

Accelerated Bayesian imaging by relaxed proximal-point Langevin sampling

Teresa Klatzer^{1,3} Paul Dobson^{1,3} Yoann Altmann⁴ Marcelo Pereyra^{2,3}
 Jesús María Sanz-Serna⁵ Konstantinos C. Zygalakis^{1,3}

August 21, 2023

Abstract

This paper presents a new accelerated proximal Markov chain Monte Carlo methodology to perform Bayesian inference in imaging inverse problems with an underlying convex geometry. The proposed strategy takes the form of a stochastic relaxed proximal-point iteration that admits two complementary interpretations. For models that are smooth or regularised by Moreau-Yosida smoothing, the algorithm is equivalent to an implicit midpoint discretisation of an overdamped Langevin diffusion targeting the posterior distribution of interest. This discretisation is asymptotically unbiased for Gaussian targets and shown to converge in an accelerated manner for any target that is κ -strongly log-concave (i.e., requiring in the order of $\sqrt{\kappa}$ iterations to converge, similarly to accelerated optimisation schemes), comparing favorably to [M. Pereyra, L. Vargas Miele, K.C. Zygalakis, SIAM J. Imaging Sciences, 13,2 (2020), pp. 905-935] which is only provably accelerated for Gaussian targets and has bias. For models that are not smooth, the algorithm is equivalent to a Leimkuhler–Matthews discretisation of a Langevin diffusion targeting a Moreau-Yosida approximation of the posterior distribution of interest, and hence achieves a significantly lower bias than conventional unadjusted Langevin strategies based on the Euler-Maruyama discretisation. For targets that are κ -strongly log-concave, the provided non-asymptotic convergence analysis also identifies the optimal time step which maximizes the convergence speed. The proposed methodology is demonstrated through a range of experiments related to image deconvolution with Gaussian and Poisson noise, with assumption-driven and data-driven convex priors.

1 Introduction

The problem of estimating an unknown image from noisy measurements is central to imaging sciences (see, e.g., [17, 7, 3], for examples related to image denoising, inpainting, tomographic reconstruction, medical imaging, resolution enhancement and computer vision). Estimation by direct inversion of the forward model relating the unknown image to the data is not usually possible, since the inverse problem is often severely ill-conditioned or ill-posed. The literature describes a range of mathematical frameworks to incorporate regularisation and formulate well-posed solutions (see, e.g., [17, 7, 3]).

We consider imaging methodology rooted in the Bayesian statistical framework [17, Chapter 3]. The Bayesian statistical framework is an intrinsically probabilistic paradigm in which statistical models represent the data observation process and the prior information available, and solutions are delivered using inference techniques from Bayesian decision theory [36]. In particular, the Bayesian decision-theoretic approach provides a range of powerful strategies to quantify and summarise the uncertainty in the solutions

¹School of Mathematics, University of Edinburgh, Edinburgh, Scotland, UK

²School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh, Scotland, UK

³Maxwell Institute for Mathematical Sciences, Bayes Centre, 47 Potterrow, Edinburgh, Scotland, UK

⁴School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh Scotland, UK

⁵Departamento de Matemáticas, Universidad Carlos III de Madrid, Leganés (Madrid), Spain

delivered, which is important for applications that use the restored images as evidence for decision-making or in subsequent scientific analysis (see, e.g., [11, 28, 2, 34, 44, 43, 24]).

Applying Bayesian strategies to imaging inverse problems involves some significant modelling and computational challenges. This paper focuses on the computational aspects of performing Bayesian inference in imaging problems with an underlying convex geometry (i.e., with a posterior distribution that is log-concave), for which we can provide detailed theoretical guarantees. We consider both assumption-driven [7] and data-driven strategies [13, 19, 26], with particular attention to situations with no or very limited ground truth data available, where adopting a large data-driven model would not be possible.

