(m, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$. Then there exists $C_{1,\varepsilon} > 0$ such that for for any C convex compact with $\bar{B}(0, R_C) \subset C$ and $R_C > 0$ there exists $C_{2,\varepsilon}$ such that for any $h : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable with $\sup_{x \in \mathbb{R}^d} \{ |h(x)| (1 + \|x\|^2)^{-1} \} \leq 1$, $n \in \mathbb{N}^*$, $\delta \in (0, \bar{\delta}]$ we have*

$$\left| n^{-1} \sum_{k=1}^n \mathbb{E} [h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) d\pi_\varepsilon(\tilde{x}) \right| \leq \left\{ C_{1,\varepsilon} R_C^{-1} + C_{2,\varepsilon, C} (\delta^{1/2} + M_R + \exp[-R] + (n\delta)^{-1}) \right\} (1 + \|x\|^4) .$$

Proof. The proof is straightforward combining Proposition 14 and Proposition 15. \square

5.2.3 Convergence guarantees for PPnP-ULA

In our experiments Section 4 we notice that the bound on the step-size $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$ is conservative. Indeed, the MMSE obtained with PnP-ULA with $\delta \leq \bar{\delta}$ and $\delta = 8\bar{\delta}$ are similar. (however instability is reached for $15\bar{\delta}$). This suggests that our convergence results could be tightened and that PnP-ULA is stable for larger step-sizes. In addition, in practice we observe that the projection constraint is never activated. More precisely, let $(X_n)_{n \in \mathbb{N}}$ be given by PnP-ULA. In all of our experiments $\sum_{k=1}^N \mathbf{1}_C(X_k)/N \geq 1$. In this section, we propose an alternative projected algorithm which is equal to PnP-ULA up to the first exit time $T = \inf\{n \in \mathbb{N} : X_n \notin C\}$. We call this algorithm Projected *Plug & Play* Unadjusted Langevin Algorithm (PPnP-ULA). It is given by the following recursion: $X_0 \in C$ and for any $k \in \mathbb{N}$

$$\begin{aligned} X_{k+1} &= \Pi_C(X_k + \delta b_\varepsilon(X_k) + \sqrt{2\delta} Z_{k+1}) , \\ b_\varepsilon(x) &= \nabla \log(q_y(x)) + \alpha P_\varepsilon(x) , \quad P_\varepsilon(x) = (D_\varepsilon(x) - x)/\varepsilon , \end{aligned} \tag{32}$$

where $\delta > 0$ is a step-size, $\alpha, \varepsilon, \lambda > 0$ are hyperparameters of the algorithm, $C \subset \mathbb{R}^d$ is a closed convex set, $\{Z_k : k \in \mathbb{N}\}$ a family of i.i.d. Gaussian random variables with zero mean and identity covariance matrix and where Π_C is the projection onto C . In what follows, for any $\delta > 0$ and $C \subset \mathbb{R}^d$ closed and convex, we denote by $Q_{\varepsilon, \delta} : \mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d) \rightarrow [0, 1]$ the Markov kernel associated with the recursion (32) and given for any $x \in \mathbb{R}^d$ and $A \in \mathcal{B}(\mathbb{R}^d)$ by

$$Q_{\varepsilon, \delta}(x, A) = (2\pi)^{-d/2} \int_{\mathbb{R}^d} \mathbf{1}_{\Pi_C^{-1}(A)}(x + \delta b_\varepsilon(x) + \sqrt{2\delta} z) \exp[-\|z\|^2/2] dz .$$

Note that for ease of notation, we do not explicitly highlight the dependency of $Q_{\varepsilon, \delta}$ and b_ε with respect to the hyperparameter $\alpha, \lambda > 0$ and C .

First, we have the following result which ensures that PPnP-ULA is geometrically ergodic for all step-sizes.

Proposition 17. *Assume H1, H2, H3(R) for some $R > 0$. Let $\alpha, \lambda, \varepsilon, \bar{\delta} > 0$. Then for any $C \subset \mathbb{R}^d$ convex and compact with $0 \in C$, there exist $\tilde{A}_C \geq 0$ and $\tilde{\rho}_C \in [0, 1)$ such that for any $\delta \in (0, \bar{\delta}]$, $x_1, x_2 \in C$ and $k \in \mathbb{N}$ we have*

$$\|\delta_{x_1} Q_{\varepsilon, \delta}^k - \delta_{x_2} Q_{\varepsilon, \delta}^k\|_{\text{TV}} \leq \tilde{A}_C \tilde{\rho}_C^{k\delta} .$$

Proof. The proof is postponed to Appendix C.4. \square

In particular $Q_{\varepsilon, \delta}$ admits an invariant probability measure $\pi_{\varepsilon, \delta}^C$. The next proposition ensures that for small enough step-size δ the invariant measures of PnP-ULA and PPnP-ULA are close if the compact convex set C has a large diameter.

Proposition 18. *Assume **H1**, **H2**, **H3**(R) for some $R > 0$ and **H5**. In addition, assume that there exists $\tilde{m}, c > 0$ such that for $C = \mathbb{R}^d$ and for any $\varepsilon > 0$ and $x \in \mathbb{R}^d$, $\langle b_\varepsilon(x), x \rangle \leq -\tilde{m} \|x\|^2 + c$. Let $\alpha, \lambda > 0$, $\varepsilon \in (0, \varepsilon_0]$ such that $2\lambda(L_y + \alpha L/\varepsilon - \min(m, 0)) \leq 1$. Then there exist $\bar{A} \geq 0$ and $\eta, \bar{\delta} > 0$ such that for any $C \subset \mathbb{R}^d$ convex and compact with $0 \in C$ and $\bar{B}(0, R_C/2) \subset C \subset \bar{B}(0, R_C)$ and $\delta \in (0, \bar{\delta}]$ we have*

$$\|\pi_{\varepsilon, \delta} - \pi_{\varepsilon, \delta}^C\|_{\text{TV}} \leq \bar{A} \exp[-\eta R_C],$$

where $\pi_{\varepsilon, \delta}$ is the invariant measure of $R_{\varepsilon, \delta}$ and $\pi_{\varepsilon, \delta}^C$ is the invariant measure of $Q_{\varepsilon, \delta}$.

Proof. The proof is postponed to Appendix C.5. \square

5.3 Posterior approximation

We consider the following general regularity assumption.

H6 (α). *There exist $\kappa \geq 0$, $\beta > 0$ and $q : \mathbb{R}^d \rightarrow (0, +\infty)$ such that $\int_{\mathbb{R}^d} q(\tilde{x}) d\tilde{x} = 1$, $\|q\|_\infty < +\infty$ and for almost every $x \in \mathbb{R}^d$, $\int_{\mathbb{R}^d} |p(\tilde{x}) - p(x - \tilde{x})| q^{\min(1-1/\alpha, 0)}(\tilde{x}) d\tilde{x} \leq e^{\kappa(1+\|x\|^2)} \|x\|^\beta$.*

In the case where $\alpha \geq 1$, **H6**(α) is equivalent to the following assumption: there exist $\kappa \geq 0$ and $\beta > 0$ such that for almost every $x \in \mathbb{R}^d$, $\|\mu - (\tau_x)_\# \mu\|_{\text{TV}} \leq e^{\kappa(1+\|x\|^2)} \|x\|^\beta$, where we recall that μ is the probability distribution with density with respect to the Lebesgue measure proportional to p and that for any $\tilde{x} \in \mathbb{R}^d$, $\tau_x(\tilde{x}) = \tilde{x} - x$. Note that since $p \in L^1(\mathbb{R}^d)$ we have $\lim_{x \rightarrow 0} \|\mu - (\tau_x)_\# \mu\|_{\text{TV}} = 0$. In **H6**(α) for $\alpha < 1$ we assume more regularity for $x \mapsto (\tau_x)_\# \mu$ in total variation in order to obtain explicit bounds between π_ε and π .

In the following proposition we provide easy-to-check conditions on the density of the prior distribution μ so that **H6**(α) holds.

Proposition 19. *Assume that there exists $U : \mathbb{R}^d \rightarrow \mathbb{R}$ such that for any $x \in \mathbb{R}^d$, $p(x) = e^{-U(x)} / \int_{\mathbb{R}^d} e^{-U(\tilde{x})} d\tilde{x}$. Assume that U is γ -Hölder, i.e. there exists $C_\gamma > 0$ such that for any $x_1, x_2 \in \mathbb{R}^d$, i.e. $\|U(x_1) - U(x_2)\| \leq C_\gamma \|x_1 - x_2\|^\gamma$. Then **H6**(α) is satisfied for $\alpha \geq 1$. In addition, assume that $\gamma \leq 2$ and that there exist $c_1, \varpi > 0$ and $c_2 \in \mathbb{R}$ such that for any $x \in \mathbb{R}^d$, $U(x) \geq c_1 \|x\|^\varpi + c_2$ then **H6**(α) holds for any $\alpha > 0$.*

Proof. The proof is postponed to Appendix D.1. \square

Under **H6**(α) we establish the following result which ensures that π_ε is close to π in total variation for small values of ε .

Proposition 20. *Assume **H1** then the following hold:*

- (a) *If $\alpha = 1$, then $\lim_{\varepsilon \rightarrow 0} \|\pi_\varepsilon - \pi\|_{\text{TV}} = 0$.*
- (b) *Assume that $\|p\|_\infty < +\infty$ then for any $\alpha \geq 1$, $\lim_{\varepsilon \rightarrow 0} \|\pi_\varepsilon - \pi\|_{\text{TV}} = 0$.*
- (c) *Assume that $\|p\|_\infty < +\infty$ and **H6**(α) then there exist $\varepsilon_1 > 0$ and $A_0 \geq 0$ such that for any $\varepsilon \in (0, \varepsilon_1]$ we have $\|\pi_\varepsilon - \pi\|_{\text{TV}} \leq A_0 \varepsilon^{\beta \min(\alpha, 1)/2}$.*

Proof. The proof is postponed to Appendix D.2. \square

Note that a related result in the case where $p(x) = e^{-U(x)} / \int_{\mathbb{R}^d} e^{-U(\tilde{x})} d\tilde{x}$ with U Lipschitz continuous and $\alpha = 1$ can be found in [86, Corollary 1] with explicit dependency with respect to the dimension d . However, note that Proposition 20 differs from [86, Corollary 1] since the Gaussian smoothing approximation is applied to the prior distribution and the estimate is given on the posterior distribution in Proposition 20, whereas in [86, Corollary 1] the Gaussian smoothing approximation is applied to the posterior distribution and the estimate is given on the posterior distribution as well.

The following proposition is an extension of Proposition 13 and Proposition 16. The main difference is that the approximation is expressed with respect to the true posterior π and not π_ε for some value $\varepsilon > 0$. Let $\varepsilon_1 > 0$ be given by Proposition 20.

Proposition 21. *Assume **H1**, **H2**, **H3**, **H4** and **H5**. Let $\alpha > 0$ and assume that for any $\varepsilon \in (0, \min(\varepsilon_0, \varepsilon_1)]$, $\int_{\mathbb{R}^d} (1 + \|\tilde{x}\|^4)(p_\varepsilon^\alpha + p^\alpha)(\tilde{x}) d\tilde{x} < +\infty$ and **H6**(α). Then there exists $C_0 \geq 0$ such that for any $\varepsilon > 0$ and $\lambda > 0$ such that $2\lambda(L_y + (\alpha/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon) - \min(m, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$, there exists $C_{1,\varepsilon} \geq 0$ such that for any C convex compact with $\bar{B}(0, R_C) \subset C$ and $R_C > 0$, there exists $C_{2,\varepsilon,C} \geq 0$ such that for any $h : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable with $\sup_{x \in \mathbb{R}^d} \{|h(x)| (1 + \|x\|^2)^{-1}\} \leq 1$, $n \in \mathbb{N}^*$, $\delta \in (0, \bar{\delta}]$ and $R > 0$ we have*

$$\begin{aligned} & \left| n^{-1} \sum_{k=1}^n \mathbb{E} [h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) d\pi(\tilde{x}) \right| \\ & \leq \left\{ C_0 \varepsilon^{\beta \min(\alpha, 1)/4} + C_{1,\varepsilon} R_C^{-1} + C_{2,\varepsilon,C} (\delta^{1/2} + M_R + \exp[-R] + (n\delta)^{-1}) \right\} (1 + \|x\|^4). \end{aligned}$$

In addition, if there exists $m > 0$ such that $\log(q_y)$ is m -concave with $m \geq 2(\alpha/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon)$ and $\bar{\delta} = m(L_y + \alpha L/\varepsilon)^{-2}/2$, then there exists $C_{1,\varepsilon} \geq 0$ such that for any $h : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable with $\sup_{x \in \mathbb{R}^d} \{|h(x)| (1 + \|x\|^2)^{-1}\} \leq 1$, $n \in \mathbb{N}^$, $\delta \in (0, \bar{\delta}]$ and $R > 0$ we have*

$$\begin{aligned} & \left| n^{-1} \sum_{k=1}^n \mathbb{E} [h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) d\pi(\tilde{x}) \right| \\ & \leq C_0 \varepsilon^{\beta \min(\alpha, 1)/4} + C_{1,\varepsilon} (\delta^{1/2} + M_R + \exp[-R] + (n\delta)^{-1}) (1 + \|x\|^4). \end{aligned}$$

Proof. In the general case where $\log(q_y)$ is not assumed to be m -concave with $m > 0$, the proof is completed upon combining Proposition 16, Proposition 20 and the fact that for any probability distribution ν_1, ν_2 , $\|\nu_1 - \nu_2\|_V \leq \|\nu_1 - \nu_2\|_{TV}^{1/2} (\nu_1[V^2] + \nu_2[V^2])^{1/2}$. The proof is similar in the case where $\log(q_y)$ is m -concave upon replacing Proposition 16 by Proposition 13. \square

5.4 Convergence of PnP-SGD

We now turn to the proof of convergence of PnP-SGD. We recall that PnP-SGD is given by the following recursion: $X_0 \in \mathbb{R}^d$ and for any $k \in \mathbb{N}$

$$\begin{aligned} X_{k+1} &= X_k + \delta_k (b_\varepsilon(X_k) + Z_{k+1}), \\ b_\varepsilon(x) &= \nabla \log(q_y(x)) + \alpha(D_\varepsilon(x) - x)/\varepsilon, \end{aligned}$$

where $(\delta_k)_{k \in \mathbb{N}} \in (\mathbb{R}_+)^{\mathbb{N}}$ is a family of step-sizes, $\alpha, \varepsilon > 0$ are hyperparameters of the algorithm and $\{Z_k : k \in \mathbb{N}\}$ a family of i.i.d. Gaussian random variables with zero mean and identity covariance matrix. We recall that the sequences $(X_k)_{k \in \mathbb{N}}$ and $(Z_k)_{k \in \mathbb{N}}$ are defined on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

Note that contrary to (39), we do not consider an additive term of the form $x \mapsto (x - \Pi_C(x))/\lambda$ which ensures the stability of the numerical scheme. As a result, our asymptotic estimates are only valid for sequences which remain in a compact set K , which is a classical assumption in stochastic approximation [80, 33, 32, 61]. Under tighter conditions on $x \mapsto \log((d\pi_\varepsilon/d\text{Leb})(x))$ this limitation can be circumvented using the global asymptotic results of [80, Theorem A1.1]. We leave this analysis for future work.

