

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

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

January 12, 2022

Abstract

Since the seminal work of Venkatakrishnan et al. [80] in 2013, *Plug & Play* (PnP) methods have become ubiquitous in Bayesian imaging. These methods derive estimators for inverse problems in imaging by combining an explicit likelihood function with a prior that is implicitly defined by an image denoising algorithm. In the case of optimisation schemes, some recent works guarantee the convergence to a fixed point, albeit not necessarily a maximum-a-posteriori Bayesian estimate. In the case of Monte Carlo sampling schemes for general Bayesian computation, to the best of our knowledge there is no known proof of convergence. Algorithm convergence issues aside, there are 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 efficient Bayesian computation schemes. This paper develops theory for Bayesian analysis and computation with PnP priors. We introduce PnP-ULA (Plug & Play Unadjusted Langevin Algorithm) for Monte Carlo sampling and minimum mean squared error estimation. Using recent results on the quantitative convergence of Markov chains, we establish detailed convergence guarantees for this algorithm 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 and meaningful from a frequentist viewpoint. PnP-ULA is demonstrated on several canonical problems such as image deblurring and inpainting, where it is used for point estimation as well as for uncertainty visualisation and quantification.

*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 acknowledges support by EPSRC grant EP/T007346/1. JD and AA acknowledge support from the French Research Agency through the PostProdLEAP project (ANR-19-CE23-0027-01). 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).

[†]These authors contributed equally

[‡]Université de Paris, MAP5 UMR 8145, F-75006 Paris, France

[§]Department of Statistics University of Oxford 24-29 St Giles OX1 3LB, Oxford United Kingdom

[¶]Institut Universitaire de France (IUF)

^{||}Centre Borelli, UMR 9010, École Normale Supérieure Paris-Saclay

^{**}School of Mathematical and Computer Sciences, Heriot-Watt University & Maxwell Institute for Mathematical Sciences, Edinburgh, United Kingdom

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 $y \in \mathbb{R}^m$ under some assumptions on their relationship. For example, many works consider observation models of the form $y = \mathbf{A}(x) + n$, where $\mathbf{A} : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is a degradation operator modelling deterministic instrumental aspects of the observation process, and n is an unknown (stochastic) noise term taking values in \mathbb{R}^m . 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.).

The estimation of x from y is usually 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. 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 *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 $F(x, y) = -\log p(y|x)$. 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 $x \in \mathbb{R}^d, y \in \mathbb{R}^m$

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

which underpins all inference about x given the observation y . Most imaging methods seek to derive estimators reaching some kind of consensus between prior and likelihood, as for instance the Minimum Mean Square Error (MMSE) or Maximum A Posteriori (MAP) estimators

$$\hat{x}_{\text{MAP}} = \arg \max_{x \in \mathbb{R}^d} p(x|y) = \arg \min_{x \in \mathbb{R}^d} \{F(x, y) + U(x)\} , \quad (1)$$

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

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 [72, 21, 58, 64]), Markov random fields [13], or learning-based priors like patch-based Gaussian or Gaussian mixture models [91, 87, 1, 79, 44]. 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, [64, 33, 69, 39, 24].

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)

$$d\mathbf{X}_t = \nabla \log p(\mathbf{X}_t|y) + \sqrt{2}d\mathbf{B}_t = \nabla \log p(y|\mathbf{X}_t) + \nabla \log p(\mathbf{X}_t) + \sqrt{2}d\mathbf{B}_t , \quad (3)$$

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 (3) has a

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 .

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

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 (3) in continuous time is generally not possible, we can use discrete time approximations of (3) 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 (3), given by $X_0 \in \mathbb{R}^d$ and the following recursion for all $k \in \mathbb{N}$

$$X_{k+1} = X_k + \delta \nabla \log p(y|X_k) + \delta \nabla \log p(X_k) + \sqrt{2\delta} Z_{k+1}, \quad (4)$$

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 $\delta > 0$ is a step-size which controls a trade-off between asymptotic accuracy and convergence speed [27, 32]. The approximation error involved in discretising (3) can be asymptotically removed at the expense of additional computation by combining (4) with a Metropolis-Hastings correction step, leading to the so-called Metropolis-adjusted Langevin Algorithm (MALA) [71].

When the prior density $p(x)$ is log-concave but not smooth, one can still use ULA by approximating the gradient of $U(x) = -\log p(x)$ in (4) by the gradient of the smooth Moreau-Yosida envelope $U_\lambda(x)$, given for any $x \in \mathbb{R}^d$ and $\lambda > 0$ by $\nabla U_\lambda(x) = \frac{1}{\lambda}(x - \text{prox}_U^\lambda(x))$.³ For example, one could use the Moreau-Yosida ULA [33], given by $X_0 \in \mathbb{R}^d$ and the following recursion for all $k \in \mathbb{N}$

$$X_{k+1} = X_k + \delta \nabla \log p(y|X_k) + \frac{\delta}{\lambda} [\text{prox}_U^\lambda(X_k) - X_k] + \sqrt{2\delta} Z_{k+1}. \quad (5)$$

Notice that prox_U^λ is equivalent to MAP denoising under the prior $p(x)$, for additive white Gaussian noise with noise variance λ . The *Plug & Play* ULA methods studied in this paper are closely related to (5), with a state-of-the-art Gaussian denoiser “plugged” in lieu of prox_U^λ . However, instead of approximating ∇U via a Moreau-Yosida envelope as above, we use Tweedie’s identity (7) relating ∇U to an MMSE denoiser (see Section 2.1).

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 [29, 88, 90, 37, 74, 36] or exploit the knowledge of \mathbf{A} in the network architecture via unrolled optimization techniques [41, 25, 28, 38]. 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 exist concerns regarding the stability of such approaches for general reconstruction problem [5, 4].

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 [6]. More

³Recall: The Moreau-Yosida envelope is defined as $U_\lambda(x) = \inf_{\tilde{x}} U(\tilde{x}) + \frac{1}{2\lambda} \|x - \tilde{x}\|^2$ and the proximal operator is defined as $\text{prox}_U^\lambda(x) = \arg \min_{\tilde{x} \in \mathbb{R}^d} U(\tilde{x}) + \frac{1}{2\lambda} \|x - \tilde{x}\|_2^2$.

precisely, using a denoising algorithm D_ε , *Plug & Play* approaches seek to derive an approximation of the gradient ∇U (called the Stein score) [11, 12] or prox_U [59, 89, 22, 48, 73], which can for instance be used within an iterative minimisation scheme to approximate \hat{x}_{MAP} , or within a Monte Carlo sampling scheme to approximate \hat{x}_{MMSE} [3, 42, 47]. To the best of our knowledge, the idea of leveraging a denoising algorithm to approximate the score ∇U within a iterative Monte Carlo scheme was first proposed in the seminal paper [3] in the context of generative modelling with denoising auto-encoders, where the authors present a Monte Carlo scheme that can be viewed as an approximate *Plug & Play* MALA. This scheme was recently combined with an expectation maximisation approach and applied to Bayesian inference for inverse problems in imaging in [42]. Similarly, the recent work [47] proposes to solve imaging inverse problems by using a *Plug & Play* stochastic gradient strategy that has close connections to an unadjusted version of the MALA scheme of [3]. While these approaches have shown some remarkable empirical performance, they rely on hybrid algorithms that are not always well 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 that is not a gradient mapping. These algorithms are better understood when interpreted as fixed-point algorithms seeking to reach a set of equilibrium equations between the denoiser and the data fidelity term [19]. Our understanding of the convergence properties of hybrid optimisation methods has advanced significantly recently [73, 86, 78, 45], but these questions remain largely unexplored in the context of stochastic Bayesian algorithms, to compute \hat{x}_{MMSE} or perform other forms of statistical inference.

The use of *Plug & Play* operators has also been investigated in the context of Approximate Message Passing (AMP) computation methods (see [30] for an introduction to AMP focused on compressed sensing and [2] for a survey on PnP-AMP in the context of magnetic resonance imaging), particularly for applications involving randomised forward operators where it is possible to characterise AMP schemes in detail (see, e.g., [9, 46, 60, 23]). This is an active area of research, and recent works have extended the approach to Vector AMP (VAMP) strategies and characterised their behaviour for a wider class of problems [35].

Approaches based on score matching techniques [76, 43] have also shown promising results recently [51, 50]. These methods are linked with *Plug & Play* approaches as they also estimate a Stein score. However, they do not rely on the asymptotic convergence of a diffusion, but instead aim at inverting a noising process stemming from an optimal transport problem [16]. The recent work [50] is particularly relevant in this context as it considers a range of imaging inverse problems, where it exploits the structure of the forward operator to perform posterior sampling in a coarse-to-fine manner. This also allows the use of multivariate step-sizes that are specific to each scale and ensure stability. However, to the best of our knowledge, the convergence properties of [50] have not been studied yet.

1.4 Contributions summary

This paper presents a formal framework for Bayesian analysis and computation with *Plug & Play* priors. We propose two *Plug & Play* ULAs, with detailed convergence guarantees under realistic assumptions on the denoiser used. We also study important questions regarding whether the underlying Bayesian models and estimators are well defined, well-posed, and have the basic regularity properties required to support efficient Bayesian computation schemes. We pay particular attention to denoisers based on deep neural networks, and report extensive numerical experiments with a specific neural network denoiser [73] shown to satisfy our convergence guarantees.

The remainder of the paper is organized as follows. Section 2 defines notation, introduces our framework for studying Bayesian inference methods with *Plug & Play* priors, and presents

two *Plug & Play* ULAs for Bayesian computation in imaging problems. This is then followed by a detailed theoretical analysis of *Plug & Play* Bayesian models and algorithms in Section 3. Section 4 demonstrates the proposed approach with experiments related to non-blind image deblurring and image inpainting, where we perform point estimation and uncertainty visualisation analyses, and report comparisons with the *Plug & Play* Stochastic Gradient Descent method of [55]. Conclusions and perspectives for future work are finally reported in Section 5.

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

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

This section presents a formal framework for Bayesian analysis and computation with *Plug & Play* priors. 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 . The Bayesian framework addresses this difficulty by using prior knowledge about the marginal distribution of x in order to reduce the uncertainty about $x|y$ and make the estimation problem well posed. In the Bayesian *Plug & Play* approach, instead of explicitly specifying the marginal distribution of x , we introduce prior knowledge about x by specifying an image denoising operator D_ε for recovering x from a noisy observation $x_\varepsilon \sim \mathcal{N}(x, \varepsilon \text{Id})$ with noise variance $\varepsilon > 0$. A case of particular relevance in this context is when D_ε is implemented by a neural network, trained by using a set of clean images $\{x'_i\}_{i=1}^N$.

A central challenge in the formalisation of Bayesian inference with *Plug & Play* priors is that the denoiser D_ε used is generally not directly related to a marginal distribution for x , so it is not possible to derive an explicit posterior for $x|y$ from D_ε . As a result, it is not clear that plugging D_ε into gradient-based algorithms such as ULA leads to a well-defined or convergent scheme that is targeting a meaningful Bayesian model.

To overcome this difficulty, in this paper we analyse *Plug & Play* Bayesian models through the prism of *M-complete* Bayesian modelling [10]. Accordingly, there exists a true -albeit unknown and intractable- marginal distribution for x and posterior distribution for $x|y$. If it were possible, basing inferences on these true marginal and posterior distributions would be optimal both in terms of point estimation and in terms of delivering Bayesian probabilities that are valid from a frequentist viewpoint. We henceforth use μ to denote this optimal prior distribution for x on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ - where $\mathcal{B}(\mathbb{R}^d)$ denotes the Borel σ -field of \mathbb{R}^d , and when μ admits a density w.r.t. the Lebesgue measure on \mathbb{R}^d , we denote it 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

$$p^*(x|y) = p(y|x)p^*(x) / \int_{\mathbb{R}^d} p(y|\tilde{x})p^*(\tilde{x})d\tilde{x}. \quad (6)$$

⁴ Unlike most Bayesian imaging approaches that operate implicitly in an *M-closed* manner and treat their postulated Bayesian models as true models (see [10] for more details), we explicitly regard p^* (or more precisely μ) as a fundamental property of the unknown x , and models used for inference as operational approximations of p^* specified by the practitioner (either analytically, algorithmically, or from training data). This distinction will be useful for using the oracle posterior (6) as a reference, and *Plug & Play* Bayesian algorithms based on a denoiser D_ε as

⁴Strictly speaking, the true likelihood $p^*(y|x)$ may also be unknown, this is particularly relevant in the case of blind or myopic inverse imaging problems. For simplicity, we restrict our experiments and theoretical development to the case where $p(y|x)$ represents the true likelihood. Generalizations of our approach to the blind or semi-blind setting are discussed, e.g. by [42] - formalising these generalisations is an important perspective for future work.

approximations to reference algorithms to perform inference w.r.t. p^* . The accuracy of the *Plug & Play* approximations will depend chiefly on the closeness between D_ε and an optimal denoiser D_ε^* derived from p^* that we define shortly.

In this conceptual construction, the marginal μ naturally depends on the imaging application considered. It 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. And 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 required⁵.

Notice that because μ is unknown, we cannot verify that $p^*(x|y)$ satisfies the basic desiderata for gradient-based Bayesian computation: i.e., $p^*(x|y)$ need not be proper and differentiable, with $\nabla \log p^*(x|y)$ Lipschitz continuous. To guarantee that gradient-based algorithms that target approximations of $p^*(x|y)$ are well defined by construction, we introduce a regularised oracle μ_ε obtained via the convolution of μ with a Gaussian smoothing kernel with bandwidth $\varepsilon > 0$. Indeed, by construction, μ_ε has a smooth proper density p_ε given for any $x \in \mathbb{R}^d$ and $\varepsilon > 0$ by

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

Equipped with this regularised marginal distribution, 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

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

which inherits the regularity properties required for gradient-based Bayesian computation when 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.

More precisely, Proposition 1 below establishes that the regularised prior $p_\varepsilon^*(x)$ and posterior $p_\varepsilon^*(x|y)$ are proper, smooth, and that they can be made arbitrarily close to the original oracle models $p^*(x)$ and $p^*(x|y)$ by reducing ε , with the approximation error vanishing as $\varepsilon \rightarrow 0$.

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 is postponed to Appendix H.2. □

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, $x \mapsto \nabla \log p_\varepsilon^*(x|y)$ might not be Lipschitz continuous and hence the Langevin SDE (3) might not have a strong solution. This requires an additional assumption on μ .

⁵Operating without densities requires measure disintegration concepts that are technical [75].

To study the Lipschitz continuity of $x \mapsto \nabla \log p_\varepsilon^*(x|y)$, as well as to set the grounds for *Plug & Play* methods that define priors implicitly through a denoising algorithm, we introduce the oracle MMSE denoiser D_ε^* defined for any $x \in \mathbb{R}^d$ and $\varepsilon > 0$ by

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

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]. Again, this optimal denoiser is a fundamental property of x and it is generally intractable. Motivated by the fact that state-of-the-art image denoisers are close-to-optimal in terms of MSE, in Section 2.3 we will characterise the accuracy of *Plug & Play* Bayesian methods for approximate inference w.r.t. $p_\varepsilon^*(x|y)$ and $p^*(x|y)$ as a function of the closeness between the denoiser D_ε used and the reference D_ε^* .

To relate the gradient $x \mapsto \nabla \log p_\varepsilon^*(x)$ and D_ε^* , we use Tweedie's identity [34] which states that for all $x \in \mathbb{R}^d$

$$\varepsilon \nabla \log p_\varepsilon^*(x) = D_\varepsilon^*(x) - x , \quad (7)$$

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^* is well-posed in the sense of Hadamard (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 [77]). As established in Proposition 2 below, this happens when the expected MSE involved in using D_ε^* to recover 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 $\varepsilon > 0$. $\nabla \log p_\varepsilon^*$ is Lipschitz continuous if and only if there exists $C \geq 0$ such that for any $x_\varepsilon \in \mathbb{R}^d$*

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

where $g_\varepsilon(\cdot|x_\varepsilon)$ is the density of the conditional distribution of the unknown image $x \in \mathbb{R}^d$ with marginal μ , given a noisy observation $x_\varepsilon \sim \mathcal{N}(x, \varepsilon \text{Id})$. See Section 3.2 for details.

Proof. The proof is postponed to Lemma 20. □

These results can be generalised to hold under the weaker assumption that the expected MSE for D_ε^* is finite but not uniformly bounded, as in this case $x \mapsto \nabla \log p_\varepsilon^*(x|y)$ is locally instead of globally Lipschitz continuous (we postpone this technical extension to future work). The pathological case where D_ε^* 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 summary, the gradient $x \mapsto \nabla \log p_\varepsilon^*(x|y)$ is Lipschitz continuous when μ carries enough information to make the problem of Bayesian image denoising under Gaussian additive noise well posed.

Notice that by using Tweedie's identity, we can express a ULA recursion for sampling approximately from $p_\varepsilon^*(x|y)$ as follows:

$$X_{k+1} = X_k + \delta \nabla \log p(y|X_k) + (\delta/\varepsilon) (D_\varepsilon^*(X_k) - X_k) + \sqrt{2\delta} Z_{k+1} . \quad (8)$$

where we recall that $\{Z_k : k \in \mathbb{N}\}$ are i.i.d standard Gaussian random variables on \mathbb{R}^d and $\delta > 0$ is a positive step-size. Under standard assumptions on δ , the sequence generated by (8) is a Markov chain which admits an invariant probability distribution whose density is provably close to $p_\varepsilon^*(x|y)$, with δ controlling a trade-off between asymptotic accuracy and convergence speed.

In the following section we present *Plug & Play* ULAs that arise from replacing D_ε^\star in (8) with a denoiser D_ε that is tractable.