Modern approaches for Bayesian computation in imaging inverse problems can be broadly classified into the following three groups. The first group encompasses approximation methods, such as message-passing and variational Bayes algorithms (see, e.g., [30, 43, Section 3]) which include the increasingly prevalent denoising diffusion approaches and other data-driven strategies based on large machine learning models (see, e.g., [18, 2]). Strategies in this first group are computationally efficient and can support a variety of inferences but are relatively model specific, offer limited guarantees, and can exhibit convergence issues. The second group encompasses deterministic and stochastic optimisation strategies [30, Section 4], including proximal convex optimisation approaches [7]. These are used predominantly to compute maximum-a-posteriori solutions but can also be useful for some basic forms of uncertainty quantification (see, e.g., [28, 35]). Within their limited scope of action, optimisation approaches scale efficiently to large problems and offer strong guarantees on the solutions delivered. The third group encompasses Monte Carlo integration methods [30], including modern Markov chain Monte Carlo (MCMC) strategies based on approximations of the Langevin diffusion [30, 11, 40, 31, 24, 20]. Monte Carlo strategies are very versatile and can support complex Bayesian inferences, at the expense of a significantly higher computational cost. Some modern MCMC strategies also provide strong guarantees on the delivered solutions [11, 20].

In this paper, we focus on MCMC methodology for Bayesian inference in imaging problems with an underlying convex geometry. Efficient MCMC simulation in this context is challenging for three main reasons: the high dimensionality of the solution space, lack of smoothness, and poor conditioning (i.e., the posterior density is highly anisotropic). Modern MCMC methods typically address the high dimensionality involved by mimicking a Langevin diffusion process that exploits first-order (gradient) information to scale to large problems very efficiently [30, Section 2]. Unfortunately, Bayesian imaging models often violate the regularity conditions that Langevin strategies require in order to be effective. This difficulty can be addressed by exploiting ideas and techniques from non-smooth convex optimisation as proposed in [27], resulting in proximal MCMC methods such as the Moreau Yoshida Unadjusted Langevin Algorithm (MYULA) [11] and the Split-and-Augmented Gibbs sampler [40, 32] (see Section 2.3 for more details). It is also possible to derive proximal MCMC methods from a reflected Langevin diffusion that naturally incorporates constraints in the solution space [24].

MCMC methods derived from the Langevin diffusion process can become pathologically slow when the posterior distribution of interest is highly anisotropic, as often encountered in problems that are severely ill-conditioned or ill-posed. Indeed, similarly to gradient descent methods, standard Langevin sampling methods have a computational cost in the order of κ for targets that are κ -strongly log-concave. This phenomenon can be demonstrated empirically on simple Gaussian models (see [31, Figure 2]) as well as established theoretically by conducting a non-asymptotic convergence analysis (see, e.g., [11]). This analysis also reveals that the difficulties in dealing with $\kappa \gg 1$ stem from the choice of the discrete-time approximation of the Langevin diffusion used, based on a Euler-Maruyama scheme, and not from the continuous-time process itself.

Overcoming this type of behaviour with respect to the condition number κ is very important and there are currently two main strategies to improve the convergence speed when $\kappa \gg 1$. One strategy is to exploit the structure of the problem in order to use some form of preconditioning [9, 41, 32]. This can be highly effective in some cases, but it is difficult to apply generally. A promising alternative is to leverage ideas from the numerical analysis of stochastic differential equations in order to develop MCMC methods from better discrete-time approximation of the Langevin diffusion that are inherently more robust to $\kappa \gg 1$. A

prime example of this approach is SKROCK [31], which exhibits *acceleration*¹ in that it only requires in the order of $\sqrt{\kappa}$ gradient evaluations to converge and hence dramatically outperforms MYULA when κ is large. However, theoretically analysing SKROCK is notoriously difficult and current convergence results only hold for Gaussian target densities.