In what follows, we show that the asymptotic bias of PnP-SGD is controlled by $\{\mathsf{M}_R : R > 0\}$ using recent

Proposition 22. *Let $\alpha, \varepsilon > 0$. Assume **H1**, **H3**, that $\log(q_y)$ is real-analytic and that $\lim_{k \rightarrow +\infty} \delta_k = 0$, $\sum_{k \in \mathbb{N}} \delta_k = +\infty$ and $\sum_{k \in \mathbb{N}} \delta_k^2 < +\infty$. Let $R > 0$, $\mathsf{K} \subset \overline{\mathbb{B}}(0, R)$ be a compact set, $X_0 \in \mathbb{R}^d$ and $\mathsf{A}_{\varepsilon, \mathsf{K}} \in \mathcal{F}$ given by*

$$\mathsf{A}_{\varepsilon, \mathsf{K}} = \{\omega \in \Omega : \text{there exists } k_0 \in \mathbb{N} \text{ such that for any } k \geq k_0, X_k(\omega) \in \mathsf{K}\}.$$

Then there exist $C_{\varepsilon, \mathsf{K}} \geq 0$ and $r_{\varepsilon, \mathsf{K}} \in (0, 1)$ such that $\limsup_{k \rightarrow +\infty} d(X_k(\omega), \mathsf{S}_{\varepsilon, \mathsf{K}}) \leq C_{\varepsilon, \mathsf{K}} \mathsf{M}_R^{r_{\varepsilon, \mathsf{K}}}$ for any $\omega \in \mathsf{A}_{\varepsilon, \mathsf{K}}$, with $\mathsf{S}_{\varepsilon, \mathsf{K}} = \{x \in \mathsf{K} : \nabla \log(d\pi_\varepsilon/d\text{Leb})(x) = 0\}$, where $(d\pi_\varepsilon/d\text{Leb})$ is given in (21).

Proof. The proof is postponed to Appendix E.1. \square

Finally we analyse the behavior of the sets $(\mathsf{S}_{\varepsilon, \mathsf{K}})_{\varepsilon > 0}$ in the case where $\varepsilon \rightarrow 0$. More precisely, we show that each cluster point (in the sense of the Hausdorff distance) of the sequences of sets $(\mathsf{S}_{\varepsilon_n, \mathsf{K}})_{n \in \mathbb{N}}$ with $\lim_{n \rightarrow +\infty} \varepsilon_n = 0$ is contained in the set of stationary points of $\log(d\pi/d\text{Leb})$.

Denote E the set of sequences $(\varepsilon_n)_{n \in \mathbb{N}}$ such that for any $n \in \mathbb{N}$, $\varepsilon_n > 0$ and $\lim_{n \rightarrow +\infty} \varepsilon_n = 0$. For any sequence $e = (\varepsilon_n)_{n \in \mathbb{N}} \in \mathsf{E}$ denote $\mathsf{T}_{e, \mathsf{K}}$ the set of cluster points of $(\mathsf{S}_{\varepsilon_n, \mathsf{K}})_{n \in \mathbb{N}}$ with respect to the Hausdorff distance on K , d_{K} defined in (22). We also define $\mathsf{T}_{\mathsf{K}} = \bigcup_{e \in \mathsf{E}} \mathsf{T}_{e, \mathsf{K}}$, i.e. the collection of the cluster points for sequences $(\varepsilon_n)_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow +\infty} \varepsilon_n = 0$. Finally, we define $\mathsf{S}_{\mathsf{K}}^* = \bigcup_{\mathsf{S} \in \mathsf{T}_{\mathsf{K}}} \mathsf{S}$ the union of all the cluster points.

In what follows, we show that $\mathsf{S}_{\mathsf{K}}^*$ is a subset of the stationary points of $\log(d\pi/d\text{Leb})$.

Proposition 23. *Assume **H1**, **H3** and that there exists $U : \mathbb{R}^d \rightarrow \mathbb{R}_+$ such that for any $x \in \mathbb{R}^d$, $p(x) = e^{-U(x)} / \int_{\mathbb{R}^d} e^{-U(\tilde{x})} d\tilde{x}$. Assume that $U \in \mathcal{C}^1(\mathbb{R}^d, \mathbb{R})$ and that there exists $\mathsf{N} \geq 0$ such that for any $x_1, x_2 \in \mathbb{R}^d$*

$$|U(x_1) - U(x_2)| \leq \mathsf{N} \|x_1 - x_2\|, \quad \|\nabla U(x_1) - \nabla U(x_2)\| \leq \mathsf{N} \|x_1 - x_2\|.$$

Then for any compact set K , $\mathsf{S}_{\mathsf{K}}^ \subset \mathsf{S}_{\mathsf{K}}$ with $\mathsf{S}_{\mathsf{K}} = \{x \in \mathsf{K} : \nabla \log(d\pi/d\text{Leb})(x) = 0\}$ with $(d\pi/d\text{Leb})$ given in (20).*

Proof. The proof is postponed to Appendix E.2. \square

6 Conclusion

This paper presented theory, methods, and computation algorithms for performing Bayesian inference with *Plug & Play* priors. This mathematical and computational framework is rooted in the Bayesian *M-complete* paradigm and adopts the view that *Plug & Play* models approximate a regularised oracle model. We established clear conditions ensuring that the involved models and quantities of interest are well defined and well posed. Following on from this, we studied three Bayesian computation algorithms related to biased approximations of a Langevin diffusion process, for which we provide detailed convergence guarantees under easily verifiable and realistic

conditions. For example, our theory does not require the denoising algorithms representing the prior to be gradient or proximal operators. We also studied the estimation error involved in using these algorithms and models instead of the oracle model, which is decision-theoretically optimal but intractable. To the best of our knowledge, this is the first Bayesian *Plug & Play* framework with this level of insight and guarantees on the delivered solutions. We illustrated the proposed framework with three Bayesian image restoration experiments - deblurring, inpainting, and denoising - where we computed point estimates as well as uncertainty visualisation and quantification analyses.

In future work, we would like to continue our theoretical and empirical investigation of Bayesian *Plug & Play* models, methods and algorithms. From a modelling viewpoint, it would be interesting to consider other neural network based priors such as the generative ones used in [14] or the autoencoder-based priors in [44], as well as to generalise the Gaussian smoothing to other smoothings and investigate their properties in the context of Bayesian inverse problems. With regards to Bayesian analysis, it would be important to investigate the frequentist accuracy of *Plug & Play* models, as well as the adoption of robust Bayesian techniques in order to perform inference directly w.r.t. to the oracle model [89]. From a Bayesian computation viewpoint, a priority is to develop accelerated computation algorithms similar to [67]. Lastly, with regards to experimental work, we intend to study the application of this framework to uncertainty quantification problems, e.g., in the context of medical imaging.

References

- [1] *Introduction to markov random fields*, in Markov Random Fields for Vision and Image Processing, The MIT Press, 2011, <https://doi.org/10.7551/mitpress/8579.003.0001>, <https://doi.org/10.7551/mitpress/8579.003.0001>.
- [2] C. AGUERREBERE, A. ALMANSA, J. DELON, Y. GOUSSEAU, AND P. MUSE, *A Bayesian Hyperprior Approach for Joint Image Denoising and Interpolation, With an Application to HDR Imaging*, IEEE Transactions on Computational Imaging, 3 (2017), pp. 633–646, <https://doi.org/10.1109/TCI.2017.2704439>, https://nounsse.github.io/HBE_project/, <https://arxiv.org/abs/1706.03261>.
- [3] G. ALAIN AND Y. BENGIO, *What Regularized Auto-Encoders Learn from the Data-Generating Distribution*, Journal of Machine Learning Research, 15 (2014), pp. 3743–3773, <http://jmlr.org/papers/v15/alain14a.html>, <https://arxiv.org/abs/1211.4246>.
- [4] Y. ALTMANN, M. PEREYRA, AND J. BIOCAS-DIAS, *Collaborative sparse regression using spatially correlated supports - application to hyperspectral unmixing*, IEEE Transactions on Image Processing, 24 (2015), pp. 5800–5811, <https://doi.org/10.1109/TIP.2015.2487862>.
- [5] V. ANTUN, M. J. COLBROOK, AND A. C. HANSEN, *Can stable and accurate neural networks be computed?—on the barriers of deep learning and smale’s 18th problem*, arXiv preprint arXiv:2101.08286, (2021).
- [6] V. ANTUN, F. RENNA, C. POON, B. ADCKOCK, AND A. C. HANSEN, *On instabilities of deep learning in image reconstruction and the potential costs of ai*, Proceedings of the National Academy of Sciences, 117 (2020), pp. 30088–30095.
- [7] F. BACH, *Breaking the curse of dimensionality with convex neural networks*, The Journal of Machine Learning Research, 18 (2017), pp. 629–681.

- [8] H. H. BAUSCHKE, P. L. COMBETTES, ET AL., *Convex analysis and monotone operator theory in Hilbert spaces*, vol. 408, Springer, 2011.
- [9] J. BERNARDO AND A. SMITH, *Bayesian Theory*, vol. 15, 01 2000, <https://doi.org/10.2307/2983298>.
- [10] L. BIEGLER, G. BIROS, O. GHATTAS, M. HEINKENSCHLOSS, D. KEYES, B. MALLICK, Y. MARZOUK, L. TENORIO, B. VAN BLOEMEN WAANDERS, AND K. WILLCOX, eds., *Large-scale inverse problems and quantification of uncertainty*, Wiley Series in Computational Statistics, John Wiley & Sons, Ltd., Chichester, 2011.
- [11] S. A. BIGDELI, M. JIN, P. FAVARO, AND M. ZWICKER, *Deep Mean-Shift Priors for Image Restoration*, in (NIPS) Advances in Neural Information Processing Systems 30, sep 2017, pp. 763–772, <http://papers.nips.cc/paper/6678-deep-mean-shift-priors-for-image-restoration>, <https://arxiv.org/abs/1709.03749>.
- [12] S. A. BIGDELI AND M. ZWICKER, *Image Restoration using Autoencoding Priors*, tech. report, 2017, <https://arxiv.org/abs/1703.09964>.
- [13] V. I. BOGACHEV, *Measure Theory*, vol. Volume 1, Springer, 1 ed., 2007, <http://gen.lib.rus.ec/book/index.php?md5=ffbdc7e3d8e571c6cd5f9e8633cdfdc2>.
- [14] A. BORA, A. JALAL, E. PRICE, AND A. G. DIMAKIS, *Compressed sensing using generative models*, in (ICML) International Conference on Machine Learning, vol. 2, JMLR. org, 2017, pp. 537–546, <https://arxiv.org/abs/arXiv:1703.03208v1>.
- [15] V. D. BORTOLI AND A. DURMUS, *Convergence of diffusions and their discretizations: from continuous to discrete processes and back*, 2019, <https://arxiv.org/abs/1904.09808>.
- [16] V. D. BORTOLI AND A. DURMUS, *Convergence of diffusions and their discretizations: from continuous to discrete processes and back*, 2020, <https://arxiv.org/abs/1904.09808>.
- [17] V. D. BORTOLI, A. DURMUS, A. F. VIDAL, AND M. PEREYRA, *Maximum likelihood estimation of regularisation parameters in high-dimensional inverse problems: an empirical bayesian approach. part ii: Theoretical analysis*, 2020, <https://arxiv.org/abs/2008.05793>.
- [18] L. BOTTOU, F. E. CURTIS, AND J. NOCEDAL, *Optimization methods for large-scale machine learning*, Siam Review, 60 (2018), pp. 223–311.
- [19] S. BOUCHERON, G. LUGOSI, AND P. MASSART, *Concentration inequalities: A nonasymptotic theory of independence*, Oxford university press, 2013.
- [20] S. BOYD, N. PARIKH, AND E. CHU, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Now Publishers Inc, 2011.
- [21] O. BRANDIÈRE AND M. DUFLO, *Les algorithmes stochastiques contournent-ils les pièges?*, Ann. Inst. H. Poincaré Probab. Statist., 32 (1996), pp. 395–427, http://www.numdam.org/item?id=AIHPB_1996__32_3_395_0.
- [22] A. BUADES, B. COLL, AND J.-M. MOREL, *A non-local algorithm for image denoising*, in 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), vol. 2, IEEE, 2005, pp. 60–65.

- [23] A. BUADES, B. COLL, AND J.-M. MOREL, *A review of image denoising algorithms, with a new one*, Multiscale Modeling & Simulation, 4 (2005), pp. 490–530.
- [24] X. CAI, M. PEREYRA, AND J. D. MCEWEN, *Uncertainty quantification for radio interferometric imaging – i. proximal MCMC methods*, Monthly Notices of the Royal Astronomical Society, 480 (2018), pp. 4154–4169, <https://doi.org/10.1093/mnras/sty2004>, <https://doi.org/10.1093/mnras/sty2004>.
- [25] A. CHAMBOLLE, *An algorithm for total variation minimization and applications*, Journal of Mathematical Imaging and Vision, 20 (2004), pp. 89–97, <https://doi.org/10.1023/B:JMIV.0000011325.36760.1e>.
- [26] S. H. CHAN, X. WANG, AND O. A. ELGENDY, *Plug-and-play admm for image restoration: Fixed-point convergence and applications*, IEEE Transactions on Computational Imaging, 3 (2017), pp. 84–98, <https://doi.org/10.1109/TCI.2016.2629286>, <https://arxiv.org/abs/1605.01710>.
- [27] Y. CHEN AND T. POCK, *Trainable Nonlinear Reaction Diffusion: A Flexible Framework for Fast and Effective Image Restoration*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 39 (2017), pp. 1256–1272, <https://doi.org/10.1109/TPAMI.2016.2596743>, <https://arxiv.org/abs/1508.02848>.
- [28] X. CHENG, N. S. CHATTERJI, P. L. BARTLETT, AND M. I. JORDAN, *Underdamped langevin mcmc: A non-asymptotic analysis*, in Conference on Learning Theory, PMLR, 2018, pp. 300–323.
- [29] K. DABOV, A. FOI, V. KATKOVNIK, AND K. EGIAZARIAN, *Image denoising with block-matching and 3d filtering*, in Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning, vol. 6064, International Society for Optics and Photonics, 2006, p. 606414.
- [30] A. S. DALALYAN, *Theoretical guarantees for approximate sampling from smooth and log-concave densities*, J. R. Stat. Soc. Ser. B. Stat. Methodol., 79 (2017), pp. 651–676, <https://doi.org/10.1111/rssb.12183>, <https://doi.org/10.1111/rssb.12183>.
- [31] V. DE BORTOLI, A. DURMUS, M. PEREYRA, AND A. F. VIDAL, *Maximum likelihood estimation of regularization parameters in high-dimensional inverse problems: An empirical bayesian approach. part ii: Theoretical analysis*, SIAM Journal on Imaging Sciences, 13 (2020), pp. 1990–2028.
- [32] B. DELYON, *General results on the convergence of stochastic algorithms*, IEEE Transactions on Automatic Control, 41 (1996), pp. 1245–1255.
- [33] B. DELYON, M. LAVIELLE, AND E. MOULINES, *Convergence of a stochastic approximation version of the em algorithm*, Annals of statistics, (1999), pp. 94–128.
- [34] S. DIAMOND, V. SITZMANN, F. HEIDE, AND G. WETZSTEIN, *Unrolled optimization with deep priors*, (2017), <https://arxiv.org/abs/1705.08041>.
- [35] C. DONG, C. C. LOY, K. HE, AND X. TANG, *Learning a deep convolutional network for image super-resolution*, in European conference on computer vision, Springer, 2014, pp. 184–199.
- [36] R. DOUC, E. MOULINES, P. PRIOURET, AND P. SOULIER, *Markov Chains*, Springer, 2019.