Before concluding this section, we study whether the oracle $p^\star(x|y)$ is itself well-posed, i.e., if $p^\star(x|y)$ changes continuously w.r.t. y under a suitable probability metric (see [54]). 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 3 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 $y_1, y_2 \in \mathbb{R}^m$*

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

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

Proof. The proof is a straightforward application of Proposition 18. \square

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, with gradients that are explicitly related to denoising operators by Tweedie's formula. Under mild assumptions on $p(y|x)$, the approximation $p_\varepsilon^\star(x|y)$ is well-posed 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 imagining convergent gradient-based algorithms for performing Bayesian computation for $p_\varepsilon^\star(x|y)$, setting the basis for *Plug & Play* ULA schemes that mimic these idealised algorithms by using a tractable denoiser D_ε such as neural network, trained to optimise MSE performance and hence to approximate the oracle MSE denoiser D_ε^\star .

2.2 Bayesian computation with *Plug & Play* priors

We are now ready to study *Plug & Play* ULA schemes to perform approximate inference w.r.t. $p_\varepsilon^\star(x|y)$ (and hence indirectly w.r.t. $p^\star(x|y)$). We use (8) as starting point, with D_ε^\star replaced by a surrogate denoiser D_ε , but also modify (8) to guarantee geometrically fast convergence⁶ to a neighbourhood of $p_\varepsilon^\star(x|y)$. In particular, geometrically fast convergence is achieved here by modifying far-tail probabilities to prevent the Markov chain from becoming too diffusive as it explores the tails of $p_\varepsilon^\star(x|y)$. We consider two alternatives to guarantee geometric convergence with markedly different bias-variance trade-offs: one with excellent accuracy guarantees but that requires using a small step-size δ and hence has a higher computational cost, and another one that allows taking a larger step-size δ to improve convergence speed at the expense of weaker guarantees in terms of estimation bias.

First, in the spirit of Moreau-Yosida regularised ULA [33], we define *Plug & Play* ULA (PnP-ULA) as the following recursion: given an initial state $X_0 \in \mathbb{R}^d$ and for any $k \in \mathbb{N}$,

$$\begin{aligned} \text{(PnP-ULA)} \quad X_{k+1} = & X_k + \delta \nabla \log p(y|X_k) + (\delta/\varepsilon) (D_\varepsilon(X_k) - X_k) \\ & + (\delta/\lambda) (\Pi_C(X_k) - X_k) + \sqrt{2\delta} Z_{k+1} , \end{aligned}$$

⁶Geometric convergence is highly desirable property in large-scale problems and guarantees that the generated Markov chains can be used for Monte Carlo integration.

where $\mathbf{C} \subset \mathbb{R}^d$ is some large compact convex set that contains most of the prior probability mass of x , $\Pi_{\mathbf{C}}$ is the projection operator onto \mathbf{C} w.r.t the Euclidean scalar product on \mathbb{R}^d , and $\lambda > 0$ is a tail regularisation parameter that is set such that the drift in PnP-ULA satisfies a certain growth condition as $\|x\| \rightarrow \infty$ (see Section 3 for details).

An alternative strategy (which we call Projected PnP-ULA, *i.e.* PPnP-ULA, see Algorithm 2) is to modify PnP-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 $k \in \mathbb{N}$

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

where we notice that, by construction, the chain cannot exit \mathbf{C} because of the action of the projection operator $\Pi_{\mathbf{C}}$. The hard projection guarantees geometric convergence with weaker restrictions on δ and hence PPnP-ULA can be tuned to converge significantly faster than PnP-ULA, albeit with a potentially larger bias. These two schemes are summarised in Algorithm 1 and Algorithm 2 below. Note the presence of a regularisation parameter α in these algorithms, which permits to balance the weights between the prior and data terms. For the sake of simplicity, this parameter is set to $\alpha = 1$ in Section 3 and Section 4 but will be taken into account in the supplementary material Appendix A. Section 3.2 and Section 3.3 present detailed convergence results for PnP-ULA and PPnP-ULA. Implementation guidelines, including suggestions for how to set the algorithm parameters of PnP-ULA and PPnP-ULA are provided in Section 4.

Algorithm 1 PnP-ULA

Require: $n \in \mathbb{N}$, $y \in \mathbb{R}^m$, $\varepsilon, \lambda, \alpha, \delta > 0$, $\mathbf{C} \subset \mathbb{R}^d$ 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 $X_0 \in \mathbb{R}^d$ and $k = 0$.

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

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

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

end for

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

Algorithm 2 PPnP-ULA

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

Initialization: Set $X_0 \in \mathbf{C}$ and $k = 0$.

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

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

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

end for

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

Lastly, it is worth mentioning that Algorithm 1 and Algorithm 2 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 [33], analytical terms that are not smooth are involved via their proximal operator.

Before concluding this section, it is worth emphasising that, in addition to being important in their own right, Algorithm 1 and Algorithm 2 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.

3 Theoretical analysis

In this section, we provide a theoretical study of the long-time behaviour of PnP-ULA, see Algorithm 1 and PPnP-ULA, see Algorithm 2. For any $\varepsilon > 0$ we recall that p_ε^\star is given by the Gaussian smoothing of p with level ε , for any $x \in \mathbb{R}^d$ by

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

One typical example of likelihood function that we consider in our numerical illustration, see Section 4, is $p(y|x) \propto \exp[-\|\mathbf{A}x - y\|^2 / (2\sigma^2)]$ for any $x \in \mathbb{R}^d$ with $\sigma > 0$ and $\mathbf{A} \in \mathbb{R}^{m \times d}$. We define π the target posterior distribution given for any $x \in \mathbb{R}^d$ by $(d\pi/d\text{Leb})(x) = p^\star(x|y)$. We also consider the family of probability distributions $\{\pi_\varepsilon : \varepsilon > 0\}$ given for any $\varepsilon > 0$ and $x \in \mathbb{R}^d$ by

$$(d\pi_\varepsilon/d\text{Leb})(x) = p(y|x)p_\varepsilon^\star(x) \Big/ \int_{\mathbb{R}^d} p(y|\tilde{x})p_\varepsilon^\star(\tilde{x})d\tilde{x} .$$

Note that in the supplementary material Appendix A we investigate the general setting where p_ε^\star is replaced by $(p_\varepsilon^\star)^\alpha$ for some $\alpha > 0$ that acts as a regularisation parameter. We divide our study into two parts. We recall that π_ε is well-defined for any $\varepsilon > 0$ under **H1**, see Proposition 1. We start with some notation in Section 3.1. We then 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 3.2. Finally, in Section 3.3 we establish similar results for PPnP-ULA.

3.1 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, $\|f\|_\infty = \sup_{\tilde{x} \in \mathbb{R}^d} |f(\tilde{x})|$. 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 and $V : \mathbb{R}^d \rightarrow [1, \infty)$ measurable, the V -norm of f is given by $\|f\|_V = \sup_{\tilde{x} \in \mathbb{R}^d} |f(\tilde{x})|/V(\tilde{x})$. 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

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

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 $m \in \mathbb{N}$, $\mathcal{P}_m(\mathbb{R}^d) = \{\nu \in \mathcal{P}(\mathbb{R}^d) : \int_{\mathbb{R}^d} \|\tilde{x}\|^m d\nu(\tilde{x}) < +\infty\}$.

We denote by $C^k(\mathbf{U}, \mathbb{R}^m)$ and $C_c^k(\mathbf{U}, \mathbb{R}^m)$ 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 \mathbf{m} -convex with $\mathbf{m} \geq 0$ if for all $x_1, x_2 \in \mathbb{R}^d$ and $t \in [0, 1]$,

$$f(tx_1 + (1-t)x_2) \leq tf(x_1) + (1-t)f(x_2) - \mathbf{m}t(1-t)\|x_1 - x_2\|^2 / 2 .$$

For any $a \in \mathbb{R}^d$ and $R > 0$, denote $B(a, R)$ the open ball centered at a with radius R . Let (X, \mathcal{X}) and (Y, \mathcal{Y}) be two measurable spaces. A Markov kernel P is a mapping $K : X \times Y \rightarrow [0, 1]$ such that for any $\tilde{x} \in X$, $P(\tilde{x}, \cdot)$ is a probability measure and for any $A \in \mathcal{Y}$, $P(\cdot, A)$ is measurable. For any probability measure μ on (X, \mathcal{X}) and measurable function $f : Y \rightarrow \mathbb{R}_+$ we denote $\mu P = \int_X P(x, \cdot) d\mu(x)$ and $Pf = \int_Y f(y) P(\cdot, dy)$. In what follows the Dirac mass at $\tilde{x} \in \mathbb{R}^d$ is denoted by $\delta_{\tilde{x}}$. For any $\tilde{x} \in \mathbb{R}^d$, we denote $\tau_{\tilde{x}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ the translation operator given for any $\tilde{x}' \in \mathbb{R}^d$ by $\tau_{\tilde{x}}(\tilde{x}') = \tilde{x}' - \tilde{x}$. The complement of a set $A \subset \mathbb{R}^d$, is denoted by A^c . All densities are w.r.t. the Lebesgue measure (denoted Leb) unless stated otherwise. For all convex and closed set $C \subset \mathbb{R}^d$, we define Π_C the projection operator onto C w.r.t the Euclidean scalar product on \mathbb{R}^d . For any matrix $a \in \mathbb{R}^{d_1 \times d_2}$ with $d_1, d_2 \in \mathbb{N}$, we denote $a^\top \in \mathbb{R}^{d_2 \times d_1}$ its adjoint.

3.2 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 distance. 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 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

$$g_\varepsilon(x_1|x_2) = p^*(x_1) \exp[-\|x_2 - x_1\|^2/(2\varepsilon)] \Big/ \int_{\mathbb{R}^d} p^*(\tilde{x}) \exp[-\|x_2 - \tilde{x}\|^2/(2\varepsilon)] d\tilde{x} . \quad (9)$$

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 (with density p^*) 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 .

H2 (R). *We have that $\int_{\mathbb{R}^d} \|\tilde{x}\|^2 p^*(\tilde{x}) d\tilde{x} < +\infty$. In addition, 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*

$$\|(\text{Id} - D_\varepsilon)(x_1) - (\text{Id} - D_\varepsilon)(x_2)\| \leq L \|x_1 - x_2\| , \quad \|D_\varepsilon(x) - D_\varepsilon^*(x)\| \leq M_R , \quad (10)$$

where we recall that

$$D_\varepsilon^*(x_1) = \int_{\mathbb{R}^d} \tilde{x} g_\varepsilon(\tilde{x}|x_1) d\tilde{x} . \quad (11)$$

The Lipschitz continuity condition in (10) will be useful for establishing the stability and geometric convergence of the Markov chain generated by PnP-ULA. This condition can be explicitly enforced during training by using an appropriate regularization of the neural network weights [73, 62]. Regarding the second condition in (10), M_R is a bound on the error involved in using D_ε as an approximation of D_ε^* for images of magnitude R (i.e., for any $x \in \overline{B}(0, R)$), and it will be useful for bounding the bias resulting from using PnP-ULA for inference w.r.t. π_ε (recall that the bias vanishes as $M_R \rightarrow 0$ and $\delta \rightarrow 0$). For denoisers represented by neural networks, one can promote a small value of M_R during training by using an appropriate loss function. More precisely, 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 $\varepsilon > 0$, one could target empirical approximation of a loss of the form $\ell_\varepsilon : \mathcal{W} \rightarrow [0, +\infty)$ 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 [56].

With regards to the theoretical limitations stemming from representing D_ε by a deep neural network, universal approximation theorems (see e.g., [7, Section 4.7]) suggest that M_R could be arbitrarily low in principle. For a given architecture and training strategy, and if there exists $\tilde{M}_R \geq 0$ such that $\inf_{w \in \mathcal{W}} \sup_{x \in \bar{B}(0, R)} \tilde{M}_R^{-1} \|f_w(x) - D_\varepsilon^*(x)\| \leq 1$ then the second condition in (10) holds upon letting $D_\varepsilon = f_{w^\dagger}$ for an appropriate choice of weights $w^\dagger \in \mathcal{W}$. This last inequality can be established using universal approximation theorems such as [7, Section 4.7]. Moreover, for any other $w \in \mathcal{W}$, $\ell_\varepsilon(w) \geq \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - D_\varepsilon^*(x_\varepsilon)\|^2 p_\varepsilon^*(x_\varepsilon) g_\varepsilon(x|x_\varepsilon) dx dx_\varepsilon = \ell_\varepsilon^*$, 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 (11). Consider $w^\dagger \in \mathcal{W}$ obtained after numerically minimizing ℓ_ε and satisfying $\ell_\varepsilon(w^\dagger) \leq \ell_\varepsilon^* + \eta$ with $\eta > 0$. In this case, the following result ensures that (10) is satisfied with M_R of order $\eta^{1/(2d+2)}$ for any $R > 0$ and letting $D_\varepsilon = f_{w^\dagger}$.

Proposition 4. *Assume that for any $w \in \mathcal{W}$*

$$\int_{\mathbb{R}^d} (\|x\|^2 + \|f_w(x_\varepsilon)\|^2) p_\varepsilon^*(x_\varepsilon) g_\varepsilon(x|x_\varepsilon) dx dx_\varepsilon < +\infty. \quad (12)$$

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

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

where D_ε^ is given in (11). Then there exists $C_R, \bar{\eta}_R \geq 0$ such that if $\eta \in (0, \bar{\eta}_R]$ then for any $\tilde{x} \in \bar{B}(0, R)$, $\|f_{w^\dagger}(\tilde{x}) - D_\varepsilon^*(\tilde{x})\| \leq C_R \eta^{1/(2d+2)}$.*

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

Note that (12) is satisfied if for any $w \in \mathcal{W}$, $\sup_{x \in \mathbb{R}^d} \|f_w(x)\| (1 + \|x\|)^{-1} < +\infty$ and **H2** holds.

We recall that PnP-ULA, see Algorithm 1, is given by the following recursion: $X_0 \in \mathbb{R}^d$ and for any $k \in \mathbb{N}$

$$\begin{aligned} X_{k+1} &= X_k + \delta b_\varepsilon(X_k) + \sqrt{2\delta} Z_{k+1}, \\ b_\varepsilon(x) &= \nabla \log p(y|x) + P_\varepsilon(x) + (\text{prox}_\lambda(\iota_C)(x) - x)/\lambda, \quad P_\varepsilon(x) = (D_\varepsilon(x) - x)/\varepsilon, \end{aligned} \quad (13)$$

where $\delta > 0$ is a step-size, $\varepsilon, \lambda > 0$ 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 $\text{prox}_\lambda(\iota_C)(x) = \Pi_C(x)$, 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 (17) and given for any $x \in \mathbb{R}^d$ and $A \in \mathcal{B}(\mathbb{R}^d)$ by

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

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 $\lambda > 0$ and C .

Here we consider the case where $x \mapsto \log p(y|x)$ satisfies a one-sided Lipschitz condition, *i.e.* we consider the following condition.

H3. *There exists $\mathfrak{m} \in \mathbb{R}$ such that for any $x_1, x_2 \in \mathbb{R}^d$ we have*

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

We refer to the supplementary material Appendix C for refined convergence rates in the case where $x \mapsto \log p(y|x)$ is strongly \mathfrak{m} -concave. Note that if **H3** is satisfied with $\mathfrak{m} > 0$ then $x \mapsto \log p(y|x)$ is \mathfrak{m} -concave. Assume **H1** then **H3** 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 **H3**. In order to deal with **H3** in the case where $\mathfrak{m} \leq 0$, we set $\mathcal{C} \subset \mathbb{R}^d$ to be some convex compact set fixed by the user. Doing so, we ensure the stability of the Markov chain. The choice of \mathcal{C} in practice is discussed in Section 4. In our imaging experiments, we recall that for any $x \in \mathbb{R}^d$ we have, $p(y|x) \propto \exp[-\|\mathbf{A}x - y\|^2 / (2\sigma^2)]$. If \mathbf{A} is not invertible then $x \mapsto \log p(y|x)$ is not \mathfrak{m} -concave with $\mathfrak{m} > 0$. This is the case, in our deblurring experiment when the convolution kernel has zeros in the Fourier domain.

We start with the following result which ensures that the Markov chain (17) is geometrically ergodic under **H2** 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

$$V(x) = 1 + \|x\|^2 . \quad (14)$$

Proposition 5. *Assume **H1**, **H2**(R) for some $R > 0$ and **H3**. Let $\lambda > 0$, $\varepsilon \in (0, \varepsilon_0]$ such that $2\lambda(L_y + L/\varepsilon - \min(\mathfrak{m}, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + L/\varepsilon + 1/\lambda)^{-1}$. Then for any $\mathcal{C} \subset \mathbb{R}^d$ convex and compact with $0 \in \mathcal{C}$, there exist $A_{1,\mathcal{C}} \geq 0$ and $\rho_{1,\mathcal{C}} \in [0, 1)$ such that for any $\delta \in (0, \bar{\delta}]$, $x_1, x_2 \in \mathbb{R}^d$ and $k \in \mathbb{N}$ we have*

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

where V is given in (18).

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

The constants $A_{1,\mathcal{C}}$ and $\rho_{1,\mathcal{C}}$ do not depend on the dimension d but only on the parameters $\mathfrak{m}, L, L_y, \varepsilon$ and \mathcal{C} . 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 5 we have for any $\nu_1, \nu_2 \in \mathcal{P}_1(\mathbb{R}^d)$

$$\begin{aligned} \|\nu_1 R_{\varepsilon,\delta}^k - \nu_2 R_{\varepsilon,\delta}^k\|_V &\leq A_{1,\mathcal{C}} \rho_{1,\mathcal{C}}^{k\delta} \left(\int_{\mathbb{R}^d} V^2(\tilde{x}) d\nu_1(\tilde{x}) + \int_{\mathbb{R}^d} V^2(\tilde{x}) d\nu_2(\tilde{x}) \right) , \\ \mathbf{W}_1(\nu_1 R_{\varepsilon,\delta}^k, \nu_2 R_{\varepsilon,\delta}^k) &\leq A_{1,\mathcal{C}} \rho_{1,\mathcal{C}}^{k\delta} \left(\int_{\mathbb{R}^d} \|\tilde{x}\| d\nu_1(\tilde{x}) + \int_{\mathbb{R}^d} \|\tilde{x}\| d\nu_2(\tilde{x}) \right) . \end{aligned} \quad (15)$$

First, $(\mathcal{P}_1(\mathbb{R}^d), \mathbf{W}_1)$ is a complete metric space [81, 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 5. 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 (15), we have for any $\nu \in \mathcal{P}_1(\mathbb{R}^d)$

$$\begin{aligned} \|\nu R_{\varepsilon,\delta}^k - \pi_{\varepsilon,\delta}\|_V &\leq A_{1,\mathcal{C}} \rho_{1,\mathcal{C}}^{k\delta} \left(\int_{\mathbb{R}^d} V^2(\tilde{x}) d\nu(\tilde{x}) + \int_{\mathbb{R}^d} V^2(\tilde{x}) d\pi_{\varepsilon,\delta}(\tilde{x}) \right) , \\ \mathbf{W}_1(\nu R_{\varepsilon,\delta}^k, \pi_{\varepsilon,\delta}) &\leq A_{1,\mathcal{C}} \rho_{1,\mathcal{C}}^{k\delta} \left(\int_{\mathbb{R}^d} \|\tilde{x}\| d\nu(\tilde{x}) + \int_{\mathbb{R}^d} \|\tilde{x}\| d\pi_{\varepsilon,\delta}(\tilde{x}) \right) . \end{aligned}$$

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 $\sup_{x \in \mathbb{R}^d} \{(1 + \|x\|^2)^{-1} |h(x)|\} < +\infty$

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

where $(X_k)_{k \in \mathbb{N}}$ is the Markov chain given by (17) with starting point $X_0 = x \in \mathbb{R}^d$.

In the rest of this section we evaluate how close the invariant measure $\pi_{\varepsilon, \delta}$ is to π_ε . Our proof will rely on the following assumption which 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 (9).

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-UULA 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 6. Assume **H1**, **H2**(R) for some $R > 0$, **H3** and **H4**. Moreover, let $\varepsilon \in (0, \varepsilon_0]$ and assume that $\int_{\mathbb{R}^d} (1 + \|\tilde{x}\|^4) p_\varepsilon^*(\tilde{x}) d\tilde{x} < +\infty$. Let $\lambda > 0$ such that $2\lambda(L_y + (1/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon) - \min(m, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + L/\varepsilon + 1/\lambda)^{-1}$. 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 (18).

Proof. The proof is postponed to Appendix F.3. □

We now combine Proposition 5 and Proposition 6 in order to control the bias of the Monte Carlo estimator obtained using PnP-UULA. In the supplementary material Appendix D we also provide bounds on $|n^{-1} \sum_{k=1}^n \mathbb{E}[h(X_k)] - \int_{\mathbb{R}^d} h(\tilde{x}) d\pi(\tilde{x})|$ by controlling $\|\pi - \pi_\varepsilon\|_V$.

Proposition 7. Assume **H1**, **H2**(R) for some $R > 0$, **H3** and **H4**. Moreover, let $\varepsilon > 0$, $\varepsilon \in (0, \varepsilon_0]$ and assume that $\int_{\mathbb{R}^d} (1 + \|\tilde{x}\|^4) p_\varepsilon^*(\tilde{x}) d\tilde{x} < +\infty$. Let $\lambda > 0$ such that $2\lambda(L_y + (1/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon) - \min(m, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + L/\varepsilon + 1/\lambda)^{-1}$. Then there exists $C_{1, \varepsilon} > 0$ 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 $C_{2, \varepsilon}$ 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 5 and Proposition 6. \square

3.3 Convergence guarantees for PPnP-ULA

We now study the Projected *Plug & Play* Unadjusted Langevin Algorithm (PPnP-ULA). It is given by the following recursion: $X_0 \in \mathcal{C}$ and for any $k \in \mathbb{N}$

$$\begin{aligned} X_{k+1} &= \Pi_{\mathcal{C}}(X_k + \delta b_\varepsilon(X_k) + \sqrt{2\delta} Z_{k+1}) , \\ b_\varepsilon(x) &= \nabla \log p(y|x) + P_\varepsilon(x) , \quad P_\varepsilon(x) = (D_\varepsilon(x) - x)/\varepsilon , \end{aligned} \tag{16}$$

where $\delta > 0$ is a step-size, $\varepsilon > 0$ is an hyperparameter of the algorithm, $\mathcal{C} \subset \mathbb{R}^d$ 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 $\Pi_{\mathcal{C}}$ is the projection onto \mathcal{C} . In what follows, for any $\delta > 0$ and $\mathcal{C} \subset \mathbb{R}^d$ 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 (16) and given for any $x \in \mathbb{R}^d$ and $A \in \mathcal{B}(\mathbb{R}^d)$ by

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

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 \mathcal{C} .

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

Proposition 8. *Assume **H1**, **H2**(R) for some $R > 0$. Let $\lambda, \varepsilon, \bar{\delta} > 0$. Then for any $\mathcal{C} \subset \mathbb{R}^d$ convex and compact with $0 \in \mathcal{C}$, there exist $\tilde{A}_{\mathcal{C}} \geq 0$ and $\tilde{\rho}_{\mathcal{C}} \in [0, 1)$ such that for any $\delta \in (0, \bar{\delta}]$, $x_1, x_2 \in \mathcal{C}$ and $k \in \mathbb{N}$ we have*

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

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

In particular $Q_{\varepsilon,\delta}$ admits an invariant probability measure $\pi_{\varepsilon,\delta}^{\mathcal{C}}$. 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 \mathcal{C} has a large diameter.

Proposition 9. *Assume **H1**, **H2**(R) for some $R > 0$ and **H3**. In addition, assume that there exists $\tilde{m}, c > 0$ such that for $\mathcal{C} = \mathbb{R}^d$ and for any $\varepsilon > 0$ and $x \in \mathbb{R}^d$, $\langle b_\varepsilon(x), x \rangle \leq -\tilde{m} \|x\|^2 + c$. Let $\lambda > 0$, $\varepsilon \in (0, \varepsilon_0]$ such that $2\lambda(L_y + L/\varepsilon - \min(\tilde{m}, 0)) \leq 1$. Then there exist $\bar{A} \geq 0$ and $\eta, \bar{\delta} > 0$ such that for any $\mathcal{C} \subset \mathbb{R}^d$ convex and compact with $0 \in \mathcal{C}$ and $\bar{B}(0, R_{\mathcal{C}}/2) \subset \mathcal{C} \subset \bar{B}(0, R_{\mathcal{C}})$ and $\delta \in (0, \bar{\delta}]$ we have*

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

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

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

It is worth mentioning at this point that in our experiments, see Section 4, the probability of the iterates $(X_n)_{n \in \mathbb{N}}$ leaving \mathcal{C} with PnP-ULA or with PPnP-ULA is so low that the projection constraint is not activated. As a result, if implemented with the same step-size both algorithms produce the same results. We do not suggest completely removing the constraints as this is important to theoretically guarantee the geometric ergodicity of the algorithms.

Regarding the choice of the step-size, we observe that the bound $\bar{\delta} = (1/3)(L_y + L/\varepsilon + 1/\lambda)^{-1}$ used in PnP-ULA is conservative and our experiments suggest that PnP-ULA is stable for larger step-sizes.

4 Experimental study

This section illustrates the behaviour of PnP-ULA and PPnP-ULA with two classical imaging inverse problems: non-blind image deblurring and inpainting. For these two problems, we first analyse in detail the convergence of the Markov chain generated by PnP-ULA for different test images. This is then followed by a comparison between the MMSE Bayesian point estimator, as calculated by using PnP-ULA and PPnP-ULA and the MAP estimator provided by the recent PnP-SGD method [55]. We refer the reader to [55] for comparisons with PnP-ADMM [73]. To simplify comparisons, for all experiments and algorithms, the operator D_ε is chosen as the pretrained denoising neural network introduced in [73], for which $(D_\varepsilon - \text{Id})$ is L -Lipschitz with $L < 1$.

For the deblurring experiments, the observation model takes the form

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

where $x \in \mathbb{R}^d$ is the unknown original image, $y \in \mathbb{R}^m$ the observed image, n is a realization of a Gaussian i.i.d. centered noise with variance $\sigma^2 \text{Id}$ (with $\sigma^2 = (1/255)^2$), and \mathbf{A} is a 9×9 box blur operator. The log-likelihood for this case writes $\log p(y|x) = -\|\mathbf{A}x - y\|^2 / (2\sigma^2)$.

In the inpainting experiments, we seek to recover $x \in \mathbb{R}^d$ from $y = \mathbf{A}x$ where the matrix \mathbf{A} is a $m \times d$ matrix containing m randomly selected rows of the $d \times d$ identity matrix. We focus on a case where 80% of the image pixels are hidden and the observed pixels are measured without any noise. Because the posterior density for $x|y$ is degenerate, we run PnP-ULA on the posterior $\tilde{x}|y$ where $\tilde{x} := \mathbf{P}x \in \mathbb{R}^n$ denotes the vector of $n = d - m$ unobserved pixels of x , and map samples to the pixel space by using the affine mapping $f_y : \mathbb{R}^n \rightarrow \mathbb{R}^d$ defined for any $\tilde{x} \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$ by

$$f_y(\tilde{x}) = \mathbf{P}^\top \tilde{x} + \mathbf{A}^\top y.$$

Note that we can write the log-posterior $\tilde{U}_\varepsilon(\tilde{x}) = -\log p_\varepsilon(\tilde{x}|y)$ on the set \mathbb{R}^n of hidden pixels in terms of f_y and the log-prior $U_\varepsilon(x) = -\log p_\varepsilon(x)$ on the set \mathbb{R}^d :

$$\tilde{U}_\varepsilon = U_\varepsilon \circ f_y.$$

Using the chain rule and Tweedie's formula, we have that for any $x \in \mathbb{R}^d$ and $y \in \mathbb{R}^m$

$$b_\varepsilon(\tilde{x}) = -\nabla \tilde{U}_\varepsilon(\tilde{x}) = -\mathbf{P} \nabla U_\varepsilon(f_y(\tilde{x})) = (1/\varepsilon) \mathbf{P} (D_\varepsilon - \text{Id})(f_y(\tilde{x})).$$

Since \mathbf{P} and f_y are 1-Lipschitz, $b_\varepsilon = -\nabla \tilde{U}_\varepsilon$ is also Lipschitz with constant $\tilde{L} \leq (L/\varepsilon)$.

Figure 1 shows the six test images of size 256×256 pixels that were used in the experiments. We have selected these six images for their diversity in composition, content and level of detail (some images are predominantly composed of piece-wise constant regions, whereas others are rich in complex textures). This diversity will highlight strengths and limitations of the chosen denoiser as an image prior. Figure 2 depicts the corresponding blurred images and Figure 3 the images to inpaint.

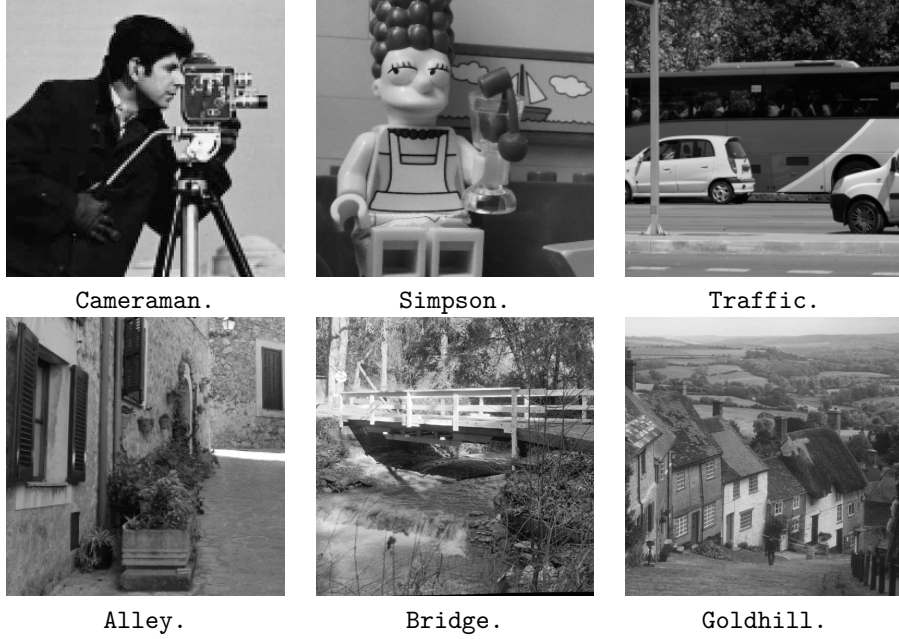


Figure 1: Original images used for the deblurring and inpainting experiments.

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 .

Choice of the denoiser The theory presented in Section 3 requires that D_ϵ satisfies **H2**(R). As default choice, we recommend using a pretrained denoising neural network such as the one described in [73]. The Lipschitz constant of the network is controlled during training by using spectral normalization and therefore the first condition of **H2**(R) holds. Moreover, the loss function used to train the network is given by ℓ_ϵ as introduced in Section 3.2. Therefore, under the conditions of Proposition 4, we get that the second condition of **H2**(R) holds.

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 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 $b_\epsilon(x) = \nabla \log p_\epsilon(x|y)$; namely, we

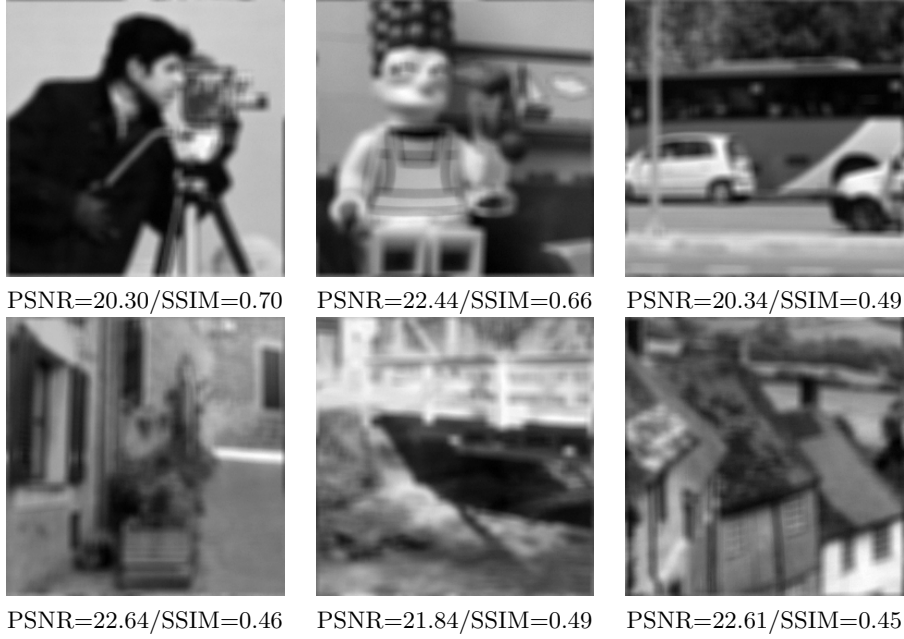


Figure 2: Images of Figure 1, blurred using a 9×9 -box-filter operator and corrupted by an additive Gaussian white noise with standard deviation $\sigma = 1/255$.

require $\delta < (1/3) \text{Lip}(b_\varepsilon)^{-1}$ where

$$\text{Lip}(b_\varepsilon) = \begin{cases} \alpha L/\varepsilon + 1/\lambda & \text{for the inpainting problem} \\ \alpha L/\varepsilon + L_y + 1/\lambda & \text{otherwise} \end{cases}$$

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 $L_y = \|\mathbf{A}^\top \mathbf{A}\|/\sigma^2$, so we choose δ just below the upper bound $\delta_{th} = 1/3(\text{Lip}(b_\varepsilon))^{-1}$ where \mathbf{A}^\top is the adjoint of \mathbf{A} . For PPnP-ULA, we set $\delta < (L/\varepsilon + L_y)^{-1}$ (resp. $\delta < (L/\varepsilon)^{-1}$ for inpainting) 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 sufficiently 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 $\lambda \in (0, 1/2(L/\varepsilon + 2L_y))$. 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 $(2L/\varepsilon + 4L_y)^{-1}$.

Other parameters The compact set \mathbf{C} is defined as $\mathbf{C} = [-1, 2]^d$, even if in practice no samples were generated 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 noise level of the denoiser D_ε to $\varepsilon = (5/255)^2$. The initialization X_0 can be set to a random vector. In our experiments (where $m = d$), we chose $X_0 = y$ in order to reduce the number of burn-in

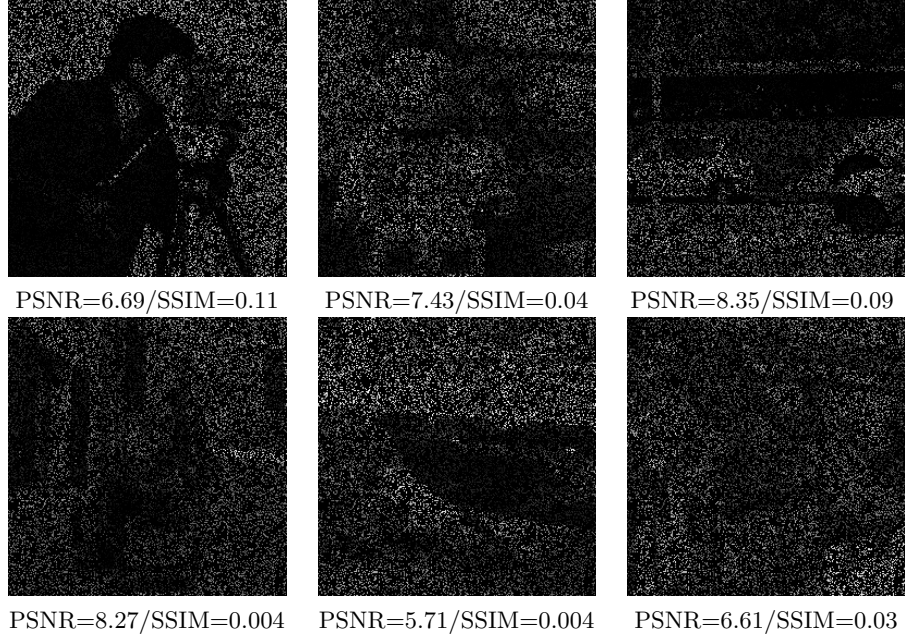


Figure 3: Images of Figure 1, with 80% missing pixels.

iterations. For $m \neq d$ we could use $X_0 = \mathbf{A}^\top y$ instead. Concerning the regularization parameter α , by default we set $\alpha = 1$, but in some cases it is possible to marginally improve the results by fine tuning it. All algorithms are implemented using Python and the PyTorch library, and run on an Intel Xeon CPU E5-2609 server with a Nvidia Titan XP graphic card or on Idris’ Jean-Zay servers featuring Intel Cascade Lake 6248 CPUs with a single Nvidia Tesla V100 SXM2 GPU. Reported running times correspond to the Xeon + Titan XP configuration.