- [37] A. DURMUS AND E. MOULINES, *Nonasymptotic convergence analysis for the unadjusted Langevin algorithm*, Ann. Appl. Probab., 27 (2017), pp. 1551–1587, <https://doi.org/10.1214/16-AAP1238>, <https://doi.org/10.1214/16-AAP1238>.
- [38] A. DURMUS, E. MOULINES, AND M. PEREYRA, *Efficient bayesian computation by proximal markov chain monte carlo: when langevin meets moreau*, SIAM Journal on Imaging Sciences, 11 (2018), pp. 473–506.
- [39] B. EFRON, *Tweedie’s formula and selection bias*, Journal of the American Statistical Association, 106 (2011), pp. 1602–1614.
- [40] H. GAO, X. TAO, X. SHEN, AND J. JIA, *Dynamic scene deblurring with parameter selective sharing and nested skip connections*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 3848–3856.
- [41] M. GHARBI, G. CHAURASIA, S. PARIS, AND F. DURAND, *Deep joint demosaicking and denoising*, ACM Transactions on Graphics (TOG), 35 (2016), p. 191.
- [42] D. GILTON, G. ONGIE, AND R. WILLETT, *Neumann networks for inverse problems in imaging*, (2019), <https://arxiv.org/abs/1901.03707>.
- [43] M. GIROLAMI AND B. CALDERHEAD, *Riemann manifold langevin and hamiltonian monte carlo methods*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 73 (2011), pp. 123–214.
- [44] M. GONZÁLEZ, M. DELBRACIO, P. MUSÉ, A. ALMANSA, P. TAN, M. DELBRACIO, P. MUSÉ, AND P. TAN, *Solving Inverse Problems by Joint Posterior Maximization with a VAE Prior*, arXiv, (2019), <https://arxiv.org/abs/1911.06379>.
- [45] K. GREGOR AND Y. LECUN, *Learning fast approximations of sparse coding*, in Proceedings of the 27th International Conference on International Conference on Machine Learning, Omnipress, 2010, pp. 399–406.
- [46] R. GRIBONVAL, *Should penalized least squares regression be interpreted as maximum a posteriori estimation?*, IEEE Transactions on Signal Processing, 59 (2011), pp. 2405–2410.
- [47] B. GUO, Y. HAN, AND J. WEN, *Agem: Solving linear inverse problems via deep priors and sampling*, in Advances in Neural Information Processing Systems, 2019, pp. 547–558.
- [48] A. HOUDARD, C. BOUVEYRON, AND J. DELON, *High-dimensional mixture models for unsupervised image denoising (hdmi)*, SIAM Journal on Imaging Sciences, 11 (2018), pp. 2815–2846.
- [49] Z. KADKHOAIE AND E. P. SIMONCELLI, *Solving linear inverse problems using the prior implicit in a denoiser*, arXiv preprint arXiv:2007.13640, (2020).
- [50] U. S. KAMILOV, H. MANSOUR, AND B. WOHLBERG, *A plug-and-play priors approach for solving nonlinear imaging inverse problems*, IEEE Signal Processing Letters, 24 (2017), pp. 1872–1876.
- [51] I. KARATZAS AND S. E. SHREVE, *Brownian motion and stochastic calculus*, vol. 113 of Graduate Texts in Mathematics, Springer-Verlag, New York, second ed., 1991, <https://doi.org/10.1007/978-1-4612-0949-2>, <https://doi.org/10.1007/978-1-4612-0949-2>.

- [52] B. KAWAR, G. VAKSMAN, AND M. ELAD, *Stochastic image denoising by sampling from the posterior distribution*, 2021, <https://arxiv.org/abs/2101.09552>.
- [53] D. P. KINGMA AND J. BA, *Adam: A method for stochastic optimization*, arXiv preprint arXiv:1412.6980, (2014).
- [54] S. KULLBACK, *Information theory and statistics*, Courier Corporation, 1997.
- [55] J. LATZ, *On the well-posedness of bayesian inverse problems*, SIAM/ASA Journal on Uncertainty Quantification, 8 (2020), pp. 451–482, <https://doi.org/10.1137/19M1247176>, <https://doi.org/10.1137/19M1247176>, <https://arxiv.org/abs/https://doi.org/10.1137/19M1247176>.
- [56] M. LEBRUN, A. BUADES, AND J.-M. MOREL, *A nonlocal bayesian image denoising algorithm*, SIAM Journal on Imaging Sciences, 6 (2013), pp. 1665–1688.
- [57] J. LEHTINEN, J. MUNKBERG, J. HASSELGREN, S. LAINE, T. KARRAS, M. AITTALA, AND T. AILA, *Noise2noise: Learning image restoration without clean data*, arXiv preprint arXiv:1803.04189, (2018).
- [58] R. S. LIPTSER AND A. N. SHIRYAEV, *Statistics of random processes. I*, vol. 5 of Applications of Mathematics (New York), Springer-Verlag, Berlin, expanded ed., 2001. General theory, Translated from the 1974 Russian original by A. B. Aries, Stochastic Modelling and Applied Probability.
- [59] C. LOUCHET AND L. MOISAN, *Posterior expectation of the total variation model: Properties and experiments*, SIAM Journal on Imaging Sciences, 6 (2013), pp. 2640–2684, <https://doi.org/10.1137/120902276>.
- [60] T. MEINHARDT, M. MOLLER, C. HAZIRBAS, AND D. CREMERS, *Learning proximal operators: Using denoising networks for regularizing inverse imaging problems*, in (ICCV) International Conference on Computer Vision, 2017, pp. 1781–1790, <https://doi.org/10.1109/ICCV.2017.198>, http://openaccess.thecvf.com/content_iccv_2017/html/Meinhardt_Learning_Proximal_Operators_ICCV_2017_paper.html.
- [61] M. METIVIER AND P. PRIOURET, *Applications of a kushner and clark lemma to general classes of stochastic algorithms*, IEEE Transactions on Information Theory, 30 (1984), pp. 140–151.
- [62] S. P. MEYN AND R. L. TWEEDIE, *Stability of Markovian processes. III. Foster-Lyapunov criteria for continuous-time processes*, Adv. in Appl. Probab., 25 (1993), pp. 518–548, <https://doi.org/10.2307/1427522>, <https://doi.org/10.2307/1427522>.
- [63] T. MIYATO, T. KATAOKA, M. KOYAMA, AND Y. YOSHIDA, *Spectral normalization for generative adversarial networks*, arXiv preprint arXiv:1802.05957, (2018).
- [64] Y. NESTEROV, *Lectures on convex optimization*, vol. 137 of Springer Optimization and Its Applications, Springer, Cham, 2018, <https://doi.org/10.1007/978-3-319-91578-4>, <https://doi.org/10.1007/978-3-319-91578-4>. Second edition of [MR2142598].
- [65] M. PEREYRA, *Proximal Markov chain Monte Carlo algorithms*, Statistics and Computing, 26 (2016), pp. 745–760, <https://doi.org/10.1007/s11222-015-9567-4>, <http://dx.doi.org/10.1007/s11222-015-9567-4>, <https://arxiv.org/abs/1306.0187>.

- [66] M. PEREYRA, P. SCHNITER, E. CHOUZENOUX, J.-C. PESQUET, J.-Y. TOURNERET, A. O. HERO, AND S. MCLAUGHLIN, *A survey of stochastic simulation and optimization methods in signal processing*, IEEE Journal of Selected Topics in Signal Processing, 10 (2015), pp. 224–241.
- [67] M. PEREYRA, L. VARGAS MIELES, AND K. C. ZYGALAKIS, *Accelerating proximal Markov chain Monte Carlo by using an explicit stabilized method*, SIAM J. Imaging Sci., 13 (2020), pp. 905–935, <https://doi.org/10.1137/19M1283719>, <https://doi.org/10.1137/19M1283719>.
- [68] E. T. REEHORST AND P. SCHNITER, *Regularization by denoising: Clarifications and new interpretations*, IEEE Transactions on Computational Imaging, 5 (2018), pp. 52–67, <https://doi.org/10.1109/TCI.2018.2880326>, <https://arxiv.org/abs/1806.02296>.
- [69] A. REPETTI, M. PEREYRA, AND Y. WIAUX, *Scalable bayesian uncertainty quantification in imaging inverse problems via convex optimization*, SIAM Journal on Imaging Sciences, 12 (2019), pp. 87–118, <https://doi.org/10.1137/18M1173629>.
- [70] C. ROBERT, *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*, Springer Texts in Statistics, Springer New York, 2007, <https://books.google.fr/books?id=6oQ4s8Pq9pYC>.
- [71] G. O. ROBERTS, R. L. TWEEDIE, ET AL., *Exponential convergence of langevin distributions and their discrete approximations*, Bernoulli, 2 (1996), pp. 341–363.
- [72] Y. ROMANO, M. ELAD, AND P. MILANFAR, *The little engine that could: Regularization by denoising (red)*, SIAM Journal on Imaging Sciences, 10 (2017), pp. 1804–1844.
- [73] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D: Nonlinear Phenomena, 60 (1992), pp. 259–268, [https://doi.org/10.1016/0167-2789\(92\)90242-F](https://doi.org/10.1016/0167-2789(92)90242-F).
- [74] E. K. RYU, J. LIU, S. WANG, X. CHEN, Z. WANG, AND W. YIN, *Plug-and-play methods provably converge with properly trained denoisers*, in Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, 2019, pp. 5546–5557, <http://proceedings.mlr.press/v97/ryu19a.html>, <https://arxiv.org/abs/1905.05406>.
- [75] E. SCHWARTZ, R. GIRYES, AND A. M. BRONSTEIN, *Deepisp: Toward learning an end-to-end image processing pipeline*, IEEE Transactions on Image Processing, 28 (2018), pp. 912–923.
- [76] L. SCHWARTZ, *Désintégration d’une mesure*, Séminaire Équations aux dérivées partielles (Polytechnique), pp. 1–10.
- [77] S. SREEHARI, S. V. VENKATAKRISHNAN, B. WOHLBERG, G. T. BUZZARD, L. F. DRUMMY, J. P. SIMMONS, AND C. A. BOUMAN, *Plug-and-Play Priors for Bright Field Electron Tomography and Sparse Interpolation*, IEEE Transactions on Computational Imaging, 2 (2016), pp. 1–1, <https://doi.org/10.1109/TCI.2016.2599778>, <https://arxiv.org/abs/1512.07331>.
- [78] A. M. STUART, *Inverse problems: A bayesian perspective*, Acta Numerica, 19 (2010), p. 451–559, <https://doi.org/10.1017/S0962492910000061>.

- [79] Y. SUN, Z. WU, B. WOHLBERG, AND U. S. KAMILOV, *Scalable plug-and-play admm with convergence guarantees*, arXiv preprint arXiv:2006.03224, (2020).
- [80] V. B. TADIĆ, A. DOUCET, ET AL., *Asymptotic bias of stochastic gradient search*, The Annals of Applied Probability, 27 (2017), pp. 3255–3304.
- [81] A. M. TEODORO, J. M. BIOUCAS-DIAS, AND M. A. FIGUEIREDO, *A convergent image fusion algorithm using scene-adapted gaussian-mixture-based denoising*, IEEE Transactions on Image Processing, 28 (2018), pp. 451–463.
- [82] A. M. TEODORO, J. M. BIOUCAS-DIAS, AND M. A. T. FIGUEIREDO, *Scene-Adapted Plug-and-Play Algorithm with Guaranteed Convergence: Applications to Data Fusion in Imaging*, jan 2018, <https://arxiv.org/abs/1801.00605>.
- [83] S. V. VENKATAKRISHNAN, C. A. BOUMAN, AND B. WOHLBERG, *Plug-and-play priors for model based reconstruction*, in 2013 IEEE Global Conference on Signal and Information Processing, IEEE, 2013, pp. 945–948.
- [84] A. F. VIDAL, V. DE BORTOLI, M. PEREYRA, AND A. DURMUS, *Maximum likelihood estimation of regularization parameters in high-dimensional inverse problems: An empirical bayesian approach part i: Methodology and experiments*, SIAM Journal on Imaging Sciences, 13 (2020), pp. 1945–1989.
- [85] C. VILLANI, *Optimal transport*, vol. 338 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], Springer-Verlag, Berlin, 2009, <https://doi.org/10.1007/978-3-540-71050-9>, <https://doi.org/10.1007/978-3-540-71050-9>. Old and new.
- [86] M. VONO, N. DOBIGEON, AND P. CHAINAIS, *Asymptotically exact data augmentation: models, properties and algorithms*, arXiv preprint arXiv:1902.05754, (2019).
- [87] Z. WANG AND A. C. BOVIK, *Mean squared error: Love it or leave it? a new look at signal fidelity measures*, IEEE Signal Processing Magazine, 26 (2009), pp. 98–117, <https://doi.org/10.1109/MSP.2008.930649>.
- [88] Z. WANG, A. C. BOVIK, H. R. SHEIKH, AND E. P. SIMONCELLI, *Image quality assessment: from error visibility to structural similarity*, IEEE transactions on image processing, 13 (2004), pp. 600–612.
- [89] J. WATSON AND C. HOLMES, *Approximate models and robust decisions*, Statistical Science, 31 (2016), pp. 465–489, <https://doi.org/10.1214/16-sts592>, <https://doi.org/10.1214/16-sts592>.
- [90] X. XU, Y. SUN, J. LIU, B. WOHLBERG, AND U. S. KAMILOV, *Provable Convergence of Plug-and-Play Priors with MMSE denoisers*, (2020), pp. 1–10, <http://arxiv.org/abs/2005.07685>, <https://arxiv.org/abs/2005.07685>.
- [91] K. YAN, X. WANG, L. LU, AND R. M. SUMMERS, *Deeplesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning*, Journal of Medical Imaging, 5 (2018), p. 036501.
- [92] G. YU, G. SAPIRO, AND S. MALLAT, *Solving inverse problems with piecewise linear estimators: From gaussian mixture models to structured sparsity*, IEEE Transactions on Image Processing, 21 (2011), pp. 2481–2499.

- [93] K. ZHANG, W. ZUO, Y. CHEN, D. MENG, AND L. ZHANG, *Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising*, IEEE Transactions on Image Processing, 26 (2017), pp. 3142–3155.
- [94] K. ZHANG, W. ZUO, S. GU, AND L. ZHANG, *Learning Deep CNN Denoiser Prior for Image Restoration*, in (CVPR) IEEE Conference on Computer Vision and Pattern Recognition, IEEE, apr 2017, pp. 2808–2817, <https://doi.org/10.1109/CVPR.2017.300>, http://openaccess.thecvf.com/content_cvpr_2017/html/Zhang_Learning_Deep_CNN_CVPR_2017_paper.html, <https://arxiv.org/abs/1704.03264>.
- [95] K. ZHANG, W. ZUO, AND L. ZHANG, *Ffdnet: Toward a fast and flexible solution for cnn-based image denoising*, IEEE Transactions on Image Processing, 27 (2018), pp. 4608–4622.
- [96] D. ZORAN AND Y. WEISS, *From learning models of natural image patches to whole image restoration*, in 2011 International Conference on Computer Vision, IEEE, nov 2011, pp. 479–486, <https://doi.org/10.1109/ICCV.2011.6126278>, <http://people.csail.mit.edu/danielzoran/EPLLICCVCameraReady.pdf>.
- [97] F. ZOU, L. SHEN, Z. JIE, W. ZHANG, AND W. LIU, *A sufficient condition for convergences of adam and rmsprop*, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 11127–11135.

A Organization of the appendix

In this appendix we gather the proofs of the main document. We first derive technical results in Appendix B. Proofs of Section 5.2 are presented in Appendix C. Proofs of Section 5.3 are given in Appendix D. Finally, the proofs of Section 5.4 can be found in Appendix E.