4.2 Convergence analysis of PnP-ULA in non-blind image deblurring and inpainting

When using a sampling algorithm such as PnP-ULA on a new problem, it is essential to check that the state space is correctly explored. In order to provide a thorough convergence study, we first run the algorithm for 25×10^6 iterations. We use a burn-in period of 2.5×10^6 iterations, and consider only the samples computed after this burn-in period to study the Markov chain in close-to-stationary regime. In section 4.3, we will see that much less iterations are required if the goal is only to compute point estimators with PnP-ULA. For simplicity, the algorithm is always initialized with the observation y in our experiments with PnP-ULA (for inpainting, this means that unknown pixels are initialized to the value 0).

There is no fully comprehensive way to empirically characterise the convergence properties of a high-dimensional Markov chain, as different statistics computed from the same chain align differently with the eigenfunctions of the Markov kernel and hence exhibit different convergence speeds. In problems of small dimension, we would calculate and analyse the d -dimensional multivariate autocorrelation function (ACF) of the Markov chain, but this is not feasible in imaging problems. In problems of moderate dimension, one could characterise the range of convergence speeds by first estimating the posterior covariance matrix (which, for 256×256

images, would be a $256^2 \times 256^2$ matrix) and then performing a principal component analysis on this matrix to identify the directions with smallest and largest uncertainty, as these would provide a good indication of the subspaces where the chain converges the fastest and the slowest. However, computing the posterior covariance matrix is also not possible in imaging problems because of the dimensionality involved. Here we focus on approximations of the posterior covariance which make sense for the particular inverse problem we study. More precisely, we use the diagonalization basis of the inverse operator, *i.e.* the Fourier basis for the deblurring experiments, and the basis formed by the unknown pixels for the inpainting experiments. Under the assumption that the posterior covariance is mostly determined by the likelihood, this strategy allows broadly identifying the linear statistics that converge fastest and slowest, without requiring the estimation and manipulation of prohibitively large matrices.

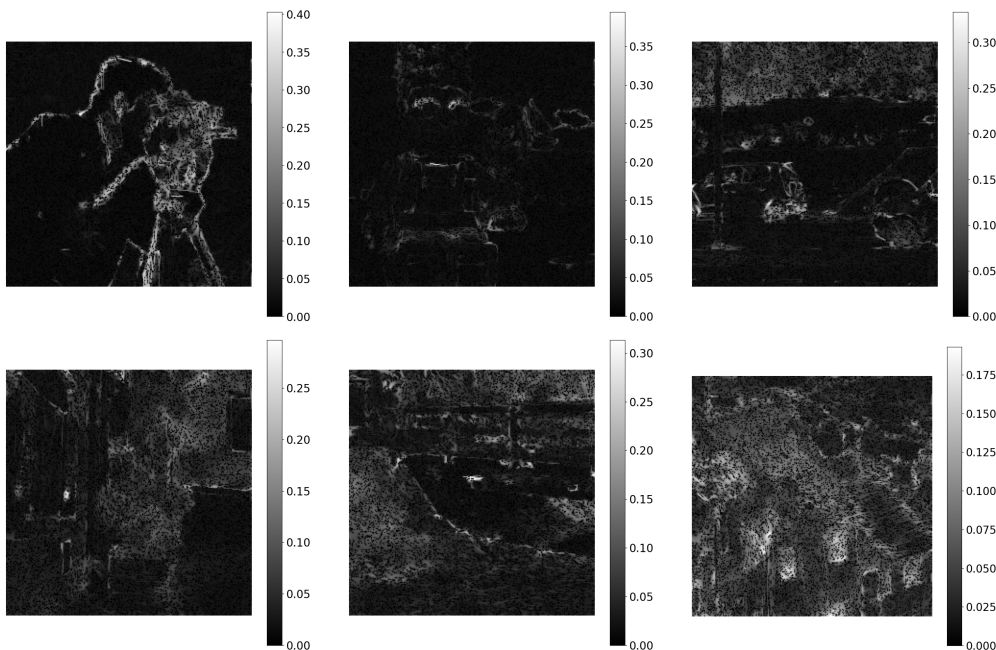


Figure 4: Marginal posterior standard deviation of the unobserved pixels for the inpainting problem. Uncertainty is located around edges and in textured areas.

Inpainting We first focus on the inpainting problem. Figure 4 shows a map of the pixel-wise marginal standard deviations, for all images. We observe that pixels in homogeneous regions have low uncertainty, while pixels on textured regions, edges, or complex structures (a reflection on the window shutter in the Alley image for instance) are the most uncertain.

For the same experiments, Figure 5 shows the Euclidean distance between the final MMSE estimate (computed using all samples) and the samples of the chain, every 2500 samples (after the burn-in period, and hence in what is considered to be a close-to-stationary regime). Fluctuations around the posterior mean and the absence of temporal structure in the plots of **Alley** or **Goldhill** are a first indication that the chain explores the solution space with ease. However, in some other cases such as the **Simpson** image, we observe meta-stability, where the chain stays in a region of the space for millions of iterations and then jumps to a different region, again for

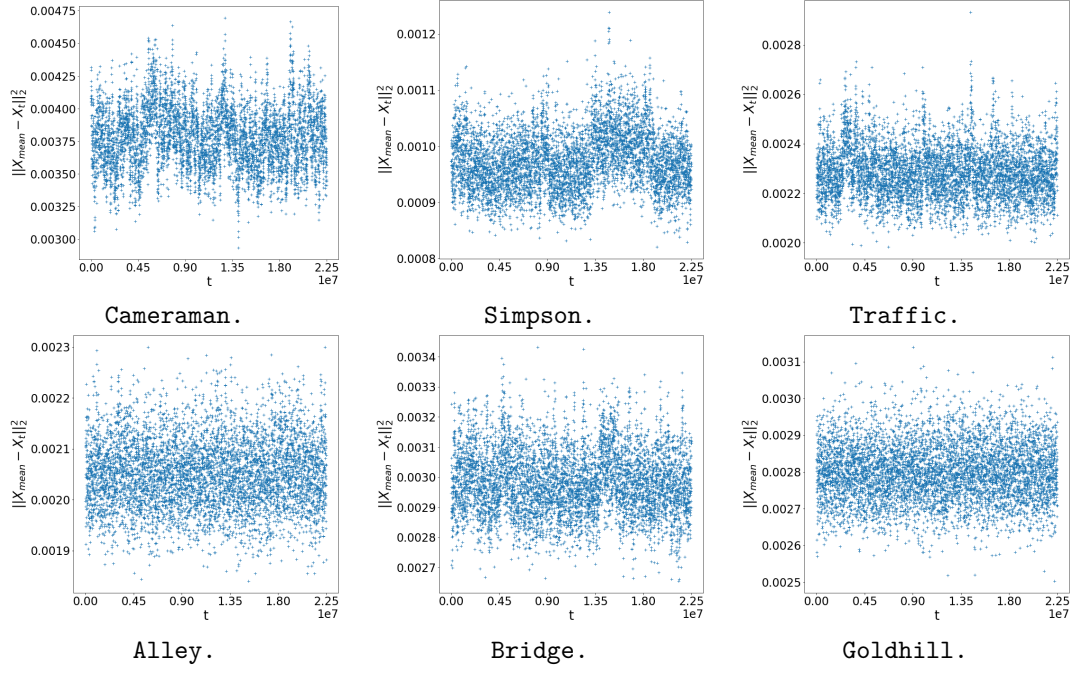


Figure 5: Evolution of the L_2 distance between the final MMSE estimate and the samples generated by PnP-ULA for the inpainting problem after the burn-in phase. Samples randomly oscillate around the MMSE. It means that they are uncorrelated. For the images **Cameraman**, **Simpson** or **Bridge**, we note a change of range for the L_2 distance. It could be interpreted as a mode switching as our posterior is likely not log-concave.

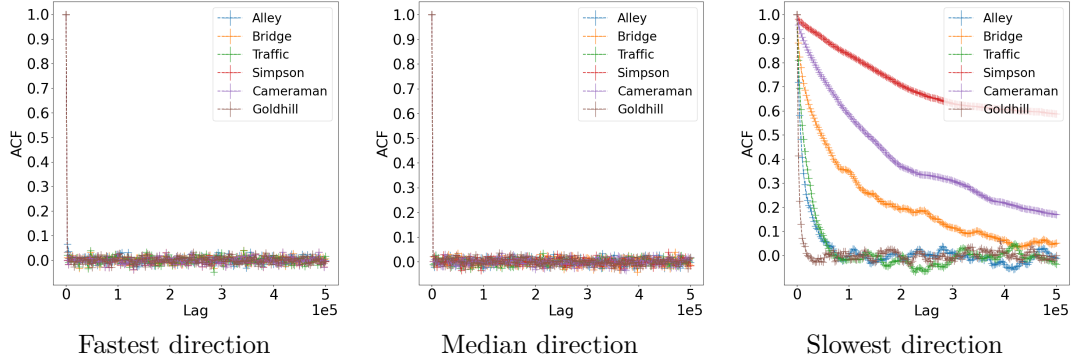


Figure 6: ACF for the inpainting problem. The ACF are shown for lags up to $5e5$ for all images in the pixel domain. After $5e5$ iterations, sample pixels are nearly uncorrelated in all spatial directions for the images **Traffic**, **Alley**, **Bridge** and **Goldhill**. For the images **Cameraman** and **Simpson**, in the slowest direction, samples need more iterations to become uncorrelated.

millions of iterations. This is one of the drawbacks of operating with a posterior distribution that is not log-concave and that may exhibit several modes.

Lastly, Figure 6 displays the sample ACFs of the fastest and slowest converging statistics associated with the inpainting experiments (as estimated by identifying, for each image, the unknown pixels with lowest and highest uncertainty). These ACF plots measure how fast samples become uncorrelated. A fast decay of the ACF is associated with good Markov chain mixing, which in turn implies accurate Monte Carlo estimates. On the contrary, a slow decay of the ACF indicates that the Markov chain is moving slowly, which leads to Monte Carlo estimates with high variance. As mentioned previously, because computing and visualising a multivariate ACF is difficult, here we show the ACF of the chain along the slowest and the fastest directions in the spatial domain (for completeness, we also show the ACF for a pixel with median uncertainty). We see that independence is reached very fast in the subspaces of low or median uncertainty, and is much slower for the few very uncertain pixels.

Deblurring We now focus on the non-blind image deblurring experiments, where, as explained previously, we perform our convergence analysis by using statistics associated with the Fourier domain. Figure 7 depicts the marginal standard deviation of the Fourier coefficients (in absolute value), for all images. For the three images **Cameraman**, **Simpsons** and **Traffic**, all the standard deviations have a similar range of values, and the largest values are observed around frequencies in the kernel of the blur filter (shown on the right of the same figure) and for high frequencies. Conversely, for the three images **Alley**, **Bridge** and **Goldhill**, very high uncertainty is observed in the vicinity of four specific frequencies. This suggests that the denoiser used is struggling to regularise these specific frequencies, and consequently the posterior distribution is very spread along these directions and difficult to explore by Markov chain sampling as a result. Interestingly, this phenomenon is only observed in the images that are rich in texture content.

Moreover, Figure 8 depicts the Euclidean distance between the MMSE estimator computed from entire chain (*i.e.* all samples) and each sample (we show one point every 2500 samples). We notice that many of the images exhibit some degree of meta-stability or slow convergence because of the presence of directions in the solution space with very high uncertainty. Again, this is consistent with our convergence theory, which identifies posterior multimodality and anisotropy as key challenges that future work should seek to overcome.

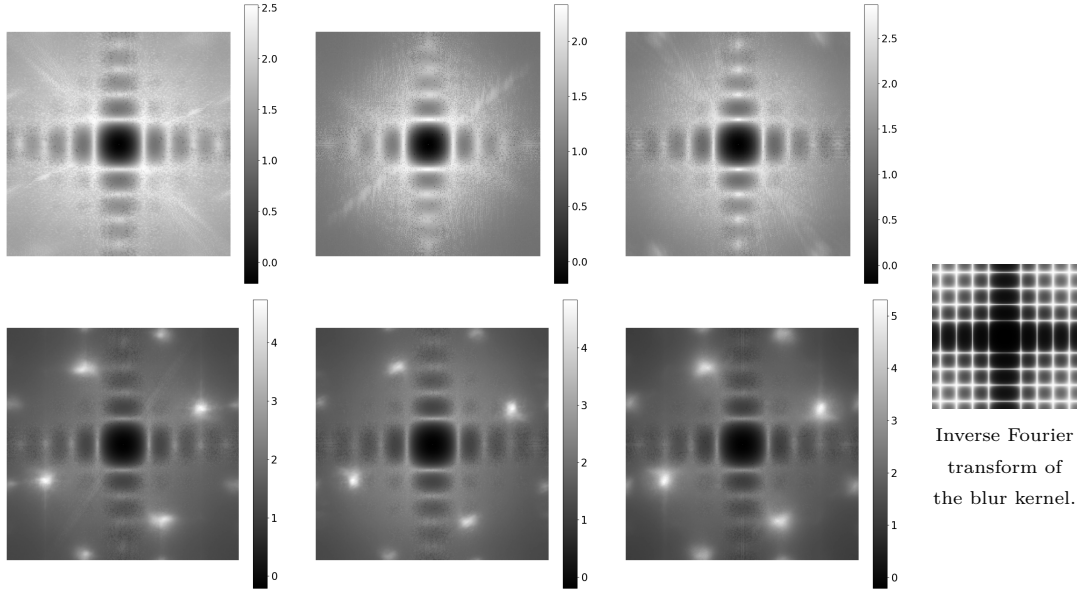


Figure 7: Log-standard deviation maps in the Fourier domain for the Markov chains defined by PnP-ULA for the deblurring problem. First line: images **Cameraman**, **Simpson**, **Traffic**. Second line: images **Alley**, **Bridge** and **Goldhill**. For the first three images, we clearly see that uncertainty is observed on frequencies that are near the kernel of the blur filter (shown on the right), and is also higher around high frequencies (*i.e.* around edges and textured areas in images). For the last three images, very high uncertainty is observed around some specific frequencies. In the direction of these frequencies, the Markov chain is moving very slowly and the mixing time of the chain is particularly slow, as shown on Figure 9.

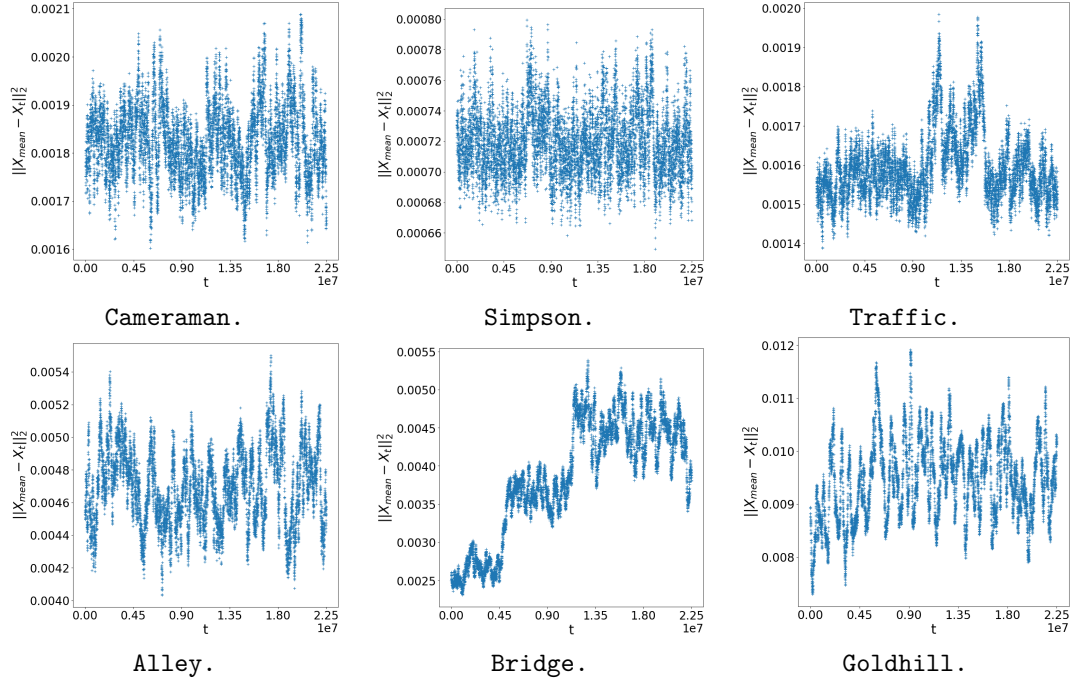


Figure 8: Evolution of the L_2 distance between the final MMSE estimate and the samples generated by PnP-ULA for the deblurring problem after the burn-in phase. For images as **Cameraman** or **Simpson**, samples randomly oscillate around the MMSE. On the contrary, for images as **Bridge** or **Goldhill**, the plot is structured, meaning that samples are still correlated.

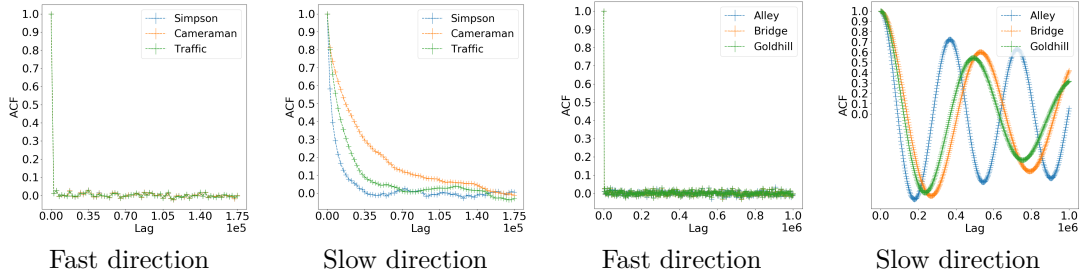


Figure 9: ACF for the deblurring problem. The ACF are shown for lags up to $1.75e5$ for the three images **Cameraman**, **Simpson** and **Traffic** (see the two plots to the left) and independence seems to be achieved in all directions. For the three other images, independence is not achieved in the slowest direction (corresponding to the most uncertain frequency of the samples in the Fourier domain) even after $1e6$ iterations.

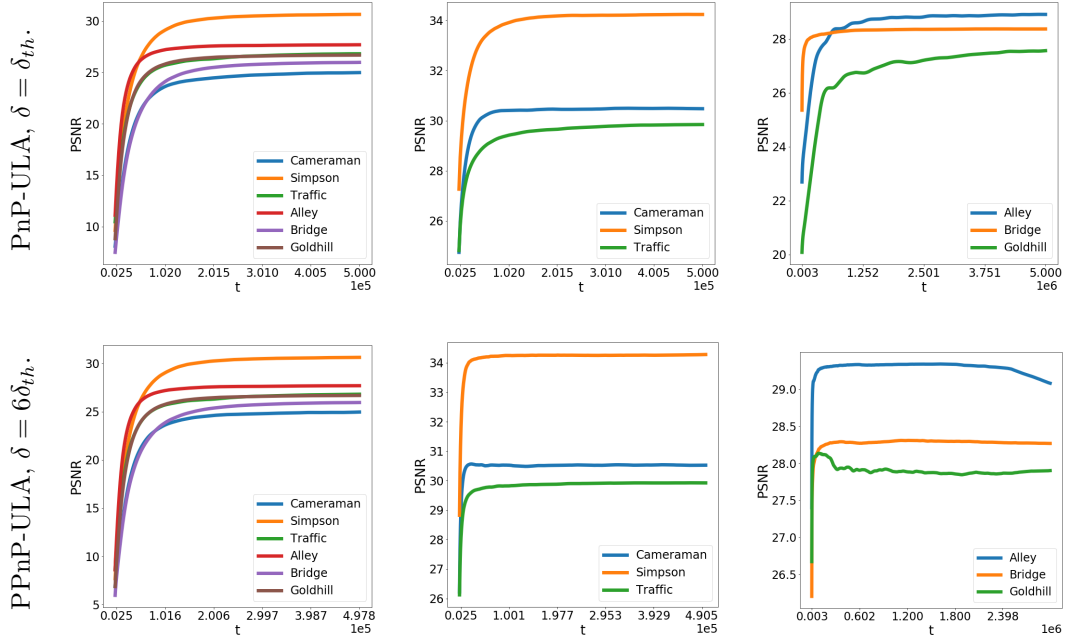


Figure 10: Left: PSNR evolution of the estimated MMSE for the inpainting problem. After $5e5$ iterations, the convergence of the first order moment of the posterior distribution seems to be achieved for all images. Middle and right: PSNR evolution of the estimated MMSE for the deblurring problem. The convergence for the posterior mean can be fast for simple images such as **Cameraman**, **Simpson**, and **Traffic** (for these images the PSNR evolution is shown for the first $5e5$ iterations). Increasing the δ increases the convergence speed for these images by a factor close to 2. For more complex images, such as **Alley** or **Goldhill**, the convergence is much slower and is still not achieved after $3e6$ iterations with PPnP-ULA for $\delta = 6\delta_{th}$.

Lastly, we show on Figure 9 the sample ACFs for the slowest and the fastest directions in the Fourier domain⁷. Again, in all experiments, independence is achieved quickly in the fastest direction. The behaviours of the slowest direction for the three images **Alley**, **Bridge** and **Goldhill** suggest that the Markov chain is close to the stability limit and exhibits highly oscillatory behaviour as well as poor mixing.

4.3 Point estimation for non-blind image deblurring and inpainting

We are now ready to study the quality of the MMSE estimators delivered by PnP-ULA and PPnP-ULA and report comparisons with MAP estimation by PnP-SGD [55].

Quantitative results Figure 10 illustrates the evolution of the PSNR of the mean of the Markov chain (the Monte Carlo estimate of the MMSE solution), as a function of the number of iterations, for the six images of Figure 1. These plots have been computed by using a step-size $\delta = \delta_{th}$ that is just below the stability limit and a 1-in-2500 thinning. We observe that the PSNR between the MMSE solution as computed by the Markov chain and the truth stabilises

⁷The slowest direction corresponds to the Fourier coefficient with the highest (real or imaginary) variance.

in approximately 10^5 iterations in the experiments where the chain exhibits fast convergence, whereas over 10^6 are required in experiments that suffer from slow convergence (e.g., deblurring of *Alley*, *Bridge* and *Goldhill*). Moreover, we observe that using PnP-ULA with a larger step-size can noticeably reduce the number of iterations required to obtain a stable estimate of the posterior mean, particularly in the image deblurring experiments.

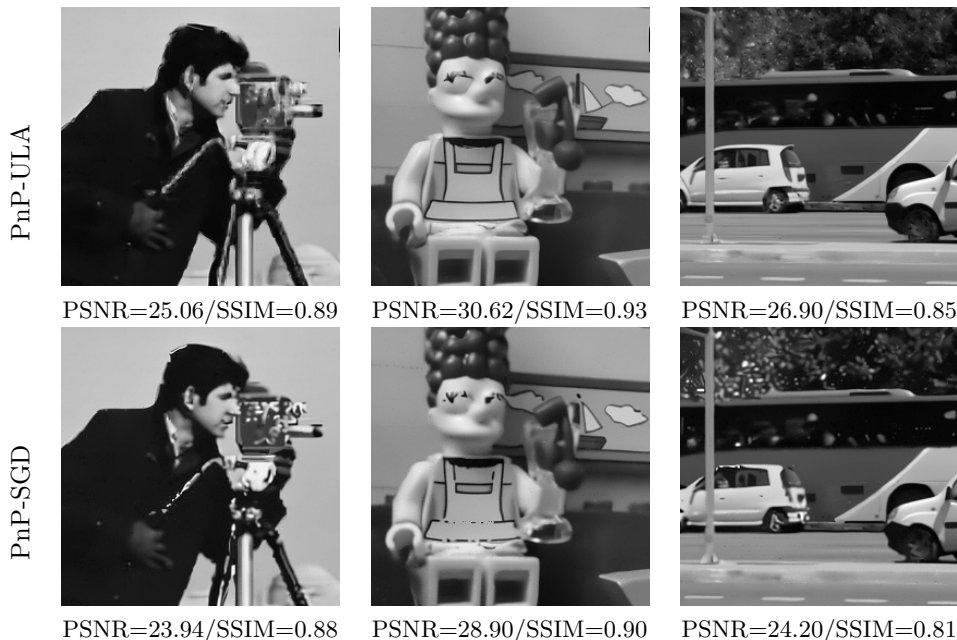


Figure 11: Results comparison for the inpainting task of the images presented in Figure 3 using PnP-ULA (first row) and PnP-SGD initialized with a TVL2 restoration (second row).

Visual results Figures 11, 12, 13 and 14 show the MMSE estimate computed by PnP-ULA on the whole chain including the burn-in for the 6 images, for the inpainting and deblurring experiments. We also provide the MAP estimation results computed by using PnP-SGD [55], which targets the same posterior distributions. We report the *Peak Signal-To Noise Ratio* (PSNR) and the *Structural Similarity Index* (SSIM) [83, 84] for all these experiments.

For the inpainting experiments, PnP-SGD struggles to converge when initialized with the observed image (see [55]). For this reason, we warm start PnP-SGD by using an estimate of x obtained by minimizing the Total Variation pseudo-norm under the constraint of the known pixels. For simplicity, PnP-ULA is initialized with the observation y . We observe in Figure 11 and Figure 12 that the results obtained by computing the MMSE Bayesian estimator with PnP-ULA are visually and quantitatively superior to the ones delivered by MAP estimation with PnP-SGD. In particular, the sampling approach seems to better recover the continuity of fine structures and lines in the different images.

For the deblurring experiments, the results of PnP-SGD are provided by using a regularisation parameter $\alpha = 0.3$ (which was shown to yield optimal results on this set of images in [55]) and for $\alpha = 1$, which recovers the model used by PnP-ULA. Observe that for the three first images (shown on Figure 13), the MMSE result is much sharper than the best MAP result, and the PSNR / SSIM results also show a clear advantage for the MMSE. For the other three images (results are

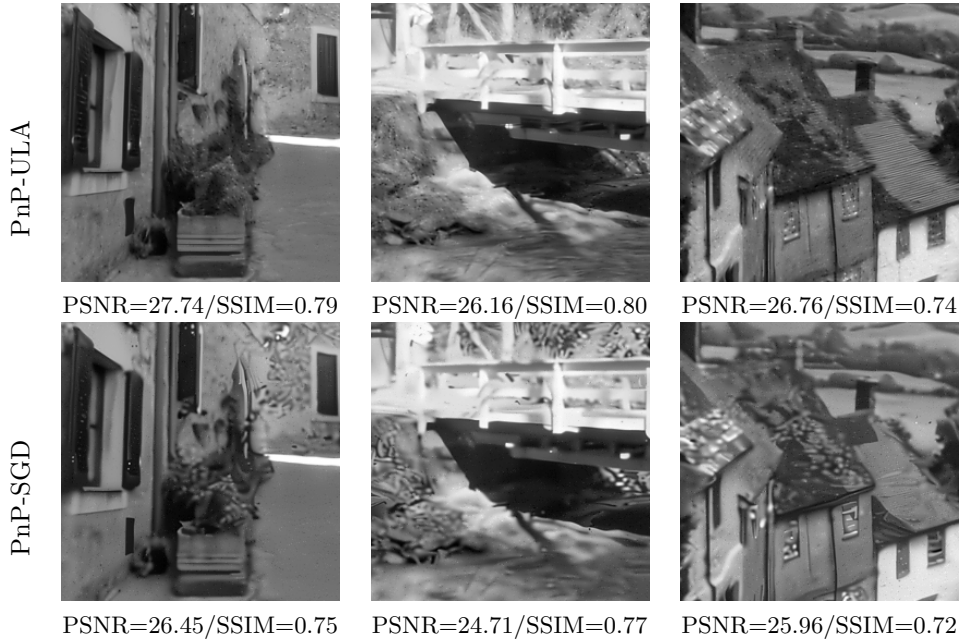


Figure 12: Results comparison for the inpainting task of the images presented in Figure 3 using PnP-ULA (first row) and PnP-SGD initialized with a TVL2 restoration (second row).

shown on Figure 14), the quality of the MMSE solutions delivered is slightly deteriorated by the slow convergence of the Markov chain and the poor regularisation of some specific frequencies, which leads to a common visual artefact (a rotated rectangular pattern). Using a different denoiser more suitable for handling textures, or combining a learnt denoiser with an analytic regularisation term, might correct this behaviour and will be the topic of future work.

A partial conclusion from this set of comparisons is that the sampling approach of PnP-ULA, when it samples the space correctly, seems to provide much better results than the MAP estimator for the same posterior. Of course, this increase in quality comes at the cost of a much higher computation time.

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 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 [20] for more details). The analysis proceeds as follows. First, Figure 4 and Figure 15 show the marginal posterior standard deviation associated with each image pixel, as computed by PnP-ULA over all samples, for the inpainting and deblurring problems. As could be expected, we observe for both problems that highly uncertain pixels are concentrated around the edges of the reconstructed images, but also on textured areas. The dynamic range of the pixel standard deviations is larger for the inpainting problem than for deblurring, which suggests that the problem has a higher level of intrinsic uncertainty.

Figure 16 shows the evolution of the RMSE between the standard deviation computed along the samples and its asymptotic value, respectively for the inpainting and deblurring problems.



Figure 13: Results comparison for the deblurring task of the images presented in Figure 2 using PnP-ULA with $\alpha = 1$ (first row), PnP-SGD with $\alpha = 0.3$ (second row) and $\alpha = 1$ (third row). PnP-ULA was initialized with the observation y (see Figure 2) whereas PnP-SGD was initialised with a TVL2 restoration.

Estimating these standard deviation maps necessitates to run the chain longer than to estimate the MMSE, as could be expected for second order statistical moment.

Following on from this, to explore the uncertainty for structures that are larger than one pixel, Figure 17 and Figure 18 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 quantifying the uncertainty of larger or lower-frequency structures, such as the bottom of the glass in **Simpson** for the deblurring experiment. 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 degradations involves.

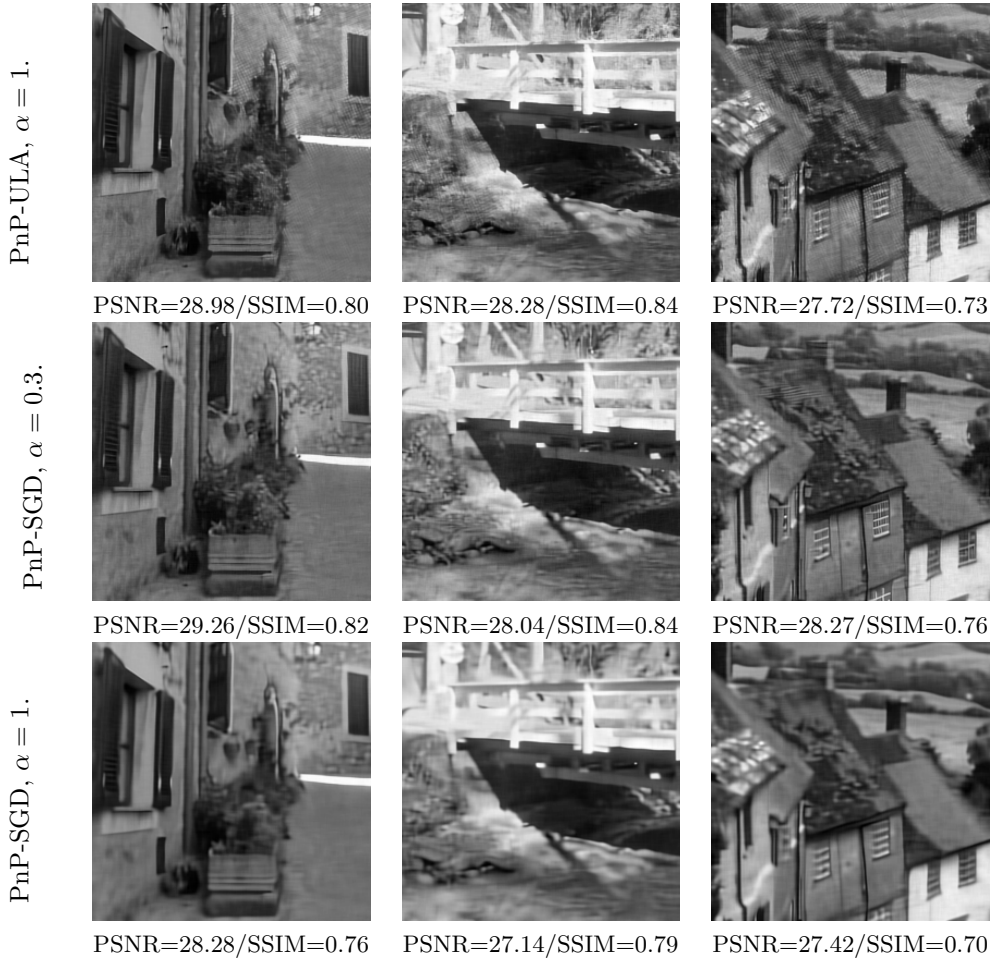


Figure 14: Results comparison for the deblurring task of the images presented in Figure 2 using PnP-ULA with $\alpha = 1$ (first row), PnP-SGD with $\alpha = 0.3$ (second row) and $\alpha = 1$ (third row). PnP-ULA was initialized with the observation y (see Figure 2) whereas PnP-SGD was initialised with a TVL2 restoration.