B Technical results

In this section, we gather technical results which will be used throughout our analysis. Let $b \in C(\mathbb{R}^d, \mathbb{R}^d)$ such that for any $x \in \mathbb{R}^d$, the following Stochastic Differential Equation admits a unique strong solution

$$d\mathbf{X}_t = b(\mathbf{X}_t)dt + \sqrt{2}d\mathbf{B}_t, \quad (33)$$

where $(\mathbf{B}_t)_{t \geq 0}$ is a d -dimensional Brownian motion and $\mathbf{X}_0 = x$. In this case, (33) defines a Markov semi-group $(P_t)_{t \geq 0}$ for any $x \in \mathbb{R}^d$ and $A \in \mathcal{B}(\mathbb{R}^d)$ by $P_t(x, A) = \mathbb{P}(\mathbf{X}_t \in A)$ where $(\mathbf{X}_t)_{t \geq 0}$ is the solution of (33) with $\mathbf{X}_0 = x$. Consider now the generator of $(P_t)_{t \geq 0}$, defined for any $f \in C^2(\mathbb{R}^d, \mathbb{R})$ by

$$\mathcal{A}f = \langle \nabla f, b(x) \rangle + \Delta f.$$

We say that a Markov semi-group $(P_t)_{t \geq 0}$ on $\mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d)$ with extended infinitesimal generator $(\mathcal{A}, D(\mathcal{A}))$ (see e.g. [62] for the definition of $(\mathcal{A}, D(\mathcal{A}))$) satisfies a continuous drift condition $\mathbf{D}_c(W, \zeta, \beta)$ if there exist $\zeta > 0$, $\beta \geq 0$ and a measurable function $W : \mathbb{R}^d \rightarrow [1, +\infty)$ with $W \in D(\mathcal{A})$ such that for all $x \in \mathbb{R}^d$

$$\mathcal{A}W(x) \leq -\zeta W(x) + \beta.$$

Similarly, we consider the Markov chain $(X_k)_{k \in \mathbb{N}}$ given by the following recursion for any $k \in \mathbb{N}$ and $x \in \mathbb{R}^d$

$$X_{k+1} = X_k + \gamma b(X_k) + \sqrt{2\gamma}Z_k,$$

with $X_0 = x$, $\gamma > 0$ and $\{Z_k : k \in \mathbb{N}\}$ a family of i.i.d Gaussian random variables with zero mean and identity covariance matrix. We define its associated Markov kernel $R_\gamma : \mathbb{R}^d \times \mathcal{B}(\mathbb{R}^d) \rightarrow [0, 1]$ as follows for any $x \in \mathbb{R}^d$ and $A \in \mathcal{B}(\mathbb{R}^d)$

$$R_\gamma(x, A) = \int_{\mathbb{R}^d} \mathbf{1}_A(x + \gamma b(x) + \sqrt{2\gamma}z) \exp[-\|z\|^2/2] dz .$$

We say that R_γ satisfies a discrete drift condition $\mathbf{D}_d(W, \lambda, c)$ if there exist $\lambda \in [0, 1)$, $c \geq 0$ and a measurable function $W : \mathbb{R}^d \rightarrow [1, +\infty)$ such that for all $x \in \mathbb{R}^d$

$$R_\gamma W(x) \leq \lambda W(x) + c .$$

The following two lemmas are classical, see for instance [17, Lemma 18, Lemma 19]. We recall these results and their proofs for the sake of completeness.

Lemma 24. *Assume that there exist $L, c \geq 0$ and $m > 0$ such that for any $x_1, x_2 \in \mathbb{R}^d$ we have*

$$\langle b(x_1), x_1 \rangle \leq -m \|x_1\|^2 + c , \quad \|b(x_1) - b(x_2)\| \leq L \|x_1 - x_2\| . \quad (34)$$

Let $\bar{\gamma} = m/L^2$. Then the following results hold:

(a) For any $\varpi \in \mathbb{N}^*$ there exist $\lambda \in (0, 1]$, $c, \beta \geq 0$ and $\zeta > 0$ such that for any $\gamma \in (0, \bar{\gamma}]$, R_γ satisfies $\mathbf{D}_d(W, \lambda^\gamma, c\gamma)$ and $(P_t)_{t \geq 0}$ satisfies $\mathbf{D}_c(W, \zeta, \beta)$ with $W(x) = 1 + \|x\|^{2\varpi}$.

(b) For any $\varpi > 0$, there exist $\lambda \in (0, 1]$, $c, \beta \geq 0$ and $\zeta > 0$ such that for any $\gamma \in (0, \bar{\gamma}]$, R_γ satisfies $\mathbf{D}_d(W, \lambda^\gamma, c\gamma)$ and $(P_t)_{t \geq 0}$ satisfies $\mathbf{D}_c(W, \zeta, \beta)$ with $W(x) = \exp[\varpi \sqrt{1 + \|x\|^2}]$.

Proof. We divide the proof into two parts.

(a) Let $\varpi \in \mathbb{N}^*$ and $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} = m/(4L^2)$. Let $\mathcal{T}_\gamma(x) = x - \gamma b(x)$. In the sequel, for any $k \in \{1, \dots, \varpi\}$, $c, \tilde{c}_k \geq 0$ and $\lambda, \tilde{\lambda}_k \in [0, 1)$ are constants independent of γ which may take different values at each appearance. Let $\varepsilon \in (0, 1/2)$. Using (34), the fact that for any $a, b \geq 0$, $(a+b)^2 \leq (1+\varepsilon)a^2 + (1+\varepsilon^{-1})b^2$ and the fact that for any $a, b \geq 0$ we have $(a+b)^{1/2} \leq a^{1/2} + b^{1/2}$, we get that for any $x \in \mathbb{R}^d$ with $\|x\| \geq (2c/(\varepsilon m))^{1/2}$

$$\begin{aligned} \|\mathcal{T}_\gamma(x)\| &= \left(\|x\|^2 + 2\gamma \langle b(x), x \rangle + \gamma^2 \|b(x)\|^2 \right)^{1/2} \\ &\leq \left((1 - 2\gamma m + (1 + \varepsilon)\gamma^2 L^2) \|x\|^2 + 2\gamma c + (1 + \varepsilon^{-1})\gamma^2 \|b(0)\|^2 \right)^{1/2} \\ &\leq \left((1 - \gamma m + (1 + \varepsilon)\gamma^2 L^2) \|x\|^2 + (1 + \varepsilon^{-1})\gamma^2 \|b(0)\|^2 \right)^{1/2} \\ &\leq \exp[-\gamma((2 - \varepsilon)m - (1 + \varepsilon)L^2\bar{\gamma})/2] \|x\| + (1 + \varepsilon^{-1/2})\gamma \|b(0)\| . \end{aligned} \quad (35)$$

Note that $(2 - \varepsilon)m - (1 + \varepsilon)L^2\bar{\gamma} < 0$ since $\varepsilon \in (0, 1/2)$ and $\bar{\gamma} = m/L^2$. On the other hand using (34) and the fact that for any $a, b \geq 0$ with $a \geq b$ and $e^a - e^b \leq e^a(a - b)$, we have for any $x \in \mathbb{R}^d$ with $\|x\| \leq (2c/(\varepsilon m))^{1/2}$

$$\begin{aligned} \|\mathcal{T}_\gamma(x)\| &\leq (1 + \gamma L) \|x\| + \gamma \|b(0)\| \\ &\leq \exp[-\gamma((2 - \varepsilon)m - (1 + \varepsilon)L^2\bar{\gamma})/2] \|x\| \\ &\quad + (2c/(\varepsilon m))^{1/2} \{ \exp[\gamma L] - \exp[-\gamma((2 - \varepsilon)m - (1 + \varepsilon)L^2\bar{\gamma})/2] \} + \gamma \|b(0)\| \\ &\leq \exp[-\gamma((2 - \varepsilon)m - (1 + \varepsilon)L^2\bar{\gamma})/2] \|x\| + \gamma(2c/(\varepsilon m))^{1/2} \exp[\bar{\gamma}L](L + 2m) + \gamma \|b(0)\| . \end{aligned} \quad (36)$$

Combining (35) and (36), there exist $\lambda \in [0, 1)$ and $c \geq 0$ such that for any $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$,

$$\|\mathcal{T}_\gamma(x)\| \leq \lambda^\gamma \|x\| + \gamma c. \quad (37)$$

Note that using (37), for any $k \in \{1, \dots, 2\varpi\}$ there exist $\tilde{\lambda}_k \in (0, 1)$ and $\tilde{c}_k \geq 0$ such that

$$\begin{aligned} \|\mathcal{T}_\gamma(x)\|^k &\leq \{\tilde{\lambda}_k^\gamma \|x\| + \gamma \tilde{c}_k\}^k \\ &\leq \tilde{\lambda}_k^{\gamma k} \|x\|^k + \gamma 2^k \max(\tilde{c}_k, 1)^k \max(\bar{\gamma}, 1)^{k-1} \{1 + \|x\|^{k-1}\} \\ &\leq \tilde{\lambda}_k^\gamma \|x\|^k + \tilde{c}_k \gamma \{1 + \|x\|^{k-1}\} \leq (1 + \|x\|^k)(1 + \tilde{c}_k \gamma). \end{aligned} \quad (38)$$

Therefore, combining (38) and the Cauchy-Schwarz inequality we obtain that for any $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$

$$\begin{aligned} \int_{\mathbb{R}^d} (1 + \|y\|^{2\varpi}) R_\gamma(x, dy) &= 1 + \mathbb{E}[\|\mathcal{T}_\gamma(x)\|^2 + 2\sqrt{2\gamma} \langle \mathcal{T}_\gamma(x), Z \rangle + 2\gamma \|Z\|^2]^\varpi \\ &= 1 + \sum_{k=0}^{\varpi} \sum_{\ell=0}^k \binom{\varpi}{k} \binom{k}{\ell} \|\mathcal{T}_\gamma(x)\|^{2(\varpi-k)} 2^{(3k-\ell)/2} \gamma^{(k+\ell)/2} \mathbb{E}[\langle \mathcal{T}_\gamma(x), Z \rangle^{k-\ell} \|Z\|^{2\ell}] \\ &\leq 1 + \|\mathcal{T}_\gamma(x)\|^{2\varpi} \\ &\quad + 2^{3\varpi/2} \sum_{k=1}^{\varpi} \sum_{\ell=0}^k \binom{\varpi}{k} \binom{k}{\ell} \|\mathcal{T}_\gamma(x)\|^{2(\varpi-k)} \gamma^{(k+\ell)/2} \mathbb{E}[\langle \mathcal{T}_\gamma(x), Z \rangle^{k-\ell} \|Z\|^{2\ell}] \mathbf{1}_{\{(1,0)\}^c}(k, \ell) \\ &\leq 1 + \|\mathcal{T}_\gamma(x)\|^{2\varpi} \\ &\quad + \gamma 2^{3\varpi/2} \sum_{k=1}^{\varpi} \sum_{\ell=0}^k \binom{\varpi}{k} \binom{k}{\ell} \|\mathcal{T}_\gamma(x)\|^{2\varpi-k-\ell} \bar{\gamma}^{(k+\ell)/2-1} \mathbb{E}[\|Z\|^{k+\ell}] \mathbf{1}_{\{(1,0)\}^c}(k, \ell) \\ &\leq 1 + \tilde{\lambda}_{2\varpi}^\gamma \|x\|^{2\varpi} + \tilde{c}_{2\varpi} \gamma \{1 + \|x\|^{2\varpi-1}\} \\ &\quad + \gamma 2^{3\varpi/2} 2^{2\varpi} \max(\bar{\gamma}, 1)^{2\varpi} \sup_{k \in \{1, \dots, \varpi\}} \{(1 + \tilde{c}_k \bar{\gamma}) \mathbb{E}[\|Z\|^k]\} (1 + \|x\|^{2\varpi-1}) \\ &\leq 1 + \lambda^\gamma \|x\|^{2\varpi} + \gamma c (1 + \|x\|^{2\varpi-1}) \\ &\leq \lambda^{\gamma/2} (1 + \|x\|^{2\varpi}) + \gamma c (1 + \|x\|^{2\varpi-1}) + \lambda^\gamma (1 + \|x\|^{2\varpi}) - \lambda^{\gamma/2} (1 + \|x\|^{2\varpi}). \end{aligned}$$

Using that $\lambda^\gamma - \lambda^{\gamma/2} \leq -\log(1/\lambda) \gamma \lambda^{\gamma/2} / 2$, we get that for any $\gamma \in (0, \bar{\gamma}]$, R_γ satisfies $\mathbf{D}_d(W, \lambda^\gamma, c\gamma)$. We now show that there exist $\zeta > 0$ and $\beta \geq 0$ such that $(P_t)_{t \geq 0}$ satisfies $\mathbf{D}_c(W, \zeta, \beta)$. First, for any $x \in \mathbb{R}^d$ we have

$$\nabla W(x) = 2\varpi \|x\|^{2(\varpi-1)} x, \quad \Delta W(x) = 2\varpi(2\varpi - 1) \|x\|^{2(\varpi-1)}$$

Combining this result, the Cauchy-Schwarz inequality and (34), we obtain that for any $x \in \mathbb{R}^d$

$$\begin{aligned} \mathcal{A}W(x) &= \langle \nabla W(x), b(x) \rangle + \Delta W(x) \\ &\leq -2\mathfrak{m}\varpi \|x\|^{2\varpi} + 2\varpi c \|x\|^{2\varpi-1} + 2\varpi(2\varpi - 1) \|x\|^{2(\varpi-1)} \\ &\leq -\mathfrak{m}\varpi \|x\|^{2\varpi} + \sup_{x \in \mathbb{R}^d} \{2\varpi(c + 2\varpi - 1) \|x\|^{2\varpi-1} - \mathfrak{m}\varpi \|x\|^{2\varpi}\} \\ &\leq -\mathfrak{m}\varpi W(x) + \sup_{x \in \mathbb{R}^d} \{2\varpi(c + 2\varpi - 1) \|x\|^{2\varpi-1} - \mathfrak{m}\varpi \|x\|^{2\varpi}\} + \mathfrak{m}\varpi. \end{aligned}$$

Hence letting $\zeta = \mathfrak{m}\varpi$ and $\beta = \sup_{x \in \mathbb{R}^d} \{2\varpi(c + 2\varpi - 1) \|x\|^{2\varpi-1} - \mathfrak{m}\varpi \|x\|^{2\varpi}\} + \mathfrak{m}\varpi$, we obtain that $(P_t)_{t \geq 0}$ satisfies $\mathbf{D}_c(W, \zeta, \beta)$.