5 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

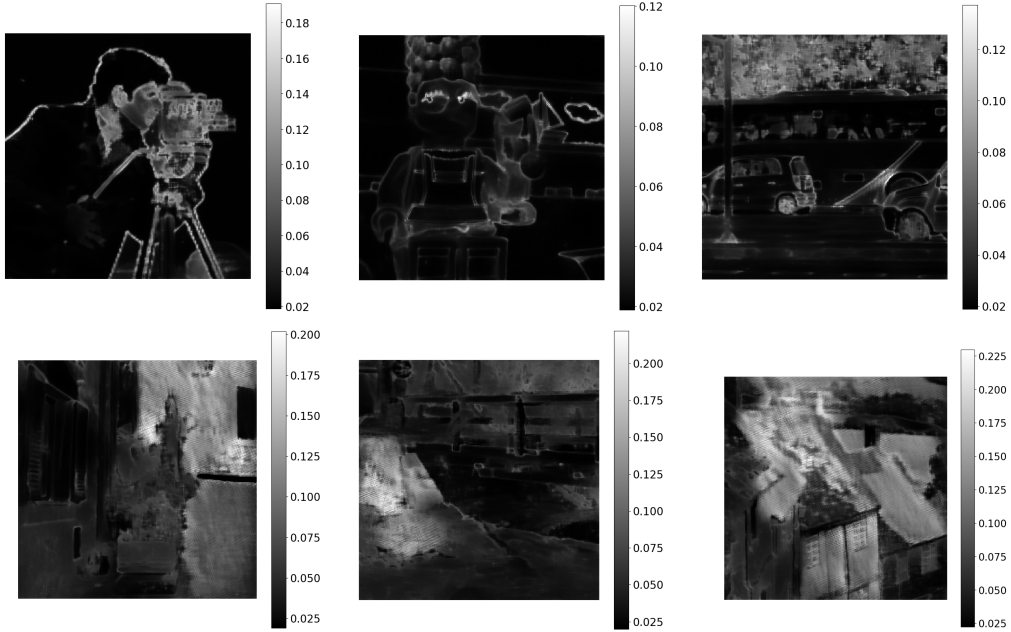


Figure 15: Marginal posterior standard deviation for the deblurring problem. On simple images such as Simpson (see figure 1), most of the uncertainty is located around the edges. For the images Alley, Bridge and Goldhill, associated with a highly correlated Markov chain in some directions, some areas are very uncertain. They correspond to the zones where the rotated rectangular pattern appears in the MMSE estimate.

with this level of insight and guarantees on the delivered solutions. We illustrated the proposed framework with two Bayesian image restoration experiments - deblurring and inpainting - where we computed point estimates as well as uncertainty visualisation and quantification analyses and highlighted how the limitations of the chosen denoiser manifest in the resulting Bayesian model and estimates.

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 priors that combine a denoiser with an analytic regularisation term, and other neural network based priors such as the generative ones used in [15] or the autoencoder-based priors in [40], as well as to generalise the Gaussian smoothing to other smoothings and investigate their properties in the context of Bayesian inverse problems. We are also very interested in strategies for training denoisers that automatically verify the conditions required for exponentially fast convergence of the Langevin SDE, for example by using the framework recently proposed in [68] to learn maximally monotone operators, or the data-driven regularisers described in [52, 63]. In addition, we would like to understand when the projected RED estimator [26] - or its relaxed variant - are the MAP estimators for well-defined Bayesian models, as well as to study the interplay between the geometric aspects of the loss defining this estimator [65] and the geometry of the set of fixed points of the denoiser defining the model. 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 [85]. From a Bayesian computation viewpoint, a priority is

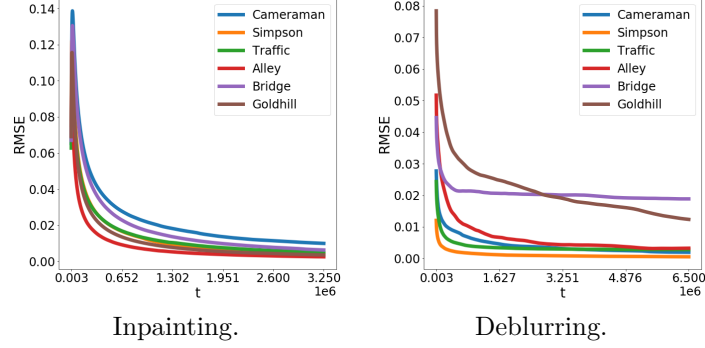


Figure 16: Evolution of the Root Mean Squared Error (RMSE) between the final standard deviation and the estimated current standard deviation for the inpainting and deblurring problems.

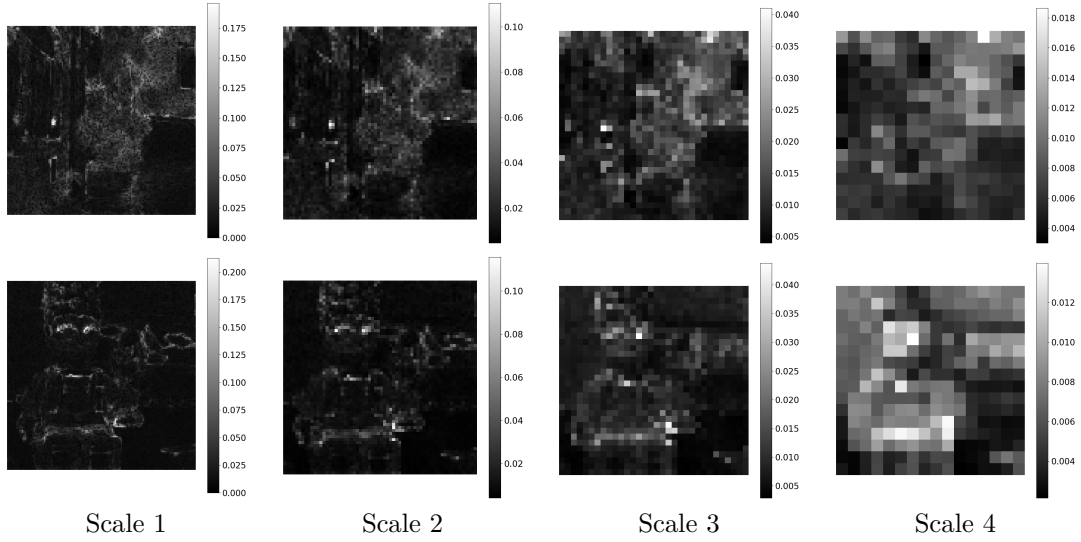


Figure 17: Marginal posterior standard deviation of the **Alley** and **Simpson** images for the inpainting problem at different scales. The scale i corresponds to a downsampling by a factor $2i$ of the original sample size.

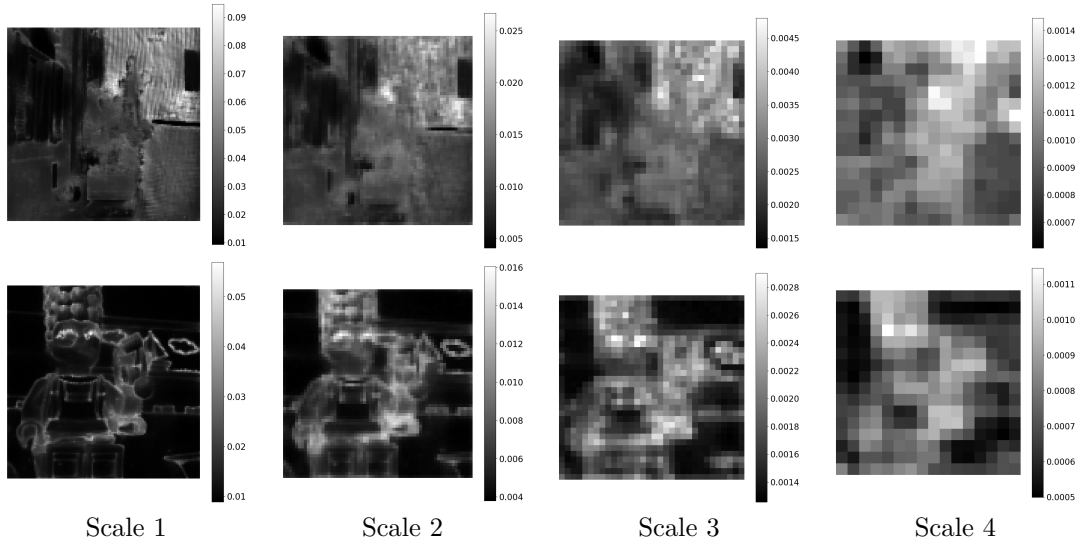


Figure 18: Marginal posterior standard deviation of the images **Alley** and **Simpson** for the deblurring problem at different scales. The scale i corresponds to a downsampling by a factor 2^i of the original sample size.

to develop accelerated 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] 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>.
- [2] R. AHMAD, C. A. BOUMAN, G. T. BUZZARD, S. CHAN, S. LIU, E. T. REEHORST, AND P. SCHNITER, *Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery*, IEEE Signal Processing Magazine, 37 (2020), pp. 105–116, <https://doi.org/10.1109/MSP.2019.2949470>.
- [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] 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).
- [5] V. ANTUN, F. RENNA, C. POON, B. ADCOCK, 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, <https://doi.org/10.1073/pnas.1907377117>.
- [6] S. ARRIDGE, P. MAASS, O. ÖKTEM, AND C.-B. SCHÖNLIEB, *Solving inverse problems using data-driven models*, Acta Numerica, 28 (2019), p. 1–174, <https://doi.org/10.1017/S0962492919000059>.
 - [7] F. BACH, *Breaking the Curse of Dimensionality with Convex Neural Networks*, The Journal of Machine Learning Research, 18 (2017), pp. 629–681, <https://jmlr.org/papers/volume18/14-546/14-546.pdf>.
 - [8] H. H. BAUSCHKE, P. L. COMBETTES, ET AL., *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, vol. 408, Springer, 2011, <https://doi.org/10.1007/978-3-319-48311-5>.
 - [9] M. BAYATI AND A. MONTANARI, *The dynamics of message passing on dense graphs, with applications to compressed sensing*, IEEE Transactions on Information Theory, 57 (2011), pp. 764–785, <https://doi.org/10.1109/TIT.2010.2094817>.
 - [10] J. BERNARDO AND A. SMITH, *Bayesian Theory*, vol. 15, 01 2000, <https://doi.org/10.2307/2983298>.
 - [11] S. A. BIGDELI, M. JIN, P. FAVARO, AND M. ZWICKER, *Deep Mean-Shift Priors for Image Restoration*, in Advances in Neural Information Processing Systems, vol. 30, Curran Associates, Inc., 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] A. BLAKE, P. KOHLI, AND C. ROTHER, *Markov Random Fields for Vision and Image Processing*, EBSCO ebook academic collection, MIT Press, 2011, <https://doi.org/10.7551/mitpress/8579.003.0001>.
 - [14] V. I. BOGACHEV, *Measure Theory*, vol. Volume 1, Springer, 1 ed., 2007, <http://gen.lib.rus.ec/book/index.php?md5=ffbdc7e3d8e571c6cd5f9e8633cdfdc2>.
 - [15] A. BORA, A. JALAL, E. PRICE, AND A. G. DIMAKIS, *Compressed Sensing using Generative Models*, in Proceedings of the 34th International Conference on Machine Learning, vol. 70 of Proceedings of Machine Learning Research, PMLR, 06–11 Aug 2017, pp. 537–546, <https://proceedings.mlr.press/v70/bora17a.html>, <https://arxiv.org/abs/arXiv:1703.03208v1>.
 - [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] S. BOUCHERON, G. LUGOSI, AND P. MASSART, *Concentration inequalities: A nonasymptotic theory of independence*, Oxford university press, 2013.

- [19] G. T. BUZZARD, S. H. CHAN, S. SREEHARI, AND C. A. BOUMAN, *Plug-and-Play unplugged: Optimization-free Reconstruction Using Consensus Equilibrium*, SIAM Journal on Imaging Sciences, 11 (2018), pp. 2001–2020, <https://doi.org/10.1137/17M1122451>.
- [20] 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>.
- [21] 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>.
- [22] 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>.
- [23] S. CHEN, C. LUO, B. DENG, Y. QIN, H. WANG, AND Z. ZHUANG, *BM3D vector approximate message passing for radar coded-aperture imaging*, in 2017 Progress in Electromagnetics Research Symposium-Fall (PIERS-FALL), IEEE, 2017, pp. 2035–2038, <https://doi.org/10.1109/PIERS-FALL.2017.8293472>.
- [24] T. CHEN, E. FOX, AND C. GUESTRIN, *Stochastic gradient Hamiltonian Monte Carlo*, in Proceedings of the 31st International Conference on Machine Learning, vol. 32 of Proceedings of Machine Learning Research, PMLR, 22–24 Jun 2014, pp. 1683–1691, <https://proceedings.mlr.press/v32/cheni14.html>.
- [25] 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>.
- [26] R. COHEN, M. ELAD, AND P. MILANFAR, *Regularization by Denoising via Fixed-Point Projection (RED-PRO)*, SIAM Journal on Imaging Sciences, 14 (2021), pp. 1374–1406, <https://doi.org/10.1137/20M1337168>, <https://arxiv.org/abs/2008.00226>.
- [27] 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>.
- [28] S. DIAMOND, V. SITZMANN, F. HEIDE, AND G. WETZSTEIN, *Unrolled optimization with deep priors*, (2017), <https://arxiv.org/abs/1705.08041>.
- [29] C. DONG, C. C. LOY, K. HE, AND X. TANG, *Learning a Deep Convolutional Network for Image Super-Resolution*, in Computer Vision – ECCV 2014, Springer International Publishing, 2014, pp. 184–199.
- [30] D. L. DONOHO, A. MALEKI, AND A. MONTANARI, *Message-passing algorithms for compressed sensing*, Proceedings of the National Academy of Sciences, 106 (2009), pp. 18914–18919, <https://doi.org/10.1073/pnas.0909892106>.
- [31] R. DOUC, E. MOULINES, P. PRIOURET, AND P. SOULIER, *Markov Chains*, Springer, 2019.

- [32] 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>.
- [33] 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, <https://doi.org/10.1137/16M1108340>.
- [34] B. EFRON, *Tweedie’s formula and selection bias*, Journal of the American Statistical Association, 106 (2011), pp. 1602–1614, <https://doi.org/10.1198/jasa.2011.tm11181>.
- [35] A. K. FLETCHER, P. PANDIT, S. RANGAN, S. SARKAR, AND P. SCHNITER, *Plug in estimation in high dimensional linear inverse problems a rigorous analysis*, Journal of Statistical Mechanics: Theory and Experiment, 2019 (2019), p. 124021, <https://doi.org/10.1088/1742-5468/ab321a>.
- [36] H. GAO, X. TAO, X. SHEN, AND J. JIA, *Dynamic Scene Deblurring With Parameter Selective Sharing and Nested Skip Connections*, in 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 3843–3851, <https://doi.org/10.1109/CVPR.2019.00397>.
- [37] M. GHARBI, G. CHAURASIA, S. PARIS, AND F. DURAND, *Deep Joint Demosaicking and Denoising*, ACM Transactions on Graphics (TOG), 35 (2016), p. 191, <https://doi.org/10.1145/2980179.2982399>.
- [38] D. GILTON, G. ONGIE, AND R. WILLETT, *Neumann networks for inverse problems in imaging*, (2019), <https://arxiv.org/abs/1901.03707>.
- [39] 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, <https://doi.org/10.1111/j.1467-9868.2010.00765.x>.
- [40] M. GONZÁLEZ, A. ALMANSA, AND P. TAN, *Solving Inverse Problems by Joint Posterior Maximization with Autoencoding Prior*, arXiv, (2021), <https://arxiv.org/abs/2103.01648>.
- [41] 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, <https://icml.cc/Conferences/2010/papers/449.pdf>.
- [42] B. GUO, Y. HAN, AND J. WEN, *AGEM: Solving Linear Inverse Problems via Deep Priors and Sampling*, in Advances in Neural Information Processing Systems, Curran Associates, Inc., 2019, pp. 547–558, <https://proceedings.neurips.cc/paper/2019/file/49182f81e6a13cf5eaa496d51fea6406-Paper.pdf>.
- [43] J. HO, A. JAIN, AND P. ABBEEL, *Denoising Diffusion Probabilistic Models*, Advances in Neural Information Processing Systems, 33 (2020), pp. 6840–6851, <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>.
- [44] A. HOUDARD, C. BOUYEYRON, AND J. DELON, *High-Dimensional Mixture Models For Unsupervised Image Denoising (HDMI)*, SIAM Journal on Imaging Sciences, 11 (2018), pp. 2815–2846, <https://doi.org/10.1137/17M1135694>.

- [45] S. HURAUULT, A. LECLAIRE, AND N. PAPADAKIS, *Gradient step denoiser for convergent plug-and-play*, arXiv preprint arXiv:2110.03220, (2021).
- [46] A. JAVANMARD AND A. MONTANARI, *State evolution for general approximate message passing algorithms, with applications to spatial coupling*, Information and Inference: A Journal of the IMA, 2 (2013), pp. 115–144, <https://doi.org/10.1093/imaiai/iat004>.
- [47] Z. KADKHODAIE AND E. P. SIMONCELLI, *Stochastic Solutions for Linear Inverse Problems using the Prior Implicit in a Denoiser*, Advances in Neural Information Processing Systems, (2021), <https://proceedings.neurips.cc/paper/2021/file/6e28943943dbed3c7f82fc05f269947a-Paper.pdf>.
- [48] 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, <https://doi.org/10.1109/LSP.2017.2763583>.
- [49] 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>.
- [50] B. KAWAR, G. VAKSMAN, AND M. ELAD, *Snips: Solving noisy inverse problems stochastically*, arXiv preprint arXiv:2105.14951, (2021).
- [51] B. KAWAR, G. VAKSMAN, AND M. ELAD, *Stochastic Image Denoising by Sampling from the Posterior Distribution*, in Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops, October 2021, pp. 1866–1875, https://openaccess.thecvf.com/content/ICCV2021W/AIM/html/Kawar_Stochastic_Image_Denoising_by_Sampling_From_the_Posterior_Distribution_ICCVW_2021_paper.html.
- [52] E. KOBLER, A. EFFLAND, K. KUNISCH, AND T. POCK, *Total Deep Variation for Linear Inverse Problems*, in 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020, <https://doi.org/10.1109/CVPR42600.2020.00757>.
- [53] S. KULLBACK, *Information theory and statistics*, Courier Corporation, 1997.
- [54] 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>.
- [55] R. LAUMONT, V. DE BORTOLI, A. ALMANSA, J. DELON, A. DURMUS, AND M. PEREYRA, *On maximum-a-posteriori estimation with plug & play priors and stochastic gradient descent*, (2021), <https://hal.archives-ouvertes.fr/hal-03348735/document>.
- [56] J. LEHTINEN, J. MUNKBERG, J. HASSELGREN, S. LAINE, T. KARRAS, M. AITTALA, AND T. AILA, *Noise2noise: Learning image restoration without clean data*, 80 (2018), pp. 2965–2974, <https://proceedings.mlr.press/v80/lehtinen18a.html>.
- [57] 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.
- [58] 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>.