(b) First, we show that for any $\gamma \in (0, \bar{\gamma}]$, R_γ satisfies $\mathbf{D}_d(\Phi, \lambda^\gamma, c)$, where $\Phi(x) = (1 + \|x\|^2)^{1/2} = W_2^{1/2}(x)$ and $W_2(x) = 1 + \|x\|^2$. Using the first part of the proof, there exist $\lambda_0 \in [0, 1)$ and $c_0 \geq 0$ such that for any $\gamma \in (0, \bar{\gamma}]$ with $\bar{\gamma} = m/(4L^2)$ we have that R_γ satisfies $\mathbf{D}_d(W_2, \lambda_0^\gamma, c_0\gamma)$. Using Jensen's inequality we obtain that for any $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$ with $\|x\| \geq R$ and $R = \max(1, ((2c_0\lambda_0^{-\bar{\gamma}})/\log(1/\lambda_0))^{1/2})$ we have

$$R_\gamma\Phi(x) \leq (R_\gamma W_2(x))^{1/2} \leq \exp[(\gamma/2)\{\log(\lambda_0) + \lambda_0^{-\bar{\gamma}}c_0R^{-2}\}]\Phi(x) \leq \lambda_0^{\gamma/4}\Phi(x).$$

In addition, using that for any $a, b \geq 0$ with $a \geq b$ we have $e^a - e^b \leq e^a(b - a)$, we get for any $x \in \mathbb{R}^d$ with $\|x\| \leq R$

$$\begin{aligned} R_\gamma\Phi(x) &\leq (R_\gamma W_2(x))^{1/2} \leq \exp[(\gamma/2)\{\log(\lambda_0) + \lambda_0^{-\bar{\gamma}}c_0\}]\Phi(x) \\ &\leq \exp[(\gamma/2)\{\log(\lambda_0) + \lambda_0^{-\bar{\gamma}}c_0R^{-2}\}]\Phi(x) \\ &\quad + \lambda_0^{-\bar{\gamma}}c_0(1 - R^{-2}) \exp[(\gamma/2)\{\log(\lambda_0) + \lambda_0^{-\bar{\gamma}}c_0R^{-2}\}]\Phi(R). \end{aligned}$$

Hence, there exist $\lambda_1 \in [0, 1)$ and $c_1 \geq 0$ such that for any $\gamma \in (0, \bar{\gamma}]$ we have that R_γ satisfies $\mathbf{D}_d(\varpi\Phi, \lambda_1^\gamma, c_1\gamma)$. Now let $W(x) = \exp[\Phi(x)]$. Using the logarithmic Sobolev inequality [19, Theorem 5.5] we get for any $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$ with $\|x\| \geq R$ and $R = 1 + (\varpi^2 + c_1)^{-1} \log(1/\lambda_1)$

$$\begin{aligned} R_\gamma W(x) &\leq \exp[R_\gamma\varpi\Phi(x) + \gamma\varpi^2] \leq \exp[-(1 - \lambda_1^\gamma)\Phi(x) + \gamma(\varpi^2 + c_1)]W(x) \\ &\leq \exp[-\gamma \log(1/\lambda_1)R + \gamma(\varpi^2 + c_1)]W(x) \leq \lambda_1^\gamma W(x). \end{aligned}$$

In addition, using that for any $a, b \geq 0$ with $a \geq b$ we have $e^a - e^b \leq e^a(b - a)$, we get for any $x \in \mathbb{R}^d$ with $\|x\| \leq R$

$$\begin{aligned} R_\gamma W(x) &\leq \exp[R_\gamma\varpi\Phi(x) + \gamma] \leq \exp[\gamma(\varpi^2 + c_1)]W(x) \\ &\leq \lambda_1^\gamma W(x) + \gamma \exp[\bar{\gamma}(\varpi^2 + c_1)]((1 + c_1) + \log(1/\lambda_1))W(R). \end{aligned}$$

Therefore, there exist $\lambda \in [0, 1)$ and $c \geq 0$ such that for any $\gamma \in (0, \bar{\gamma}]$ we have that R_γ satisfies $\mathbf{D}_d(W, \lambda^\gamma, c\gamma)$. We now show that there exist $\zeta > 0$ and $\beta \geq 0$ such that $(P_t)_{t \geq 0}$ satisfies $\mathbf{D}_c(W, \zeta, \beta)$. First, for any $x \in \mathbb{R}^d$ we have

$$\nabla W(x) = \varpi x \Phi^{-1}(x)W(x), \quad \Delta W(x) = \{\varpi \Phi^{-1}(x)(1 - \|x\|^2/\Phi^2(x)) + \varpi^2 \|x\|^2/\Phi^2(x)\}W(x).$$

Therefore using (34) we obtain that for any $x \in \mathbb{R}^d$ with $\|x\| \geq \sqrt{2}(1 + (c + 1 + \varpi)/m)$

$$\mathcal{A}W(x) \leq \varpi(-m\Phi^{-1}(x)\|x\|^2 + c + 1 + \varpi)W(x) \leq -(m/2)W(x),$$

which concludes the proof. \square

Lemma 25. *Assume that there exist $\lambda \in (0, 1]$, $c, \beta \geq 0$, $\zeta, \bar{\gamma} > 0$ such that for any $\gamma \in (0, \bar{\gamma}]$, R_γ satisfies $\mathbf{D}_d(W, \lambda^\gamma, c\gamma)$ and $(P_t)_{t \geq 0}$ satisfies $\mathbf{D}_c(W, \zeta, \beta)$. Then, there exists $C \geq 0$ such that for any $x \in \mathbb{R}^d$, $t \geq 0$ and $k \in \mathbb{N}^*$ we have*

$$R_\gamma^k W(x) + P_t W(x) \leq CW(x).$$

Proof. There exists $C_c \geq 0$ such that for any $x \in \mathbb{R}^d$ and $t \geq 0$, $P_t W(x) \leq C_c W(x)$ using [15, Lemma 25-(b)]. Using that for any $t \geq 0$, $(1 - e^{-t})^{-1} \leq 1 + 1/t$ we get that for any $\gamma \in (0, \bar{\gamma}]$, $x \in \mathbb{R}^d$ and $k \in \mathbb{N}^*$

$$R_\gamma^k W(x) \leq W(x) + c\gamma \sum_{k \in \mathbb{N}} \lambda^{k\gamma} \leq (1 + c(\bar{\gamma} + \log(1/\lambda)))W(x),$$

which concludes the proof upon letting $C = C_c + 1 + c(\bar{\gamma} + \log(1/\lambda))$. \square

We conclude this section with the proof of Proposition 9.

Proof. Let $y_1, y_2 \in \mathsf{K}$ with K a compact set. Let $y_0 \in \mathsf{K}$ and D_{K} be the diameter of K . Using Lemma 29 we get that

$$\|\pi_{y_1} - \pi_{y_2}\|_{\text{TV}} \leq 2c_{y_1} \int_{\mathbb{R}^d} |q_{y_1}(x) - q_{y_2}(x)| p(x) dx ,$$

with $c_{y_1} = \int_{\mathbb{R}^d} q_{y_1}(x) p(x) dx$. Combining this result with the fact that for any $a, b \in \mathbb{R}$ we have $|e^a - e^b| \leq |a - b| \max(e^a, e^b)$ we get that

$$\begin{aligned} \|\pi_{y_1} - \pi_{y_2}\|_{\text{TV}} &\leq 2c_{y_1} \int_{\mathbb{R}^d} |q_{y_1}(x) - q_{y_2}(x)| p(x) dx \\ &\leq 2c_{y_1} \int_{\mathbb{R}^d} (\Phi_1(x) + \Phi_2(y_1) + \Phi_2(y_2)) \|y_1 - y_2\| \\ &\quad \times \exp[(2\Phi_1(x) + \Phi_2(y_1) + \Phi_2(y_0) + \Phi_2(y_2))D_{\mathsf{K}}] p(x) dx \\ &\leq 2c_{y_1} (\Phi_2(y_1) + \Phi_2(y_2)) \exp[\Phi_2(y_1) + \Phi_2(y_0) + \Phi_2(y_2)] \\ &\quad \times \int_{\mathbb{R}^d} (1 + \Phi_1(x)) \exp[2D_{\mathsf{K}}\Phi_1(x)] p(x) dx \times \|y_1 - y_2\| , \end{aligned}$$

which concludes the proof. \square

C Proofs of Section 5.2

We recall that the Markov chain $(X_k)_{k \in \mathbb{N}}$ is given by

$$\begin{aligned} X_{k+1} &= X_k + \delta b_\varepsilon(X_k) + \sqrt{2\delta} Z_{k+1} , \\ b_\varepsilon(x) &= \nabla \log(q_y(x)) + \alpha(D_\varepsilon(x) - x)/\varepsilon + (x - \Pi_{\mathsf{C}}(x))/\lambda , \end{aligned} \tag{39}$$

where $\delta > 0$ is a step-size, $\alpha, \varepsilon, \lambda > 0$ are hyperparameters of the algorithm, $\mathsf{C} \subset \mathbb{R}^d$ is a closed convex set with $0 \in \mathsf{C}$, Π_{C} is the projection on C and $\{Z_k : k \in \mathbb{N}\}$ a family of i.i.d. Gaussian random variables with zero mean and identity covariance matrix.

C.1 Proof of Proposition 10

Let $R > 0$. Let X and Z be random variables with distribution μ and zero mean Gaussian with identity covariance matrix. Let $X_\varepsilon = X + \varepsilon^{1/2}Z$. We recall that the distributions of X and X_ε have density with respect to the Lebesgue measure given by p and p_ε respectively. In addition, the conditional density of X given X_ε is given by g_ε . By definition $D_\varepsilon^*(X_\varepsilon) = \mathbb{E}[X|X_\varepsilon]$ and therefore we have

$$\begin{aligned} \ell_\varepsilon(w^*) &= \mathbb{E} \left[\|X - f_{w^*}(X_\varepsilon)\|^2 \right] \\ &= \mathbb{E} \left[\|X - D_\varepsilon^*(X_\varepsilon)\|^2 \right] + 2\mathbb{E} [\langle X - D_\varepsilon^*(X_\varepsilon), D_\varepsilon^*(X_\varepsilon) - f_{w^*}(X_\varepsilon) \rangle] + \mathbb{E} \left[\|f_{w^*}(X_\varepsilon) - D_\varepsilon^*(X_\varepsilon)\|^2 \right] \\ &= \mathbb{E} \left[\|X - D_\varepsilon^*(X_\varepsilon)\|^2 \right] + \mathbb{E} \left[\|f_{w^*}(X_\varepsilon) - D_\varepsilon^*(X_\varepsilon)\|^2 \right] = \ell_\varepsilon^* + \mathbb{E} \left[\|f_{w^*}(X_\varepsilon) - D_\varepsilon^*(X_\varepsilon)\|^2 \right] . \end{aligned}$$

Combining this result, the condition that $\ell_\varepsilon(w^*) \leq \ell_\varepsilon^* + \eta$ and the Cauchy-Schwarz inequality we get that

$$\mathbb{E}[\|f_{w^*}(X_\varepsilon) - D_\varepsilon^*(X_\varepsilon)\|] \leq \sqrt{\eta} . \tag{40}$$

Since f_{w^*} and D_ε^* are locally Lipschitz, there exists $C_R \geq 0$ such that for any $x_1, x_2 \in \overline{\mathbb{B}}(0, 2R)$ we have

$$\|f_{w^*}(x_2) - D_\varepsilon^*(x_2)\| - \|f_{w^*}(x_1) - D_\varepsilon^*(x_1)\| \leq C_R \|x_2 - x_1\|. \quad (41)$$

Assume that $\sup_{\tilde{x} \in \overline{\mathbb{B}}(0, R)} \|f_{w^*}(\tilde{x}) - D_\varepsilon^*(\tilde{x})\| > \eta^\varpi$ with $\varpi = (2d + 2)^{-1}$ and denote $x_R \in \overline{\mathbb{B}}(0, R)$ such that we have $\sup_{\tilde{x} \in \overline{\mathbb{B}}(0, R)} \|f_{w^*}(\tilde{x}) - D_\varepsilon^*(\tilde{x})\| = \|f_{w^*}(x_R) - D_\varepsilon^*(x_R)\|$. Using (41) we have

$$\begin{aligned} \mathbb{E}[\|f_{w^*}(X_\varepsilon) - D_\varepsilon^*(X_\varepsilon)\|] &\geq \int_{\overline{\mathbb{B}}(0, 2R) \cap \overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)} \|f_{w^*}(\tilde{x}) - D_\varepsilon^*(\tilde{x})\| p_\varepsilon(\tilde{x}) d\tilde{x} \\ &\geq (\|f_{w^*}(x_R) - D_\varepsilon^*(x_R)\| - \eta^\varpi) \int_{\overline{\mathbb{B}}(0, 2R) \cap \overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)} p_\varepsilon(\tilde{x}) d\tilde{x}. \end{aligned}$$

Combining this result and (40) we obtain that

$$\|f_{w^*}(x_R) - D_\varepsilon^*(x_R)\| \leq \eta^{1/2} \left(\int_{\overline{\mathbb{B}}(0, 2R) \cap \overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)} p_\varepsilon(\tilde{x}) d\tilde{x} \right)^{-1} + \eta^\varpi,$$

Setting $\mathbb{M}_R = \eta^{1/2} \left(\int_{\overline{\mathbb{B}}(0, 2R) \cap \overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)} p_\varepsilon(\tilde{x}) d\tilde{x} \right)^{-1} + \eta^\varpi$ concludes the first part of the proof. Denote v_d the volume of the unit d -dimensional ball. We have that $\text{Leb}(\overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)) = C_R^{-d} \eta^{\varpi d} v_d$. Using the Fubini theorem, the Lebesgue differentiation theorem [13, Theorem 5.6.2], the dominated convergence theorem and the fact that for $\eta \in (0, (C_R R)^{1/\varpi}]$, $\overline{\mathbb{B}}(0, 2R) \cap \overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi) = \overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)$ we get that

$$\begin{aligned} \lim_{\eta \rightarrow 0} \text{Leb}(\overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi))^{-1} \int_{\mathbb{R}^d} \mathbf{1}_{\overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi) \cap \overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)}(x) p_\varepsilon(x) dx \\ = \lim_{\eta \rightarrow 0} \int_{\mathbb{R}^d} |\overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)|^{-1} (2\pi\varepsilon)^{-d/2} \int_{\mathbb{R}^d} \mathbf{1}_{\overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)}(x) \exp[-\|x - \tilde{x}\|^2 / (2\varepsilon)] p(\tilde{x}) dx d\tilde{x} \\ = \int_{\mathbb{R}^d} (2\pi\varepsilon)^{-d/2} \exp[-\|x_R - \tilde{x}\|^2 / (2\varepsilon)] p(\tilde{x}) dx d\tilde{x} = p_\varepsilon(x_R) > 0. \end{aligned}$$

Using this result we have,

$$\begin{aligned} \limsup_{\eta \rightarrow 0} \eta^{-\varpi} \mathbb{M}_R &= 1 + \limsup_{\eta \rightarrow 0} \eta^{1/2 - \varpi(d+1)} \eta^{\varpi d} \left(\int_{\overline{\mathbb{B}}(0, 2R) \cap \overline{\mathbb{B}}(x_R, C_R^{-1}\eta^\varpi)} p_\varepsilon(\tilde{x}) d\tilde{x} \right)^{-1} \\ &= 1 + C_R^d v_d p_\varepsilon^{-1}(x_R) < +\infty, \end{aligned}$$

which concludes the proof.

C.2 Proof of Proposition 11, Proposition 14

We divide this section into two parts. First, we prove the general case where $\log(q_y)$ is not assumed to be strongly concave but only satisfying a one-sided Lipschitz condition, *i.e.* Proposition 14. Then we turn to the proof of Proposition 11.

(a) Let $\lambda > 0$ such that $2\lambda(\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon) \leq 1$ and $\bar{\delta} = (1/3)(\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon + 1/\lambda)^{-1}$. Let \mathbf{C} be a compact convex set with $0 \in \mathbf{C}$. Using H3, (39) and that $\text{Id} - \Pi_{\mathbf{C}}$ is non-expansive we have for any $x_1, x_2 \in \mathbb{R}^d$

$$\|b_\varepsilon(x_1) - b_\varepsilon(x_2)\| \leq (\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon + 1/\lambda) \|x_1 - x_2\|.$$

Denote $R_C = \sup\{\|x_1 - x_2\| : x_1, x_2 \in C\}$. Using (39), the Cauchy-Schwarz inequality and that $2\lambda(\alpha L/\varepsilon - m) \leq 1$ we have for any $x_1, x_2 \in \mathbb{R}^d$

$$\begin{aligned} \langle b_\varepsilon(x_1) - b_\varepsilon(x_2), x_1 - x_2 \rangle &\leq (-m + \alpha L/\varepsilon) \|x_1 - x_2\|^2 - \|x_1 - x_2\|^2/\lambda + R_C \|x_1 - x_2\|/\lambda \\ &\leq -\|x_1 - x_2\|^2/(2\lambda) + R_C \|x_1 - x_2\|/\lambda. \end{aligned}$$

Hence, for any $x_1, x_2 \in \mathbb{R}^d$ with $\|x_1 - x_2\| \geq 4R_C$ we obtain that $\langle b_\varepsilon(x_1) - b_\varepsilon(x_2), x_1 - x_2 \rangle \leq -\|x_1 - x_2\|^2/(4\lambda)$. We also have that for any $x \in \mathbb{R}^d$

$$\langle b_\varepsilon(x), x \rangle \leq -\|x\|^2/(4\lambda) + \sup_{\tilde{x} \in \mathbb{R}^d} \left\{ (R_C/\lambda + \|b(0)\|) \|\tilde{x}\| - \|\tilde{x}\|^2/(4\lambda) \right\}.$$

We conclude the proof of Proposition 14 upon using Lemma 24, Lemma 25, [15, Corollary 2] with $\bar{\gamma} \leftarrow (4\lambda)^{-1}(\mathbf{L}_y + \alpha L/\varepsilon + 1/\lambda)^{-2} \geq \bar{\delta}$ and the fact that for any probability distribution ν_1, ν_2 ,

$$\|\nu_1 - \nu_2\|_V \leq \|\nu_1 - \nu_2\|_{TV}^{1/2} (\nu_1[V^2] + \nu_2[V^2])^{1/2}. \quad (42)$$

(b) Using that $\log(q_y)$ is m -concave with $2\alpha L/(m\varepsilon) \leq 1$, we obtain that for any $x_1, x_2 \in \mathbb{R}^d$

$$\begin{aligned} \langle b_\varepsilon(x_1) - b_\varepsilon(x_2), x_1 - x_2 \rangle &\leq -m \|x_1 - x_2\|^2/2, \\ \|b_\varepsilon(x_1) - b_\varepsilon(x_2)\| &\leq (\mathbf{L}_y + \alpha L/\varepsilon) \|x_1 - x_2\|. \end{aligned}$$

This concludes the proof of Proposition 11 upon using [15, Corollary 2] with $\bar{\gamma} \leftarrow m(\mathbf{L}_y + \alpha L/\varepsilon)^{-2} \geq \bar{\delta}$ and (42).