- [59] 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.
- [60] C. A. METZLER, A. MALEKI, AND R. G. BARANIUK, *From denoising to compressed sensing*, IEEE Transactions on Information Theory, 62 (2016), pp. 5117–5144, <https://doi.org/10.1109/TIT.2016.2556683>.
- [61] 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>.
- [62] T. MIYATO, T. KATAOKA, M. KOYAMA, AND Y. YOSHIDA, *Spectral Normalization for Generative Adversarial Networks*, (2018), <https://openreview.net/forum?id=B1QRgziT->.
- [63] S. MUKHERJEE, S. DITTMER, Z. SHUMAYLOV, S. LUNZ, O. ÖKTEM, AND C.-B. SCHÖNLIEB, *Learned convex regularizers for inverse problems*, 2021, <https://arxiv.org/abs/2008.02839>.
- [64] 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>, <https://arxiv.org/abs/1306.0187>.
- [65] M. PEREYRA, *Revisiting Maximum-A-Posteriori Estimation in Log-Concave models*, SIAM Journal on Imaging Sciences, 12 (2019), pp. 650–670, <https://doi.org/10.1137/18M1174076>.
- [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, <https://doi.org/10.1109/JSTSP.2015.2496908>.
- [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>.
- [68] J.-C. PESQUET, A. REPETTI, M. TERRIS, AND Y. WIAUX, *Learning Maximally Monotone Operators for Image Recovery*, 2020, <https://doi.org/10.1137/20M1387961>, <https://arxiv.org/abs/2012.13247>.
- [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, <https://doi.org/10.2307/3318418>.

- [72] 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).
- [73] 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>.
- [74] 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, <https://doi.org/10.1109/TIP.2018.2872858>.
- [75] L. SCHWARTZ, *Désintégration d’une mesure*, Séminaire Équations aux dérivées partielles (Polytechnique), pp. 1–10, <http://eudml.org/doc/111551>.
- [76] Y. SONG AND S. ERMON, *Generative modeling by estimating gradients of the data distribution*, in Advances in Neural Information Processing Systems, vol. 32, 2019, <https://proceedings.neurips.cc/paper/2019/file/3001ef257407d5a371a96dcd947c7d93-Paper.pdf>.
- [77] A. M. STUART, *Inverse problems: A Bayesian perspective*, Acta Numerica, 19 (2010), p. 451–559, <https://doi.org/10.1017/S0962492910000061>.
- [78] Y. SUN, Z. WU, B. WOHLBERG, AND U. S. KAMILOV, *Scalable Plug-and-Play ADMM with Convergence Guarantees*, IEEE Transactions on Computational Imaging, 7 (2021), pp. 849–863, <https://doi.org/10.1109/TCI.2021.3094062>.
- [79] 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>.
- [80] 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, <https://doi.org/10.1109/GlobalSIP.2013.6737048>.
- [81] 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>. Old and new.
- [82] M. VONO, N. DOBIGEON, AND P. CHAINAIS, *Asymptotically exact data augmentation: models, properties and algorithms*, arXiv preprint arXiv:1902.05754, (2019).
- [83] 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>.
- [84] 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.
- [85] J. WATSON AND C. HOLMES, *Approximate Models and Robust Decisions*, Statistical Science, 31 (2016), pp. 465–489, <https://doi.org/10.1214/16-sts592>.

- [86] X. XU, Y. SUN, J. LIU, B. WOHLBERG, AND U. S. KAMILOV, *Provable Convergence of Plug-and-Play Priors with MMSE denoisers*, IEEE Signal Processing Letters, 27 (2020), pp. 1–10, <https://doi.org/10.1109/LSP.2020.3006390>, <https://arxiv.org/abs/2005.07685>.
- [87] 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, <https://doi.org/10.1109/TIP.2011.2176743>.
- [88] 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, <https://doi.org/10.1109/TIP.2017.2662206>.
- [89] 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>.
- [90] 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, <https://doi.org/10.1109/TIP.2018.2839891>.
- [91] 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/EPLLICVCameraReady.pdf>.

A Organization of the supplementary

In this supplementary document we present some extensions and gather the proofs of this paper. We first introduce a more general framework in Appendix B. Then in Appendix C we present our improved convergence results in the case where the log-likelihood is strongly log-concave. Posterior approximation bounds in our general setting are gathered in Appendix D. Then we turn to the proof of these results. We first derive technical results in Appendix E. Proofs of Section 3.2 and Section 3.3 are presented in Appendix F and Appendix G respectively. Finally, proofs of Appendix D are given in Appendix H.

B A general framework

We start by considering a slightly more general framework than the one previously introduced. More precisely, instead of p^* we consider a general distribution p and instead of considering p_ε as a prior we consider a tamed version of this density by introducing another hyperparameter $\alpha > 0$. In what follows, we describe this setting in details. We start by recalling a mild assumption on the likelihood.

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.

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

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

One typical example of likelihood function that we consider in our numerical illustration, see Section 4, is $p(y|x) \propto \exp[-\|\mathbf{A}x - y\|^2 / (2\sigma^2)]$ 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 .

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

H5. *The following hold:*

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

Note that if $\alpha = 1$, **H5-(a)** hold under **H1**, see Proposition 1. Under **H5-(a)**, define π the target probability distribution for any $x \in \mathbb{R}^d$ by

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

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

$$(d\pi_\varepsilon/d\text{Leb})(x) = p(y|x) p_\varepsilon^\alpha(x) \Big/ \int_{\mathbb{R}^d} p(y|\tilde{x}) p_\varepsilon^\alpha(\tilde{x}) d\tilde{x}.$$

We also recall the assumption on the denoiser D_ε , see Section 3.2 for details.

H2. *There exist $\varepsilon_0 > 0$, $\mathbf{M}_R \geq 0$ and $\mathbf{L} \geq 0$ such that for any $\varepsilon \in (0, \varepsilon_0]$, $x_1, x_2 \in \mathbb{R}^d$ and $x \in \mathbf{B}(0, R)$ we have*

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

where we recall that

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

C Strongly log-concave case

We now present an improvement on the results of Section 3.2 in the case where the log-likelihood $x \mapsto \log p(y|x)$ is strongly concave. We recall that the Markov chain is given by the following recursion: $X_0 \in \mathbb{R}^d$ and for any $k \in \mathbb{N}$

$$\begin{aligned} X_{k+1} &= X_k + \delta b_\varepsilon(X_k) + \sqrt{2\delta} Z_{k+1}, \\ b_\varepsilon(x) &= \nabla \log p(y|x) + \alpha P_\varepsilon(x) + (\text{prox}_\lambda(\iota_{\mathbf{C}})(x) - x)/\lambda, \quad P_\varepsilon(x) = (D_\varepsilon(x) - x)/\varepsilon, \end{aligned} \tag{17}$$

In the strongly concave setting we set $\mathbf{C} = \mathbb{R}^d$, i.e. $\forall x \in \mathbf{C}$, $\text{prox}_\lambda(\iota_{\mathbf{C}})(x) = x$. We recall that in our image processing applications, we have that for any $x \in \mathbb{R}^d$, $p(y|x) \propto \exp[-\|\mathbf{A}x - y\|^2 / (2\sigma^2)]$ and that $x \mapsto p(y|x)$ 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 (17) is geometrically ergodic under **H2** 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

$$V(x) = 1 + \|x\|^2. \quad (18)$$

The following proposition is the counterpart of Proposition 5.

Proposition 10. *Assume **H1**, **H5** and **H2**(R) for some $R > 0$. Let $\alpha > 0$ and $\varepsilon \in (0, \varepsilon_0]$. If there exists $\mathfrak{m} > 0$ such that $\log(p(y|\cdot))$ 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 (18) and $\bar{\delta} = \mathfrak{m}(L_y + \alpha L/\varepsilon)^{-2}/2$.

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

We recall the assumption on g_ε which ensures that $x \mapsto \log(p_\varepsilon(x))$ has Lipschitz gradients.

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 (9).

The following proposition is the counterpart of Proposition 6.

Proposition 11. *Assume **H1**, **H5**, **H2**(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 $\log(p(y|\cdot))$ 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 $\delta \in (0, \bar{\delta}]$*

$$\|\pi_{\varepsilon, \delta} - \pi_\varepsilon\|_V \leq B_1(\delta^{1/2} + \mathfrak{M}_R + \exp[-R]), \quad (19)$$

where V is given in (18) and B_1 does not depend on R .

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

The bound appearing in (19) depends on an extra hyperparameter $R > 0$ which may be optimized if **H2**(R) holds for any $R > 0$ and $\{\mathfrak{M}_R : R > 0\}$ can be expressed in a closed form. In particular if there exists $\mathfrak{M} \in (0, 1)$ such that for any $R > 0$, $\mathfrak{M}_R = \mathfrak{M} \times R$ then there exists $B_1 \geq 0$ such that for any $\delta \in (0, \bar{\delta}]$ and $R > 0$

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

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

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

by letting $R \rightarrow +\infty$.

We now combine Proposition 10 and Proposition 11 in order to control the bias of the Monte Carlo estimator obtained using PnP-ULA. This proposition is the counterpart of Proposition 7.

Proposition 12. Assume **H1**, **H5**, **H2(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 $\log(p(y|\cdot))$ is \mathfrak{m} -concave with $\mathfrak{m} \geq (2\alpha/\varepsilon) \max(\mathbf{L}, 1 + \mathbf{K}_\varepsilon/\varepsilon)$ and $\bar{\delta} = \mathfrak{m}(\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon)^{-2}/2$, 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

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

Proof. The proof is straightforward upon combining Proposition 10 and Proposition 11. \square

In particular, applying Proposition 12 to the family $\{h_i\}_{i=1}^d$ where for any $i \in \{1, \dots, d\}$, $h_i(x) = x_i$ we get that

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

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

D 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 13. 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 $\alpha > 0$.

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 14. 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 $\|\pi_\varepsilon - \pi\|_{\text{TV}} \leq A_0 \varepsilon^{\beta \min(\alpha, 1)/2}$.

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 [82, Corollary 1] with explicit dependency with respect to the dimension d . However, note that Proposition 14 differs from [82, Corollary 1] since the Gaussian smoothing approximation is applied to the prior distribution and the estimate is given on the posterior distribution in Proposition 14, whereas in [82, 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 12 and Proposition 7. 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 14. In order to state this proposition, we recall the following assumption which is a relaxation of the strongly log-concave condition.

H3. *There exists $\mathfrak{m} \in \mathbb{R}$ such that for any $x_1, x_2 \in \mathbb{R}^d$ we have*

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

Note that the posterior is strongly log-concave if and only if $\mathfrak{m} > 0$.

Proposition 15. *Assume H1, H5, H2, H4 and H3. 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 $2\lambda(L_y + (\alpha/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon) - \min(\mathfrak{m}, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$, there exists $C_{1,\varepsilon} \geq 0$ such that for any \mathcal{C} convex compact with $\bar{B}(0, R_C) \subset \mathcal{C}$ and $R_C > 0$, there exists $C_{2,\varepsilon,\mathcal{C}} \geq 0$ 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*

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

In addition, if there exists $\mathfrak{m} > 0$ such that $\log(p(y|\cdot))$ 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 $\sup_{x \in \mathbb{R}^d} \{|h(x)| (1 + \|x\|^2)^{-1}\} \leq 1$, $n \in \mathbb{N}^$, $\delta \in (0, \bar{\delta}]$ and $R > 0$ we have*

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

Proof. In the general case where $\log(p(y|\cdot))$ is not assumed to be \mathfrak{m} -concave with $\mathfrak{m} > 0$, the proof is completed upon combining Proposition 7, Proposition 14 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 $\log(p(y|\cdot))$ is \mathfrak{m} -concave upon replacing Proposition 7 by Proposition 12. \square

E 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

$$d\mathbf{X}_t = b(\mathbf{X}_t)dt + \sqrt{2}d\mathbf{B}_t, \quad (20)$$

where $(\mathbf{B}_t)_{t \geq 0}$ is a d -dimensional Brownian motion and $\mathbf{X}_0 = x$. In this case, (20) 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 (20) 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. [61] 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 16. *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*

$$\langle b(x_1), x_1 \rangle \leq -m \|x_1\|^2 + c , \quad \|b(x_1) - b(x_2)\| \leq L \|x_1 - x_2\| . \quad (21)$$

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 (21), 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 $\|x\| \geq (2c/(\varepsilon m))^{1/2}$

$$\|\mathcal{T}_\gamma(x)\| = \left(\|x\|^2 + 2\gamma \langle b(x), x \rangle + \gamma^2 \|b(x)\|^2 \right)^{1/2} \quad (22)$$

$$\begin{aligned}
&\leq \left((1 - 2\gamma\mathfrak{m} + (1 + \varepsilon)\gamma^2\mathbf{L}^2) \|x\|^2 + 2\gamma\mathfrak{c} + (1 + \varepsilon^{-1})\gamma^2 \|b(0)\|^2 \right)^{1/2} \\
&\leq \left((1 - \gamma\mathfrak{m} + (1 + \varepsilon)\gamma^2\mathbf{L}^2) \|x\|^2 + (1 + \varepsilon^{-1})\gamma^2 \|b(0)\|^2 \right)^{1/2} \\
&\leq \exp[-\gamma((2 - \varepsilon)\mathfrak{m} - (1 + \varepsilon)\mathbf{L}^2\bar{\gamma})/2] \|x\| + (1 + \varepsilon^{-1/2})\gamma \|b(0)\|.
\end{aligned}$$

Note that $(2 - \varepsilon)\mathfrak{m} - (1 + \varepsilon)\mathbf{L}^2\bar{\gamma} < 0$ since $\varepsilon \in (0, 1/2)$ and $\bar{\gamma} = \mathfrak{m}/\mathbf{L}^2$. On the other hand using (21) 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 $\|x\| \leq (2\mathfrak{c}/(\varepsilon\mathfrak{m}))^{1/2}$

$$\begin{aligned}
\|\mathcal{T}_\gamma(x)\| &\leq (1 + \gamma\mathbf{L}) \|x\| + \gamma \|b(0)\| \\
&\leq \exp[-\gamma((2 - \varepsilon)\mathfrak{m} - (1 + \varepsilon)\mathbf{L}^2\bar{\gamma})/2] \|x\| \\
&\quad + (2\mathfrak{c}/(\varepsilon\mathfrak{m}))^{1/2} \{ \exp[\gamma\mathbf{L}] - \exp[-\gamma((2 - \varepsilon)\mathfrak{m} - (1 + \varepsilon)\mathbf{L}^2\bar{\gamma})/2] \} + \gamma \|b(0)\| \\
&\leq \exp[-\gamma((2 - \varepsilon)\mathfrak{m} - (1 + \varepsilon)\mathbf{L}^2\bar{\gamma})/2] \|x\| + \gamma(2\mathfrak{c}/(\varepsilon\mathfrak{m}))^{1/2} \exp[\bar{\gamma}\mathbf{L}](\mathbf{L} + 2\mathfrak{m}) + \gamma \|b(0)\|.
\end{aligned} \tag{23}$$

Combining (22) and (23), there exist $\lambda \in [0, 1)$ and $c \geq 0$ such that for any $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$,

$$\|\mathcal{T}_\gamma(x)\| \leq \lambda^\gamma \|x\| + \gamma c. \tag{24}$$

Note that using (24), for any $k \in \{1, \dots, 2\varpi\}$ there exist $\tilde{\lambda}_k \in (0, 1)$ and $\tilde{c}_k \geq 0$ such that

$$\begin{aligned}
\|\mathcal{T}_\gamma(x)\|^k &\leq \{\tilde{\lambda}_k^\gamma \|x\| + \gamma\tilde{c}_k\}^k \\
&\leq \tilde{\lambda}_k^{\gamma k} \|x\|^k + \gamma 2^k \max(\tilde{c}_k, 1)^k \max(\bar{\gamma}, 1)^{k-1} \{1 + \|x\|^{k-1}\} \\
&\leq \tilde{\lambda}_k^\gamma \|x\|^k + \tilde{c}_k \gamma \{1 + \|x\|^{k-1}\} \leq (1 + \|x\|^k)(1 + \tilde{c}_k \gamma).
\end{aligned} \tag{25}$$

Therefore, combining (25) and the Cauchy-Schwarz inequality we obtain that for any $\gamma \in (0, \bar{\gamma}]$ and $x \in \mathbb{R}^d$

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

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 (21), 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 $\bar{\gamma} = \mathfrak{m}/(4L^2)$ 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 [18, 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 (21) we obtain that for any $x \in \mathbb{R}^d$ with $\|x\| \geq \sqrt{2}(1 + (c + 1 + \varpi)/\mathfrak{m})$

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

which concludes the proof. \square

Lemma 17. 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 [16, 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

Proposition 18. 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 $y_1, y_2 \in \mathbb{R}^m$

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

and for any $c > 0$, $\int_{\mathbb{R}^d} (1 + \Phi_1(\tilde{x})) \exp[c\Phi_1(\tilde{x})] p(x) dx < +\infty$. 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

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

Proof. Let $y_1, y_2 \in \mathbb{K}$ with \mathbb{K} a compact set. Let $y_0 \in \mathbb{K}$ and $D_{\mathbb{K}}$ be the diameter of \mathbb{K} . Using Lemma 22 we get that

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

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

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

which concludes the proof. \square

F Proofs of Section 3.2

We recall that the Markov chain $(X_k)_{k \in \mathbb{N}}$, defined in (17), is given by

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

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

In this section, we prove the convergence of PnP-ULA and control the bias of its invariant measure in the general framework introduced in Appendix B (*i.e.* $\alpha \neq 1$) under two different assumptions on the posterior: either the posterior is log-concave as in Appendix C or the posterior satisfies a more general one-sided Lipschitz condition as in Section 3.2. Note that in Section 3.2 the results are only stated for $\alpha = 1$. The statements of the propositions can be generalized to $\alpha > 0$ by replacing $2\lambda(L_y + L/\varepsilon - \min(\mathbf{m}, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + L/\varepsilon + 1/\lambda)^{-1}$ by $2\lambda(L_y + \alpha L/\varepsilon - \min(\mathbf{m}, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$ in Proposition 5 and $2\lambda(L_y + (a/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon) - \min(\mathbf{m}, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + L/\varepsilon + 1/\lambda)^{-1}$ by $2\lambda(L_y + (\alpha/\varepsilon) \max(L, 1 + K_\varepsilon/\varepsilon) - \min(\mathbf{m}, 0)) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$ in Proposition 6 and Proposition 7.

F.1 Proof of Proposition 4

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

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

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

$$\mathbb{E}[\|f_{w^\dagger}(X_\varepsilon) - D_\varepsilon^*(X_\varepsilon)\|] \leq \sqrt{\eta}. \quad (26)$$

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

$$\|f_{w^\dagger}(x_2) - D_\varepsilon^*(x_2)\| - \|f_{w^\dagger}(x_1) - D_\varepsilon^*(x_1)\| \leq C_R \|x_2 - x_1\|. \quad (27)$$

Assume that $\sup_{\tilde{x} \in \overline{\mathcal{B}}(0, R)} \|f_{w^\dagger}(\tilde{x}) - D_\varepsilon^*(\tilde{x})\| > \eta^\varpi$ with $\varpi = (2d + 2)^{-1}$ and denote $x_R \in \overline{\mathcal{B}}(0, R)$ such that we have $\sup_{\tilde{x} \in \overline{\mathcal{B}}(0, R)} \|f_{w^\dagger}(\tilde{x}) - D_\varepsilon^*(\tilde{x})\| = \|f_{w^\dagger}(x_R) - D_\varepsilon^*(x_R)\|$. Using (27) we have

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

Combining this result and (26) we obtain that

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

Setting $\mathbb{M}_R = \eta^{1/2} (\int_{\bar{B}(0,2R) \cap \bar{B}(x_R, C_R^{-1}\eta^\varpi)} p_\varepsilon(\tilde{x}) d\tilde{x})^{-1} + \eta^\varpi$ concludes the first part of the proof. Denote v_d the volume of the unit d -dimensional ball. We have that $\text{Leb}(\bar{B}(x_R, C_R^{-1}\eta^\varpi)) = C_R^{-d} \eta^{\varpi d} v_d$. Using the Fubini theorem, the Lebesgue differentiation theorem [14, Theorem 5.6.2], the dominated convergence theorem and the fact that for $\eta \in (0, (C_R R)^{1/\varpi}]$, $\bar{B}(0, 2R) \cap \bar{B}(x_R, C_R^{-1}\eta^\varpi) = \bar{B}(x_R, C_R^{-1}\eta^\varpi)$ we get that

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

Using this result we have,

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

which concludes the proof.

F.2 Proof of Proposition 5 and Proposition 10

We divide this section into two parts. First, we prove the general case where $\log(p(y|\cdot))$ is not assumed to be strongly concave but only satisfying a one-sided Lipschitz condition, *i.e.* Proposition 5. Then we turn to the proof of Proposition 10.

(a) Let $\lambda > 0$ such that $2\lambda(L_y + \alpha L/\varepsilon) \leq 1$ and $\bar{\delta} = (1/3)(L_y + \alpha L/\varepsilon + 1/\lambda)^{-1}$. Let \mathbf{C} be a compact convex set with $0 \in \mathbf{C}$. Using H2, (17) and that $\text{Id} - \Pi_{\mathbf{C}}$ is non-expansive we have for any $x_1, x_2 \in \mathbb{R}^d$

$$\|b_\varepsilon(x_1) - b_\varepsilon(x_2)\| \leq (L_y + \alpha L/\varepsilon + 1/\lambda) \|x_1 - x_2\| .$$

Denote $R_{\mathbf{C}} = \sup\{\|x_1 - x_2\| : x_1, x_2 \in \mathbf{C}\}$. Using (17), the Cauchy-Schwarz inequality and that $2\lambda(\alpha L/\varepsilon - \mathbf{m}) \leq 1$ we have for any $x_1, x_2 \in \mathbb{R}^d$

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

Hence, for any $x_1, x_2 \in \mathbb{R}^d$ with $\|x_1 - x_2\| \geq 4R_{\mathbf{C}}$ 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 $x \in \mathbb{R}^d$

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

We conclude the proof of Proposition 5 upon using Lemma 16, Lemma 17, [16, Corollary 2] with $\bar{\gamma} \leftarrow (4\lambda)^{-1}(\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon + 1/\lambda)^{-2} \geq \bar{\delta}$ and the fact that for any probability distribution ν_1, ν_2 ,

$$\|\nu_1 - \nu_2\|_V \leq \|\nu_1 - \nu_2\|_{TV}^{1/2} (\nu_1[V^2] + \nu_2[V^2])^{1/2}. \quad (28)$$

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

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

This concludes the proof of Proposition 10 upon using [16, Corollary 2] with $\bar{\gamma} \leftarrow \mathbf{m}(\mathbf{L}_y + \alpha\mathbf{L}/\varepsilon)^{-2} \geq \bar{\delta}$ and (28).

F.3 Proof of Proposition 6 and Proposition 11

Before proving Proposition 6 and Proposition 11, we show the following lemma which is a straightforward consequence of Girsanov's theorem [57, Theorem 7.7]. A similar version of this lemma can be found in the proof of [32, Proposition 2].

Lemma 19. *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 $(P_t^{(i)})_{t \geq 0}$ 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 $\mathbb{P}(\int_0^T \{\|b_i(t, \mathbf{X}_t^{(i)})\|^2 + \|b_i(t, \mathbf{B}_t)\|^2\} dt < +\infty) = 1$. Let $V : \mathbb{R}^d \rightarrow [0, +\infty)$ measurable, then for any $x \in \mathbb{R}^d$ we have*

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

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 [32, Lemma 24] and the transfer theorem [53, Theorem 4.1] we get that

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

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 [57, Theorem 7.7] and μ_B -almost surely for any $w \in C([0, T], \mathbb{R})$ we get

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

Hence, we obtain that

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

which concludes the proof. \square

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

Lemma 20. *Assume **H4**. Then for any $x_1, x_2 \in \mathbb{R}^d$ we have*

$$\|\nabla \log(p_\varepsilon(x_1)) - \nabla \log(p_\varepsilon(x_2))\| \leq (1 + K_\varepsilon/\varepsilon) \|x_1 - x_2\| / \varepsilon .$$

*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 $x \in \mathbb{R}^d$

$$\begin{aligned} \nabla^2 \log(p_\varepsilon(x)) &= -\varepsilon^{-1} \text{Id} + \varepsilon^{-2} \int_{\mathbb{R}^d} (x - \tilde{x})^{\otimes 2} g_\varepsilon(\tilde{x}|x) d\tilde{x} - \varepsilon^{-2} \left(\int_{\mathbb{R}^d} (x - \tilde{x}) g_\varepsilon(\tilde{x}|x) d\tilde{x} \right)^{\otimes 2} \\ &= -\varepsilon^{-1} \text{Id} + \varepsilon^{-2} \int_{\mathbb{R}^d} \left(\tilde{x} - \int_{\mathbb{R}^d} \tilde{x}' g_\varepsilon(\tilde{x}'|x) d\tilde{x}' \right)^{\otimes 2} g_\varepsilon(\tilde{x}|x) d\tilde{x} \end{aligned} \quad (29)$$

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

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

which concludes the first part of the proof. Reciprocally, since $x \mapsto \nabla \log(p_\varepsilon(x))$ is Lipschitz-continuous with constant $K \geq 0$ we get that for any basis vector $(e_i)_{i \in \{1, \dots, d\}}$ we have that $e_i^\top \nabla^2 \log(p_\varepsilon(x)) e_i \leq K$. Combining this result with (29), we get that

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

which concludes the proof. \square

In what follows we prove Proposition 6. The proof of Proposition 11 is similar and left to the reader.

Proof of Proposition 6. 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 \mathbb{C} convex compact with $0 \in \mathbb{C}$ there exists $B_{1,\mathbb{C}} \geq 0$ such that for any $\delta \in (0, \bar{\delta}]$ and $R > 0$

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

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

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

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

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

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

$$\begin{aligned} d\bar{\mathbf{X}}_t &= \bar{b}_\varepsilon(\bar{\mathbf{X}}_t)dt + \sqrt{2}d\mathbf{B}_t, \\ \bar{b}_\varepsilon(x) &= \nabla \log(p(y|x)) + \alpha \nabla \log(p_\varepsilon(x)) + \text{prox}_\lambda(\iota_C)(x), \end{aligned} \quad (30)$$

with $(\mathbf{B}_t)_{t \geq 0}$ a d -dimensional Brownian motion. \bar{b}_ε is Lipschitz continuous using Lemma 20, hence this SDE admits a unique strong solution for any initial condition \mathbf{X}_0 with $\mathbb{E}[\|\mathbf{X}_0\|^2] < +\infty$, see [49, Chapter 5, Theorem 2.9]. We denote by $(P_{t,\varepsilon})_{t \geq 0}$ the semigroup associated with the strong solutions of (30). Similarly to the proof of Proposition 10, replacing [16, Corollary 2] by [16, 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 $t \geq 0$

$$\begin{aligned} \|\delta_{x_1}P_{t,\varepsilon} - \delta_{x_2}P_{t,\varepsilon}\|_V &\leq \tilde{A}_C \tilde{\rho}_C^t (V^2(x_1) + V^2(x_2)), \\ \mathbf{W}_1(\delta_{x_1}P_{t,\varepsilon}, \delta_{x_2}P_{t,\varepsilon}) &\leq \tilde{A}_C \tilde{\rho}_C^t \|x_1 - x_2\|. \end{aligned} \quad (31)$$

Combining (31), Proposition 10, 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 $x \in \mathbb{R}^d$

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

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

$$\left\| \delta_x R_{\varepsilon,1/m}^{km}[f] - \delta_x P_{km,\varepsilon}^{km}[f] \right\| = \left\| \sum_{j=0}^{k-1} \delta_x R_{\varepsilon,1/m}^{jm} (R_{\varepsilon,1/m}^m - P_{1,\varepsilon}) P_{k-j-1,\varepsilon}[f] \right\| \quad (32)$$

Using (31), Lemma 16 and Lemma 17 there exists $B_a \geq 0$ such that for any $x \in \mathbb{R}^d$ and $k \in \mathbb{N}$ we have

$$\|\delta_x P_{k,\varepsilon,C}[f] - \tilde{\pi}_\varepsilon[f]\| \leq B_a \tilde{\rho}_C^k V^2(x). \quad (33)$$

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

$$\left\| \delta_x R_{\varepsilon,1/m,C}^m - \delta_x P_{1,\varepsilon,C} \right\|_{\text{TV}} = \left\| \delta_x P_1^{(1)} - \delta_x P_1^{(2)} \right\|_{\text{TV}}. \quad (34)$$

Using H2(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 $(w_t)_{t \in [0,1]} \in C([0, 1], \mathbb{R}^d)$

$$\begin{aligned} \|b_1(t, (w_t)_{t \in [0,1]}) - b_2(t, (w_t)_{t \in [0,1]})\|^2 &= \|b_\varepsilon(w_{j/m}) - \bar{b}_\varepsilon(w_t)\|^2 \\ &\leq 2 \|b_\varepsilon(w_{j/m}) - b_\varepsilon(w_t)\|^2 + 2 \|\bar{b}_\varepsilon(w_t) - b_\varepsilon(w_t)\|^2 \\ &\leq 2L_b^2 \|w_{j/m} - w_t\|^2 + 4\alpha^2 M_R^2 / \varepsilon^2 + 4\alpha^2 \mathbf{1}_{\bar{B}(0,R)^c}(\|w_t\|) / \varepsilon^2, \end{aligned} \quad (35)$$

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

$$\mathbb{E}[\|\mathbf{X}_t^{(1)} - \mathbf{X}_{j/m}^{(1)}\|^2] = 2\mathbb{E}[\|\int_{j/m}^t d\mathbf{B}_t\|^2] \leq 2d\delta. \quad (36)$$

Finally, using Lemma 16, Lemma 17, the logarithmic Sobolev inequality [18, 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 $x \in \mathbb{R}^d$

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

where $\ell_t = \lfloor tm \rfloor$ and $\Phi(x) = \sqrt{1 + \|x\|^2}$. Combining this result, (35), (34), (36) and Lemma 19, we obtain that there exists $B_b \geq 0$ such that for any $x \in \mathbb{R}^d$ and $R > 0$

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

Combining this result and (33) 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

$$\left| (\delta_x \mathbf{R}_{1/m, \mathbb{C}}^m - \delta_x \mathbf{P}_{1, \mathbb{C}}) \mathbf{P}_{k-j-1, \mathbb{C}}[f] \right| \leq B_a B_b (\sqrt{\delta} + \mathbf{M}_R + \exp[-R]) \tilde{\rho}_{\mathbb{C}}^{k-j-1} \exp[2\Phi(x)] .$$

Using this result, Lemma 16, Lemma 17 and (32) we obtain that there exists $B_c \geq 0$ such that for any $m \in \mathbb{N}^*$ with $m^{-1} \geq \bar{\delta}$

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

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 $\delta \in (0, \bar{\delta}]$

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

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

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

Similarly, define ρ_{ε} and π_{ε} such that for any $x \in \mathbb{R}^d$

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

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

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

with X a random variable with distribution π_{ε} . We conclude using the generalized Pinsker inequality [32, Lemma 24].

□

G Proofs of Section 3.3

G.1 Proof of Proposition 8

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 [16, Section 3] with initial condition $(x_1, x_2) \in C$. Using [16, Corollary 7-(b)] we get that for any $\ell \in \mathbb{N}$

$$\mathbb{E} \left[\mathbf{1}_{\Delta_{\mathbb{R}^d}^c}(X_{(\ell+1)\lceil 1/\delta \rceil}, Y_{(\ell+1)\lceil 1/\delta \rceil}) \right] \leq (1 - \beta) \mathbb{E} \left[\mathbf{1}_{\Delta_{\mathbb{R}^d}^c}(X_{\lceil 1/\delta \rceil}, Y_{\lceil 1/\delta \rceil}) \right], \quad (37)$$

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 (37), we get that for any $k \in \mathbb{N}$

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

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 $\tilde{A}_C = (1 - \beta)^{-1}$.

G.2 Proof of Proposition 9

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

$$\|b_\varepsilon(x_1) - b_\varepsilon(x_2)\| \leq (L_y + \alpha L/\varepsilon + 1/\lambda) \|x_1 - x_2\|.$$

Using this result, the fact that for any $x \in \mathbb{R}^d$, $\langle b_\varepsilon(x), x \rangle \leq -\tilde{m}\|x\|^2 + c$ and [31, 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 $k \in \mathbb{N}$

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

with \tilde{B} and $\tilde{\rho}$ which do not depend on R . In addition, using Lemma 16, for any $k \in \mathbb{N}$ and $\delta \in (0, \bar{\delta}_2]$ we have

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

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 $x \in \mathbb{R}^d$, $R_{\varepsilon, \delta} V(x) \leq A^\delta V(x)$. 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 $k \in \mathbb{N}$

$$\|\pi_{\varepsilon, \delta} - \pi_{\varepsilon, \delta}^{\mathbb{C}}\|_{\text{TV}} \leq (V(0) + \tilde{c}\delta_2)A^{k\delta} \exp[-R] + \tilde{B}\tilde{\rho}^{k\delta}V(0).$$

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

H Proofs of Appendix D

H.1 Proof of Proposition 13

The first part of the proposition is straightforward. Using Pinsker's inequality [18, Theorem 4.19] we have for any $x \in \mathbb{R}^d$

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

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.

H.2 Proof of Proposition 14

First we show the following technical lemma.

Lemma 21. *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 14 we state the following lemma.

Lemma 22. *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 14.

Proof. Let $\alpha > 0$. For any $\varepsilon > 0$ and $x \in \mathbb{R}^d$ denote $\bar{p}(x) = p(y|x)p^\alpha(x)$ and $\bar{p}_\varepsilon(x) = (py|x)p_\varepsilon^\alpha(x)$, 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

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

Using Lemma 21 and that $\|p_\varepsilon\|_\infty \leq \|p\|_\infty < +\infty$, we have for any $\varepsilon > 0$ and $x \in \mathbb{R}^d$

$$\begin{aligned} \int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx &\leq 2^\alpha \|p(y|\cdot)\|_\infty (1 + \|p\|_\infty^{(\alpha-1)\wedge 0}) \\ &\quad \times \left\{ \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| dx + \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)|^\alpha dx \right\}. \end{aligned} \quad (38)$$

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 (38) we get that

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

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

$$\begin{aligned} \int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx &\leq 2^\alpha \|p(y|\cdot)\|_\infty (1 + \|p\|_\infty^{(\alpha-1)\wedge 0}) \\ &\quad \times \left\{ \int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| dx + \left(\int_{\mathbb{R}^d} |p(x) - p_\varepsilon(x)| (x) dx \right)^\alpha \right\}. \end{aligned} \quad (39)$$

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, (39) and Lemma 22 concludes the first part of the proof.

Now let $\alpha > 0$ and assume **H6**(α). If $\alpha \geq 1$ then using (38) we have

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

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

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

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

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

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

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

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

$$\int_{\mathbb{R}^d} |\bar{p}(x) - \bar{p}_\varepsilon(x)| dx \leq C_1 (\varepsilon^{\beta/2} + \varepsilon^{\beta\alpha/2}), \quad (40)$$

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 (40) with Lemma 22, 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