C.3 Proof of Proposition 12, Proposition 15

Before proving Proposition 12 and Proposition 15, we show the following lemma which is a straightforward consequence of Girsanov's theorem [58, Theorem 7.7]. A similar version of this lemma can be found in the proof of [37, Proposition 2].

Lemma 26. *Let $T > 0$, $b_1, b_2 : [0, +\infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ measurable such that for any $i \in \{1, 2\}$ and $x \in \mathbb{R}^d$, $d\mathbf{X}_t^{(i)} = b_i(t, \mathbf{X}_t^{(i)})dt + \sqrt{2}d\mathbf{B}_t$ admits a unique strong solution with $\mathbf{X}_0^{(i)} = x$ with Markov semigroup $(\mathbf{P}_t^{(i)})_{t \geq 0}$ and where $(\mathbf{B}_t)_{t \geq 0}$ is a d -dimensional Brownian motion. In addition, assume that for any $x \in \mathbb{R}^d$ and $\mathbb{P}(\int_0^T \{\|b_1(t, \mathbf{X}_t^{(1)})\|^2 + \|b_2(t, \mathbf{B}_t)\|^2\} dt < +\infty) = 1$. Let $V : \mathbb{R}^d \rightarrow [0, +\infty)$ measurable, then for any $x \in \mathbb{R}^d$ we have*

$$\begin{aligned} &\left\| \delta_x \mathbf{P}_T^{(1)} - \delta_x \mathbf{P}_T^{(2)} \right\|_V \\ &\leq \left(\delta_x \mathbf{P}_t^{(1)}[V^2] + \delta_x \mathbf{P}_t^{(2)}[V^2] \right)^{1/2} \left(\int_0^T \mathbb{E} \left[\|b_1(t, \mathbf{X}_t^{(1)}) - b_2(t, \mathbf{X}_t^{(1)})\|^2 \right] dt \right)^{1/2}. \end{aligned}$$

Proof. Let $T > 0$ and $x \in \mathbb{R}^d$. For any $i \in \{1, 2\}$, denote $\mu_{(i)}^x$ the distribution of $(\mathbf{X}_t^{(i)})_{t \in [0, T]}$ on the Wiener space $(C([0, T], \mathbb{R}), \mathcal{B}(C([0, T], \mathbb{R})))$ with $\mathbf{X}_0^{(i)} = x$. Similarly denote μ_B^x the distribution of $(\mathbf{B}_t)_{t \in [0, T]}$ with $\mathbf{B}_0 = x$. Using the generalized Pinsker inequality [37, Lemma 24] and the transfer theorem [54, Theorem 4.1] we get that

$$\left\| \delta_x \mathbf{P}_T^{(1)} - \delta_x \mathbf{P}_T^{(2)} \right\|_V \leq \sqrt{2} \left(\delta_x \mathbf{P}_t^{(1)}[V^2] + \delta_x \mathbf{P}_t^{(2)}[V^2] \right)^{1/2} \text{KL}^{1/2}(\mu_{(1)} | \mu_{(2)}).$$

Since for any $i \in \{1, 2\}$ we have $\mathbb{P}(\int_0^T \{\|b_i(\mathbf{X}_t^{(i)})\|^2 + \|b_i(\mathbf{B}_t)\|^2\} dt < +\infty) = 1$, we can apply Girsanov's theorem [58, Theorem 7.7] and μ_B -almost surely for any $w \in C([0, T], \mathbb{R})$ we get

$$\begin{aligned} (d\mu_{(1)}^x/d\mu_B^x)((w_t)_{t \in [0, T]}) &= \exp \left[(1/2) \int_0^T \langle b_1(w_t), dw_t \rangle - (1/4) \int_0^T \|b_1(w_t)\|^2 dt \right], \\ (d\mu_B^x/d\mu_{(2)}^x)((w_t)_{t \in [0, T]}) &= \exp \left[-(1/2) \int_0^T \langle b_2(w_t), dw_t \rangle + (1/4) \int_0^T \|b_2(w_t)\|^2 dt \right]. \end{aligned}$$

Hence, we obtain that

$$\text{KL}(\mu_{(1)}^x | \mu_{(2)}^x) = \mathbb{E} \left[\log((d\mu_{(1)}^x/d\mu_{(2)}^x)(\mathbf{X}_t^{(1)})) \right] = (1/4) \int_0^T \mathbb{E} \left[\left\| b_1(\mathbf{X}_t^{(1)}) - b_2(\mathbf{X}_t^{(2)}) \right\|^2 \right] dt,$$

which concludes the proof. \square

In the following lemma, we show that under **H4**, $\nabla \log(p_\varepsilon)$ is Lipschitz continuous.

Lemma 27. *Assume **H4**. Then for any $x_1, x_2 \in \mathbb{R}^d$ we have*

$$\|\nabla \log(p_\varepsilon(x_1)) - \nabla \log(p_\varepsilon(x_2))\| \leq (1 + \mathbf{K}_\varepsilon/\varepsilon) \|x_1 - x_2\| / \varepsilon.$$

Reciprocally, if there $x \mapsto \nabla \log(p_\varepsilon(x))$ is Lipschitz-continuous then **H4**.

Proof. Let $\varepsilon > 0$. We recall that for any $x \in \mathbb{R}^d$ we have

$$p_\varepsilon(x) = \int_{\mathbb{R}^d} \exp[-\|x - \tilde{x}\|^2 / (2\varepsilon)] p(\tilde{x}) d\tilde{x}.$$

Using the dominated convergence theorem we obtain that $\log(p_\varepsilon) \in C^\infty(\mathbb{R}^d, \mathbb{R})$. In particular we have for any $x \in \mathbb{R}^d$

$$\begin{aligned} \nabla^2 \log(p_\varepsilon(x)) &= -\varepsilon^{-1} \text{Id} + \varepsilon^{-2} \int_{\mathbb{R}^d} (x - \tilde{x})^{\otimes 2} g_\varepsilon(\tilde{x}|x) d\tilde{x} - \varepsilon^{-2} \left(\int_{\mathbb{R}^d} (x - \tilde{x}) g_\varepsilon(\tilde{x}|x) d\tilde{x} \right)^{\otimes 2} \\ &= -\varepsilon^{-1} \text{Id} + \varepsilon^{-2} \int_{\mathbb{R}^d} \left(\tilde{x} - \int_{\mathbb{R}^d} \tilde{x}' g_\varepsilon(\tilde{x}'|x) d\tilde{x}' \right)^{\otimes 2} g_\varepsilon(\tilde{x}|x) d\tilde{x} \end{aligned} \quad (43)$$

Therefore, using **H4** we obtain that for any $x \in \mathbb{R}^d$ we have

$$\|\nabla^2 \log(p_\varepsilon(x))\|_2 \leq \varepsilon^{-1} + \varepsilon^{-2} \mathbf{K}_\varepsilon,$$

which concludes the first part of the proof. Reciprocally, since $x \mapsto \nabla \log(p_\varepsilon(x))$ is Lipschitz-continuous with constant $\mathbf{K} \geq 0$ we get that for any basis vector $(e_i)_{i \in \{1, \dots, d\}}$ we have that $e_i^\top \nabla^2 \log(p_\varepsilon(x)) e_i \leq \mathbf{K}$. Combining this result with (43), we get that

$$\varepsilon^{-2} \int_{\mathbb{R}^d} \left\| \tilde{x} - \int_{\mathbb{R}^d} \tilde{x}' g_\varepsilon(\tilde{x}'|x) d\tilde{x}' \right\|^2 g_\varepsilon(\tilde{x}|x) d\tilde{x} \leq \mathbf{K}d + \varepsilon^{-1}d,$$

which concludes the proof. \square

In what follows we prove Proposition 15. The proof of Proposition 12 is similar and left to the reader.

Proof of Proposition 15. Let $\lambda > 0$ such that $2\lambda(L_y + \alpha L/\varepsilon - m) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$. We divide the proof into two parts. First, we show that for any C convex compact with $0 \in C$ there exists $B_{1,C} \geq 0$ such that for any $\delta \in (0, \bar{\delta}]$ and $R > 0$

$$\|\pi_{\varepsilon,\delta} - \tilde{\pi}_\varepsilon\|_V \leq B_{1,C}(\delta^{1/2} + M_R + \exp[-R]) ,$$

with $\tilde{\pi}_\varepsilon$ given by

$$(d\tilde{\pi}_\varepsilon/d\text{Leb})(x) \propto \exp[-d^2(x, C)/(2\lambda)]q_y(x)p_\varepsilon^\alpha(x) ,$$

Second, we show that there exists $B_0 \geq 0$ such that for any C convex compact with $0 \in C$

$$\|\pi_\varepsilon - \tilde{\pi}_\varepsilon\|_V \leq B_0 \text{diam}^{-1/4}(C) ,$$

which concludes the proof upon using the triangle inequality.

(a) Let C convex compact with $0 \in C$. We introduce $(\bar{\mathbf{X}}_t)_{t \geq 0}$ solution of the following Stochastic Differential Equation (SDE): $\bar{\mathbf{X}}_0 = X_0$ and

$$\begin{aligned} d\bar{\mathbf{X}}_t &= \bar{b}_\varepsilon(\bar{\mathbf{X}}_t)dt + \sqrt{2}d\mathbf{B}_t , \\ \bar{b}_\varepsilon(x) &= \nabla \log(q_y(x)) + \alpha \nabla \log(p_\varepsilon(x)) + \text{prox}_\lambda(\iota_C)(x) , \end{aligned} \quad (44)$$

with $(\mathbf{B}_t)_{t \geq 0}$ a d -dimensional Brownian motion. \bar{b}_ε is Lipschitz continuous using Lemma 27, hence this SDE admits a unique strong solution for any initial condition \mathbf{X}_0 with $\mathbb{E}[\|\mathbf{X}_0\|^2] < +\infty$, see [51, Chapter 5, Theorem 2.9]. We denote by $(P_{t,\varepsilon})_{t \geq 0}$ the semigroup associated with the strong solutions of (44). Similarly to the proof of Proposition 11, replacing [15, Corollary 2] by [15, Corollary 22], there exist $\tilde{A}_C \geq 0$ and $\tilde{\rho}_C \in [0, 1)$ such that that for any $x_1, x_2 \in \mathbb{R}^d$ and $t \geq 0$

$$\begin{aligned} \|\delta_{x_1}P_{t,\varepsilon} - \delta_{x_2}P_{t,\varepsilon}\|_V &\leq \tilde{A}_C \tilde{\rho}_C^t (V^2(x_1) + V^2(x_2)) , \\ \mathbf{W}_1(\delta_{x_1}P_{t,\varepsilon}, \delta_{x_2}P_{t,\varepsilon}) &\leq \tilde{A}_C \tilde{\rho}_C^t \|x_1 - x_2\| . \end{aligned} \quad (45)$$

Combining (45), Proposition 11, the fact that $(\mathcal{P}_1(\mathbb{R}^d), \mathbf{W}_1)$ is a complete metric space and the Picard fixed point theorem we obtain that for any $\delta \in (0, \bar{\delta}]$ there exist $\pi_{\varepsilon,\delta}, \tilde{\pi}_\varepsilon \in \mathcal{P}_1(\mathbb{R}^d)$ such that $\pi_{\varepsilon,\delta}R_{\varepsilon,\delta,C} = \pi_{\varepsilon,\delta}$ and for any $t \geq 0$, $\tilde{\pi}_\varepsilon P_{t,\varepsilon} = \tilde{\pi}_\varepsilon$. Note that by [71, Theorem 2.1] we have for any $x \in \mathbb{R}^d$

$$(d\tilde{\pi}_\varepsilon/d\text{Leb})(x) \propto \exp[-d^2(x, C)/(2\lambda)]q_y(x)p_\varepsilon^\alpha(x) ,$$

since $\text{prox}_\lambda(\iota_C) = \nabla d^2(\cdot, C)/(2\lambda)$. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ measurable and such that for any $x \in \mathbb{R}^d$, $|f(x)| \leq V(x)$. Let $m \in \mathbb{N}^*$ such that $m \geq \bar{\delta}^{-1}$, $x \in \mathbb{R}^d$ and $k \in \mathbb{N}$ we have

$$\left\| \delta_x R_{\varepsilon,1/m}^{km}[f] - \delta_x P_{km,\varepsilon}^{km}[f] \right\| = \left\| \sum_{j=0}^{k-1} \delta_x R_{\varepsilon,1/m}^{jm} (R_{\varepsilon,1/m}^m - P_{1,\varepsilon}) P_{k-j-1,\varepsilon}[f] \right\| \quad (46)$$

Using (45), Lemma 24 and Lemma 25 there exists $B_a \geq 0$ such that for any $x \in \mathbb{R}^d$ and $k \in \mathbb{N}$ we have

$$\|\delta_x P_{k,\varepsilon,C}[f] - \tilde{\pi}_\varepsilon[f]\| \leq B_a \tilde{\rho}_C^k V^2(x) . \quad (47)$$

Let $T = 1$, $b_1(t, (w_t)_{t \in [0,T]}) = \sum_{j=0}^{m-1} \mathbf{1}_{[j/m, (j+1)/m)}(t) b_\varepsilon(w_j \delta)$ and $b_2(t, (w_t)_{t \in [0,T]}) = \bar{b}_\varepsilon(w_t)$. Let $\mathbf{X}_t^{(1)}$ and $\mathbf{X}_t^{(2)}$ the unique strong solution of $d\mathbf{X}_t = b(t, (\mathbf{X}_t)_{t \in [0,1]}) + \sqrt{2}\mathbf{B}_t$ with $\mathbf{X}_0 = x$ with $x \in \mathbb{R}^d$ and $b = b_1$, respectively $b = b_2$. Note that $(\mathbf{X}_t^{(2)})_{t \geq 0} = (\bar{\mathbf{X}}_t)_{t \geq 0}$ and $(\mathbf{X}_{k/m}^{(1)}) = (X_k)_{k \in \mathbb{N}}$.

For any $i \in \{1, 2\}$, denote $P_t^{(i)}$ the Markov semigroup associated with $\mathbf{X}_t^{(i)}$. For any $x \in \mathbb{R}^d$ we have

$$\left\| \delta_x R_{\varepsilon, 1/m, C}^m - \delta_x P_{1, \varepsilon, C} \right\|_{TV} = \left\| \delta_x P_1^{(1)} - \delta_x P_1^{(2)} \right\|_{TV}. \quad (48)$$

Using **H3(R)** and the fact that for any $a, b \geq 0$, $(a + b)^2 \leq 2(a^2 + b^2)$, we have for any $t \in [j/m, (j+1)/m]$, $j \in \{0, \dots, m-1\}$ and $(w_t)_{t \in [0, 1]} \in C([0, 1], \mathbb{R}^d)$

$$\begin{aligned} \left\| b_1(t, (w_t)_{t \in [0, 1]}) - b_2(t, (w_t)_{t \in [0, 1]}) \right\|^2 &= \left\| b_\varepsilon(w_{j/m}) - \bar{b}_\varepsilon(w_t) \right\|^2 \\ &\leq 2 \left\| b_\varepsilon(w_{j/m}) - b_\varepsilon(w_t) \right\|^2 + 2 \left\| \bar{b}_\varepsilon(w_t) - b_\varepsilon(w_t) \right\|^2 \\ &\leq 2L_b^2 \left\| w_{j/m} - w_t \right\|^2 + 4\alpha^2 M_R^2 / \varepsilon^2 + 4\alpha^2 \mathbf{1}_{\bar{B}(0, R)^c}(\|w_t\|) / \varepsilon^2, \end{aligned} \quad (49)$$

where L_b is the Lipschitz constant associated with b_ε . In addition using Itô's isometry we have for any $t \in [j/m, (j+1)/m]$

$$\mathbb{E}[\|\mathbf{X}_t^{(1)} - \mathbf{X}_{j/m}^{(1)}\|^2] = 2\mathbb{E}[\|\int_{j/m}^t d\mathbf{B}_t\|^2] \leq 2d\delta. \quad (50)$$

Finally, using Lemma 24, Lemma 25, the logarithmic Sobolev inequality [19, Theorem 5.5], the Cauchy-Schwarz inequality and the Markov inequality, there exists $\tilde{B}_b \geq 0$ such that for any $t \geq 0$ and $x \in \mathbb{R}^d$

$$\begin{aligned} \mathbb{P}(\|\mathbf{X}_t^{(1)}\| \geq R) &\leq \exp[-2R] \mathbb{E} \left[\exp[2\|\mathbf{X}_t^{(1)}\|] \right] \\ &\leq \exp[-2R] \mathbb{E}^{1/2} \left[\exp \left[4\sqrt{2} \left\| \int_{\ell_t/m}^t d\mathbf{B}_t \right\| \right] \right] \mathbb{E}^{1/2} \left[\exp[4\|X_{\ell_t}\|] \right] \\ &\leq \tilde{B}_b \exp[-2R] \exp[2\Phi(x)], \end{aligned}$$

where $\ell_t = \lfloor tm \rfloor$ and $\Phi(x) = \sqrt{1 + \|x\|^2}$. Combining this result, (49), (48), (50) and Lemma 26, we obtain that there exists $B_b \geq 0$ such that for any $x \in \mathbb{R}^d$ and $R > 0$

$$\begin{aligned} \left\| \delta_x R_{1/m, C}^m - \delta_x P_{1, C} \right\|_V &\leq 2B_b(\sqrt{\delta} + M_R + \exp[-R])(1 + \|x\|^4) \exp[\Phi(x)] \\ &\leq 48B_b(\sqrt{\delta} + M_R + \exp[-R]) \exp[2\Phi(x)], \end{aligned}$$

Combining this result and (47) we obtain that for any $k \in \mathbb{N}$, $j \in \{0, \dots, k-1\}$, $x \in \mathbb{R}^d$ and $R > 0$ we have

$$\left| (\delta_x R_{1/m, C}^m - \delta_x P_{1, C}) P_{k-j-1, C} [f] \right| \leq B_a B_b (\sqrt{\delta} + M_R + \exp[-R]) \tilde{\rho}_C^{k-j-1} \exp[2\Phi(x)].$$

Using this result, Lemma 24, Lemma 25 and (46) we obtain that there exists $B_c \geq 0$ such that for any $m \in \mathbb{N}^*$ with $m^{-1} \geq \bar{\delta}$

$$\left\| \pi_{\varepsilon, 1/m, C} - \tilde{\pi}_\varepsilon \right\|_V \leq \limsup_{k \rightarrow +\infty} \left\| \delta_0 R_{\varepsilon, 1/m, C}^{km} - \delta_0 P_{km, \varepsilon, C}^{km} \right\|_V \leq B_c(\sqrt{\delta} + M_R + \exp[-R]).$$

The proof in the general case where $\delta \in (0, \bar{\delta}]$ is similar and we obtain that there exists $B_c \geq 0$ such that for any $\delta \in (0, \bar{\delta}]$

$$\left\| \pi_{\varepsilon, \delta} - \tilde{\pi}_\varepsilon \right\|_V \leq B_c(\sqrt{\delta} + M_R + \exp[-R]).$$

(b) For any C compact convex with $0 \in C$ we define $\tilde{\pi}_\varepsilon$ and $\rho_{\varepsilon,C}$ such that for any $x \in \mathbb{R}^d$

$$\rho_{\varepsilon,C}(x) = \exp[-d^2(x, C)/(2\lambda)]q_y(x)p_\varepsilon^\alpha(x), \quad (d\tilde{\pi}_\varepsilon/d\text{Leb})(x) = \rho_{\varepsilon,C}(x) \Big/ \int_{\mathbb{R}^d} \rho_{\varepsilon,C}(\tilde{x})d\tilde{x} .$$

Similarly, define ρ_ε and π_ε such that for any $x \in \mathbb{R}^d$

$$\rho_\varepsilon(x) = q_y(x)p_\varepsilon^\alpha(x), \quad (d\pi_\varepsilon/d\text{Leb})(x) = \rho_\varepsilon(x) \Big/ \int_{\mathbb{R}^d} \rho_\varepsilon(\tilde{x})d\tilde{x} .$$

Since for any $x \in \mathbb{R}^d$, $\rho_{\varepsilon,C}(x) \leq \rho_\varepsilon(x)$ we get $\int_{\mathbb{R}^d} \rho_{\varepsilon,C}(\tilde{x})d\tilde{x} \leq \int_{\mathbb{R}^d} \rho_\varepsilon(\tilde{x})d\tilde{x}$. Hence we obtain using the Cauchy-Schwarz inequality and the Markov inequality

$$\begin{aligned} \text{KL}(\pi_\varepsilon|\pi_C) &\leq \int_{\mathbb{R}^d} \log(\rho_\varepsilon(\tilde{x})/\rho_{\varepsilon,C}(\tilde{x}))d\pi_\varepsilon(\tilde{x}) \\ &\leq \int_{C^c} \|\tilde{x}\|^2 d\pi_\varepsilon(\tilde{x}) \leq \mathbb{P}^{1/2}(X \notin C) \mathbb{E}^{1/2}[\|X\|^4] \leq \mathbb{E}[\|X\|^4]R_C^{-2} . \end{aligned}$$

with X a random variable with distribution π_ε . We conclude using the generalized Pinsker inequality [37, Lemma 24].

□

C.4 Proof of Proposition 17

Let $\alpha, \lambda, \varepsilon, \bar{\delta} > 0$, $\delta \in (0, \bar{\delta}]$ and $C \subset \mathbb{R}^d$ convex and compact with $0 \in C$. For any $x_1, x_2 \in \mathbb{R}^d$ we have

$$\|b_\varepsilon(x_1) - b_\varepsilon(x_2)\| \leq (L_y + \alpha L/\varepsilon) \|x_1 - x_2\| .$$

Denote $(X_n, Y_n)_{n \in \mathbb{N}}$ the Markov chain obtained using the coupling described in [15, Section 3] with initial condition $(x_1, x_2) \in C$. Using [15, Corollary 7-(b)] we get that for any $\ell \in \mathbb{N}$

$$\mathbb{E} \left[\mathbf{1}_{\Delta_{\mathbb{R}^d}^c}(X_{(\ell+1)\lceil 1/\delta \rceil}, Y_{(\ell+1)\lceil 1/\delta \rceil}) \right] \leq (1 - \beta) \mathbb{E} \left[\mathbf{1}_{\Delta_{\mathbb{R}^d}^c}(X_{\ell\lceil 1/\delta \rceil}, Y_{\ell\lceil 1/\delta \rceil}) \right] , \quad (51)$$

where $\Delta_{\mathbb{R}^d} = \{(x, x) : x \in \mathbb{R}^d\}$ and $\beta \in (0, 1)$ with

$$\beta = 2\Phi\{-(1 + \bar{\delta})(1 + L_y + (\alpha L/\varepsilon))\text{diam}(C)\} ,$$

where Φ is the cumulative distribution function of the univariate Gaussian distribution with zero mean and unit variance. In addition, using that the coupling is absorbing, we have that for any $k \in \mathbb{N}$,

$$\mathbb{E} \left[\mathbf{1}_{\Delta_{\mathbb{R}^d}^c}(X_k, Y_k) \right] \leq \mathbb{E} \left[\mathbf{1}_{\Delta_{\mathbb{R}^d}^c}(X_{\lfloor k/\lceil 1/\delta \rceil \rceil \lceil 1/\delta \rceil}, Y_{\lfloor k/\lceil 1/\delta \rceil \rceil \lceil 1/\delta \rceil}) \right] ,$$

Combining this result and (51), we get that for any $k \in \mathbb{N}$

$$\|\delta_{x_1} Q_{\varepsilon, \delta}^k - \delta_{x_2} Q_{\varepsilon, \delta}^k\|_{\text{TV}} \leq \mathbb{E} \left[\mathbf{1}_{\Delta_{\mathbb{R}^d}^c}(X_k, Y_k) \right] \leq (1 - \beta)^{\lfloor k/\lceil 1/\delta \rceil \rceil} .$$

Using that $\lfloor k/\lceil 1/\delta \rceil \rceil \geq k\delta/(1 + \delta) - 1$ concludes the proof upon letting $\tilde{\rho}_C = (1 - \beta)^{1/(1 + \bar{\delta})}$ and $\tilde{A}_C = (1 - \beta)^{-1}$.

C.5 Proof of Proposition 18

Let $\alpha, \lambda > 0, \varepsilon \in (0, \varepsilon_0]$ such that $2\lambda(\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon - \min(\mathbf{m}, 0)) \leq 1$ and $\bar{\delta}_1 = (1/3)(\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon + 1/\lambda)^{-1}$. Recall that for any $x_1, x_2 \in \mathbb{R}^d$

$$\|b_\varepsilon(x_1) - b_\varepsilon(x_2)\| \leq (\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon + 1/\lambda) \|x_1 - x_2\| .$$

Using this result, the fact that for any $x \in \mathbb{R}^d$, $\langle b_\varepsilon(x), x \rangle \leq -\tilde{\mathbf{m}} \|x\|^2 + c$ and [36, Theorem 19.4.1] there exist $\bar{\delta}_2 > 0, \tilde{B} \geq 0$ and $\tilde{\rho} \in (0, 1]$ such that for any $\delta \in (0, \bar{\delta}_2], x \in \mathbb{R}^d$ and $k \in \mathbb{N}$

$$\|\delta_x \mathbf{R}_{\varepsilon, \delta}^k - \pi_{\varepsilon, \delta}\|_V + \|\delta_x \mathbf{Q}_{\varepsilon, \delta}^k - \pi_{\varepsilon, \delta}^C\|_V \leq \tilde{B} \tilde{\rho}^{k\delta} V(x) ,$$

with \tilde{B} and $\tilde{\rho}$ which do not depend on R . In addition, using Lemma 24, for any $k \in \mathbb{N}$ and $\delta \in (0, \bar{\delta}_2]$ we have

$$\mathbf{R}_{\varepsilon, \delta}^k V(x) \leq \tilde{\lambda}^{k\delta} V(x) + \tilde{c}\delta ,$$

with $\tilde{\lambda} \in [0, 1)$ and $\tilde{c} > 0$ which do not depend on $R \geq 0$. For any $\delta \in (0, \bar{\delta}_2]$ we have

$$\lambda^\delta + c\delta \leq \lambda^\delta (1 + c\delta \lambda^{-\bar{\delta}_2}) \leq (\lambda \exp[c\lambda^{-\bar{\delta}_2}])^\delta .$$

Let $A = \lambda \exp[c\lambda^{-\bar{\delta}_2}]$, we have that for any $x \in \mathbb{R}^d$, $\mathbf{R}_{\varepsilon, \delta} V(x) \leq A^\delta V(x)$. Therefore we get that $(V(X_n)A^{-n})_{n \in \mathbb{N}}$ is a supermartingale. Hence using Doob maximal inequality and Markov inequality we get that

$$\mathbb{P} \left(\sup_{k \in \{0, \dots, n\}} \|X_k\| \geq R \right) \leq V(x) A^{n\delta} \exp[-R] .$$

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

$$\|\pi_{\varepsilon, \delta} - \pi_{\varepsilon, \delta}^C\|_{TV} \leq (V(0) + \tilde{c}\bar{\delta}_2) A^{k\delta} \exp[-R] + \tilde{B} \tilde{\rho}^{k\delta} V(0) .$$

We conclude upon letting $k = \lfloor r/(2 \log(A)\delta) \rfloor$.

D Proofs of Section 5.3

D.1 Proof of Proposition 19

The first part of the proposition is straightforward. Using Pinsker's inequality [19, Theorem 4.19] we have for any $x \in \mathbb{R}^d$

$$\|\mu - (\tau_x) \# \mu\|_{TV}^2 \leq 2\text{KL}((\tau_x) \# \mu | \mu) \leq 2 \int_{\mathbb{R}^d} \|U(\tilde{x} + x) - U(\tilde{x})\| d\mu(\tilde{x}) \leq 2C_\gamma \|x\|^\gamma .$$

For the second part of the proof, since there exist $c_1, \varpi > 0$ and $c_2 \in \mathbb{R}$ such that for any $x \in \mathbb{R}^d$, $U(x) \geq c_1 \|x\|^\varpi + c_2$ then for any $k \in \mathbb{N}^*$ and $\alpha > 0$, $\int_{\mathbb{R}^d} (1 + \|x\|)^k p(x) < +\infty$. Let $q(x) = (1 + \|x\|)^{-(d+1)} / \int_{\mathbb{R}^d} (1 + \|\tilde{x}\|)^{-(d+1)} d\tilde{x}$. Then using that for any $t \geq 0$, $|e^t - 1| \leq |t| e^{|t|}$ we get that for any $x \in \mathbb{R}^d$

$$\begin{aligned} & \int_{\mathbb{R}^d} |p(\tilde{x}) - p(x - \tilde{x})| q^{1-1/\alpha}(\tilde{x}) d\tilde{x} \\ & \leq C_\gamma \|x\|^\gamma \exp[C_\gamma \|x\|^\gamma] \int_{\mathbb{R}^d} (1 + \|\tilde{x}\|)^{(d+1)(1/\alpha-1)} p(\tilde{x}) d\tilde{x} \left(\int_{\mathbb{R}^d} (1 + \|\tilde{x}\|)^{-(d+1)} d\tilde{x} \right)^{1-1/\alpha} , \end{aligned}$$

which concludes the proof.

D.2 Proof of Proposition 20

First we show the following technical lemma.

Lemma 28. *For any $x, y \geq 0$ and $\beta > 0$, $(x + y)^\beta - x^\beta \leq 2^\beta(y^\beta + x^{(\beta-1) \wedge 0}y)$.*

Proof. The result is straightforward if $\beta \in (0, 1]$, since in this case $(x + y)^\beta \leq x^\beta + y^\beta$. Assume that $\beta > 1$. If $x = 0$ the result holds. Now assume that $x > 0$. If $y \geq x$ then $(x + y)^\beta - x^\beta \leq 2^\beta y^\beta$. Assume that $y \leq x$. Since $f : t \mapsto (1 + t)^\beta - 1$ is convex we obtain that for any $t \in [0, 1]$, $f(t) \leq 2^\beta t$. Using this result we have

$$(x + y)^\beta - x^\beta \leq x^\beta f(y/x) \leq 2^\beta x^{\beta-1} y,$$

which concludes the proof. \square

Before proving Proposition 20 we state the following lemma.

Lemma 29. *Let π_1, π_2 two probability measures and $q_1, q_2 : \mathbb{R}^d \rightarrow [0, +\infty)$ two measurable functions such that for any $x \in \mathbb{R}^d$, $(d\pi_i/d\text{Leb})(x) = q_i(x)/c_i$ with $c_i = \int_{\mathbb{R}^d} q_i(\tilde{x})d\tilde{x}$. Denote $D = \int_{\mathbb{R}^d} |q_1(x) - q_2(x)|$. We have*

$$\|\pi_1 - \pi_2\|_{\text{TV}} \leq 2c_1^{-1}D.$$

Proof. We have

$$\|\pi_1 - \pi_2\|_{\text{TV}} = \int_{\mathbb{R}^d} \left| \frac{q_1(x)}{c_1} - \frac{q_2(x)}{c_2} \right| dx \leq c_1^{-1}(D + |c_2 - c_1|),$$

which concludes the proof using that $|c_2 - c_1| \leq D$. \square

We now give the proof of Proposition 20.

Proof. Let $\alpha > 0$. For any $\varepsilon > 0$ and $x \in \mathbb{R}^d$ denote $\bar{p}(x) = q_y(x)p^\alpha(x)$ and $\bar{p}_\varepsilon(x) = q_y(x)p_\varepsilon^\alpha(x)$, where we recall that for any $x \in \mathbb{R}^d$

$$p_\varepsilon(x) = (2\pi\varepsilon)^{-d/2} \int_{\mathbb{R}^d} p(\tilde{x}) \exp[-\|x - \tilde{x}\|^2/(2\varepsilon)]d\tilde{x}.$$

For any $\varepsilon > 0$ we have

$$\int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx \leq \|q_y\|_\infty \int_{\mathbb{R}^d} |p^\alpha(x) - p_\varepsilon^\alpha(x)| dx.$$

Using Lemma 28 and that $\|p_\varepsilon\|_\infty \leq \|p\|_\infty < +\infty$, we have for any $\varepsilon > 0$ and $x \in \mathbb{R}^d$

$$\begin{aligned} \int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx &\leq 2^\alpha \|q_y\|_\infty (1 + \|p\|_\infty^{(\alpha-1) \wedge 0}) \\ &\times \left\{ \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| dx + \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)|^\alpha dx \right\}. \end{aligned} \quad (52)$$

Using Jensen's inequality, for any $q : \mathbb{R}^d \rightarrow (0, +\infty)$ with $\int_{\mathbb{R}^d} q(\tilde{x})d\tilde{x} = 1$ we have

$$\int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)|^\alpha dx \leq \left(\int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)q^{1-1/\alpha}(x)| dx \right)^\alpha.$$

Combining this result with (52) we get that

$$\int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx \leq 2^\alpha \|q_y\|_\infty (1 + \|p\|_\infty^{(\alpha-1)\wedge 0}) \times \left\{ \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| dx + \left(\int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| q^{1-1/\alpha}(x) dx \right)^\alpha \right\} .$$

If $\alpha \geq 1$, choosing q such that $\|q\|_\infty \leq 1$ we get

$$\int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx \leq 2^\alpha \|q_y\|_\infty (1 + \|p\|_\infty^{(\alpha-1)\wedge 0}) \times \left\{ \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| dx + \left(\int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| (x) dx \right)^\alpha \right\} . \quad (53)$$

Hence since $p \in L^1(\mathbb{R}^d)$ and $\{\tilde{x} \mapsto (2\pi\varepsilon)^{-d/2} \exp[-\|\tilde{x}\|^2/(2\varepsilon)] : \varepsilon > 0\}$ is a family of mollifiers, we have $\lim_{\varepsilon \rightarrow 0} \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| dx = 0$. Combining this result, (53) and Lemma 29 concludes the first part of the proof.

Now let $\alpha > 0$ and assume **H6**(α). If $\alpha \geq 1$ then using (52) we have

$$\int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx \leq 2^\alpha (1 + 2^{\alpha-1}) \|q_y\|_\infty (1 + \|p\|_\infty^{(\alpha-1)\wedge 0}) \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| dx .$$

If $\alpha < 1$ then using that $\|q\|_\infty < +\infty$, we get that

$$\int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx \leq 2^\alpha \|q_y\|_\infty (1 + \|q\|_\infty^{1/\alpha-1}) (1 + \|p\|_\infty^{(\alpha-1)\wedge 0}) \times \left\{ \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| q^{1-1/\alpha}(x) dx + \left(\int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| q^{1-1/\alpha}(x) dx \right)^\alpha \right\} .$$

Hence, in any case, there exists $\tilde{C}_0 \geq 0$ such that

$$\int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| \leq \tilde{C}_0 \left\{ \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| q^{\min(1-1/\alpha, 0)}(x) dx + \left(\int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| q^{\min(1-1/\alpha, 0)}(x) dx \right)^\alpha \right\} .$$

Using Jensen's inequality and the change of variable $\tilde{x} \mapsto \varepsilon^{1/2} \tilde{x}$, we have for any $\varepsilon \in (0, (4\kappa)^{-1}]$

$$\begin{aligned} & \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| q^{\min(1-1/\alpha, 0)}(x) dx \\ & \leq \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} |p(x) - p(x - \tilde{x})| q^{\min(1-1/\alpha, 0)}(x) (2\pi\varepsilon)^{-d/2} \exp[-\|\tilde{x}\|^2/(2\varepsilon)] dx d\tilde{x} \\ & \leq \int_{\mathbb{R}^d} \exp[\kappa \|\tilde{x}\|^2] \|\tilde{x}\|^\beta (2\pi\varepsilon)^{-d/2} \exp[-\|\tilde{x}\|^2/(2\varepsilon)] d\tilde{x} \\ & \leq \varepsilon^{\beta/2} (2\pi)^{-d/2} \int_{\mathbb{R}^d} \exp[\kappa \|\tilde{x}\|^2] \|\tilde{x}\|^\beta \exp[-\|\tilde{x}\|^2/2] d\tilde{x} \\ & \leq \varepsilon^{\beta/2} (2\pi)^{-d/2} \int_{\mathbb{R}^d} \|\tilde{x}\|^\beta \exp[-\|\tilde{x}\|^2/4] d\tilde{x} \leq C_0 \varepsilon^{\beta/2} , \end{aligned}$$

with $C_0 = (2\pi)^{-d/2} \int_{\mathbb{R}^d} \|\tilde{x}\|^\beta \exp[-\|\tilde{x}\|^2/4] d\tilde{x}$. Hence, we have

$$\int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx \leq C_1(\varepsilon^{\beta/2} + \varepsilon^{\beta\alpha/2}), \quad (54)$$

with $C_1 = \tilde{C}_0(C_0 + C_0^\alpha)$. Let $\varepsilon_1 = \min((cC_1)^{-2/\beta}/2, (cC_1)^{-2/(\beta\alpha)}/2, (4\kappa)^{-1})$ and $c = \int_{\mathbb{R}^d} \bar{p}(x) dx$. Combining (54) with Lemma 29, we get that for any $\varepsilon \in (0, \varepsilon_1]$

$$\|\pi - \pi_\varepsilon\|_{\text{TV}} \leq 2c^{-1}C_1(\varepsilon^{\beta/2} + \varepsilon^{\beta\alpha/2}),$$

which concludes the proof upon letting $A_0 = 2c^{-1}C_1$. \square

E Proofs of Section 5.4

E.1 Proof of Proposition 22

Let $\varepsilon > 0$. For any $k \in \mathbb{N}$, let $\zeta_k = Z_{k+1}$ and $\eta_k = b_\varepsilon(X_k) + \nabla \log q_y(X_k) + \nabla \log p_\varepsilon(X_k)$. For any $k \in \mathbb{N}$ using **H3** we have $\|b(X_k) - \nabla \log p_\varepsilon(X_k) - \nabla \log p_\varepsilon(X_k)\| = \varepsilon^{-1} \|D_\varepsilon(X_k) - D_\varepsilon^*(X_k)\| \leq M_R$. Hence, we obtain that [80, Assumption 2.1, Assumption 2.2] are satisfied. In what follows, we show that [80, Assumption 2.3.c] holds. Denoting $G_\varepsilon(x) = (2\pi\varepsilon)^{-d/2} \exp[-\|x\|^2/(2\varepsilon)]$ we have that for any $x \in \mathbb{R}^d$, $p_\varepsilon(x) = (p * G_\varepsilon)(x)$, where $*$ denotes the convolution product. Since $p, G_\varepsilon \in L^1(\mathbb{R}^d)$ we get that for any $\xi \in \mathbb{R}^d$, $\widehat{p * G_\varepsilon}(\xi) = \hat{p}(\xi)\hat{G}_\varepsilon(\xi)$. Since $p \in L^1(\mathbb{R}^d)$, $\|\hat{p}\|_\infty < +\infty$ using Riemann-Lebesgue theorem and in addition $\hat{G}_\varepsilon(\xi) = \exp[-\varepsilon\|\xi\|^2/2]$. Hence, $\widehat{p * G_\varepsilon} \in L^1(\mathbb{R}^d)$ and we obtain that almost everywhere for any $x \in \mathbb{R}^d$

$$p_\varepsilon(x) = \int_{\mathbb{R}^d} \hat{p}(\xi)\hat{G}_\varepsilon(\xi) \exp[i\langle x, \xi \rangle] d\xi.$$

Denote $\bar{p}_\varepsilon : \mathbb{C}^d \rightarrow \mathbb{C}$ give for any $z \in \mathbb{C}^d$ by $\bar{p}_\varepsilon(z) = \int_{\mathbb{R}^d} \hat{p}(\xi)\hat{G}_\varepsilon(\xi) \exp[i\langle z, \xi \rangle] d\xi$. We have that \bar{p}_ε is analytic. Since for any $x \in \mathbb{R}^d$, $p_\varepsilon(x) > 0$ and $\bar{p}_\varepsilon \in C(\mathbb{C}^d, \mathbb{C})$, there exists an open set $\mathbf{U} \subset \mathbb{C}^d$ such that for any $z \in \mathbf{U}$, $\Re(\bar{p}_\varepsilon(z)) > 0$. Since $\log : \mathbb{C} \setminus (\{t \in \mathbb{C} : \Re(t) \leq 0\}) \rightarrow \mathbb{C}$ is analytic we obtain that $z \mapsto \log(\bar{p}_\varepsilon)$ is analytic on \mathbf{U} . Hence, $x \mapsto \log(q_y(x)) + \log(p_\varepsilon(x))$ is real-analytic on \mathbb{R}^d . We conclude using [80, Theorem 2.1].

E.2 Proof of Proposition 23

Let $(\varepsilon_n)_{n \in \mathbb{N}} \in \mathbf{E}$ such that $\lim_{n \rightarrow +\infty} S_{K, \varepsilon_n} = S$. Let $x^* \in S$. For any $\eta > 0$ there exists $\varepsilon_\eta \in \mathbb{N}$ such that for any $n \in \mathbb{N}$, $S \subset S_{K, \varepsilon_n} + \bar{\mathbf{B}}(0, \eta)$. Hence, for any $n \in \mathbb{N}^*$ there exist an increasing sequence $(k_n)_{n \in \mathbb{N}} \in \mathbb{N}^{\mathbb{N}}$ and $x_n \in S_{K, \varepsilon_{k_n}}$ and $z_n \in \bar{\mathbf{B}}(0, 1/n)$ such that $x^* = x_{k_n} + z_n$. Up to extraction, we can assume that $\lim_{n \rightarrow +\infty} x_{k_n} = x^*$ since for any $n \in \mathbb{N}$, $x_{k_n} \in K$.

In what follows, we show that $x^* \in S_K$. For any $n \in \mathbb{N}$ we have using the dominated convergence theorem

$$\nabla \log(p_{\varepsilon_{k_n}}(x_{k_n})) = (\nabla p_{\varepsilon_{k_n}}/p_{\varepsilon_{k_n}})(x_{k_n}) = (((\nabla p) * G_{\varepsilon_{k_n}})/(p * G_{\varepsilon_{k_n}}))(x_{k_n}),$$

where for any $x \in \mathbb{R}^d$ and $\varepsilon > 0$, $G_\varepsilon(x) = (2\pi\varepsilon)^{-d/2} \exp[-\|x\|^2/(2\varepsilon)]$. Using that for any $t \in \mathbb{R}$, $|1 - e^t| \leq |t|e^{|t|}$ we get that for any $x_1, x_2 \in \mathbb{R}^d$

$$|p(x_1) - p(x_2)| \leq |1 - \exp[U(x_1) - U(x_2)]|/c \leq \mathbf{N} \|x_1 - x_2\| \exp[\mathbf{N} \|x_1 - x_2\|]/c,$$

where $c = \int_{\mathbb{R}^d} \exp[-U(\tilde{x})] d\tilde{x}$. Similarly we have for any $x_1, x_2 \in \mathbf{K}$

$$\begin{aligned} \|\nabla p(x_1) - \nabla p(x_2)\| &\leq c^{-1} (\|\nabla U(x_1)\| |\exp[-U(x_1)] - \exp[-U(x_2)]| + \|\nabla U(x_1) - \nabla U(x_2)\|) \\ &\leq c^{-1} \mathbf{N} (1 + \sup_{\tilde{x} \in \mathbf{K}} \|\nabla U(\tilde{x})\| \exp[\mathbf{N} \|x_1 - x_2\|]) \|x_1 - x_2\| . \end{aligned}$$

Therefore, there exists $\mathbf{N}_{\mathbf{K}} \geq 0$ such that for any $x_1, x_2 \in \mathbf{K}$ we have

$$\max(|p(x_1) - p(x_2)|, \|\nabla p(x_1) - \nabla p(x_2)\|) \leq \mathbf{N}_{\mathbf{K}} \|x_1 - x_2\| .$$

Using this result, we obtain that $\lim_{n \rightarrow +\infty} |p(x_{k_n}) - p(x^*)| = 0$ and $\lim_{n \rightarrow +\infty} \|\nabla p_{\varepsilon_{k_n}}(x_{k_n}) - \nabla p(x^*)\| = 0$. Hence we obtain that $\lim_{n \rightarrow +\infty} \nabla \log(p(x_{k_n})) = \nabla \log p(x^*)$ since $p(x^*) > 0$. We also have that

$$\lim_{\varepsilon \rightarrow 0} \max(|p - p_{\varepsilon_{k_n}}|_{\infty, \mathbf{K}}, \|\nabla p - \nabla p_{\varepsilon_{k_n}}\|_{\infty, \mathbf{K}}) = 0 ,$$

since $\{G_{\varepsilon} : \varepsilon > 0\}$ is a family of mollifiers and $p \in C^1(\mathbb{R}^d, \mathbb{R})$. Therefore, we obtain that $\lim_{n \rightarrow +\infty} \nabla \log(p_{\varepsilon_{k_n}}(x_{k_n})) = \nabla \log p(x^*)$. Finally, we obtain that

$$\lim_{n \rightarrow +\infty} \{\nabla \log(q_y(x_{k_n})) + \log(p_{\varepsilon_{k_n}}(x_{k_n}))\} = \nabla \log(q_y(x^*)) + \log(p(x^*)) = 0 .$$

Hence, $x^* \in \mathbf{S}_{\mathbf{K}}$ and therefore $\mathbf{S}_{\mathbf{K}}^* \subset \mathbf{S}_{\mathbf{K}}$.