In this paper, we develop and study the Implicit Midpoint Langevin Algorithm (IMLA) for computation in Bayesian imaging models that are log-concave. Each iteration of IMLA can be viewed as a relaxed proximal-point step with a stochastic perturbation. We show that IMLA is asymptotically unbiased for Gaussian target distributions and converges in a provably accelerated manner for all distributions that are κ -strongly log-concave and smooth (or that are regularised by smoothing, e.g., by using a Moreau-Yosida approximation as described in Section 2), providing a rigorous alternative to SKROCK [31] which is only provably accelerated for Gaussian models and has bias.

The remainder of this paper is structured as follows: Section 2 introduces the class of Bayesian models considered as well as the overdamped Langevin diffusion process and its numerical approximations for models that are smooth and non-smooth. Section 3 presents the proposed IMLA method and provides detailed theoretical convergence guarantees for Gaussian and strongly log-concave models, including the optimal time-step choice to optimise convergence speed; the presented theory is illustrated with a range of toy examples. Following on from this, Section 4 demonstrates IMLA on two challenging image restoration problems and by comparisons with SKROCK. Conclusions and perspectives for future work are finally reported in Section 5.

2 Problem Statement

2.1 Bayesian inference and imaging inverse problems

We consider imaging inverse problems where we seek to estimate the unknown image $x \in \mathcal{X}$ from an observation y , related to x by a forward statistical model with likelihood function $p(y|x)$. Typical examples include:

- Gaussian likelihood $y \sim \mathcal{N}(Ax, \sigma^2\mathbb{I})$:

$$p(y|x) \propto \exp\left(-\frac{\|Ax - y\|_2^2}{2\sigma^2}\right), \quad \mathcal{X} = \mathbb{R}^d, \quad (1)$$

- Poisson likelihood $y \sim \mathcal{P}(Ax, \beta)$:

$$p(y|x) \propto \exp\left\{\sum_{i=1}^p [(Ax)_i + \beta - y_i \log((Ax)_i + \beta)]\right\}, \quad \mathcal{X} = \{x : (Ax)_i \geq 0\}, \quad (2)$$

where $A \in \mathbb{R}^{p \times d}$ is a linear operator which models physical properties of the observation process, $\sigma > 0$ in (1) is related to the intensity of the additive observation noise, and $\beta > 0$ in (2) represents a background Poisson noise level.

The operator $A \in \mathbb{R}^{p \times d}$ is usually rank-deficient or $A^T A$ has a poor condition number $\kappa \gg 1$, making the imaging problem ill-conditioned or ill-posed. To make the problem well-posed and deliver meaningful solutions one needs to incorporate regularisation [17]. The Bayesian statistical framework introduces regularisation by defining a prior distribution $p(x)$ describing expected properties of the unknown image x . Observed and prior information are then combined by using Bayes' Theorem to obtain the model's posterior distribution, with density given by [17]

$$\pi(x) := p(x|y) = \frac{p(y|x)p(x)}{\int_{\mathcal{X}} p(y|x)p(x)dx}. \quad (3)$$

¹We use the term acceleration in a manner akin to the convex optimisation literature; other works, e.g., [25], use this term to describe an improvement w.r.t. the problem dimension.

For Bayesian computation, it is convenient to reformulate π as follows:

$$\pi(x) \propto e^{-U(x)} = e^{-f_y(x)-g(x)}, \quad (4)$$

with $f_y : \mathcal{X} \rightarrow \mathbb{R}$ and $g : \mathcal{X} \rightarrow \mathbb{R}$; representing the data likelihood and the prior information respectively.

In this paper, we assume that f_y and g satisfy the following conditions:

1. f_y is convex with Lipschitz continuous gradient.
2. g is proper, convex and lower semi-continuous, but potentially non-smooth.

Requiring g to be convex on \mathcal{X} significantly restricts the class of data-driven priors that can be considered. However, convexity provides some significant benefits in terms of guaranteeing the well-posedness of $p(x|y)$ and of posterior moments of interest (i.e., existence and stability w.r.t. changes in y , see [16] for details). Convexity is also central to the analysis of the convergence properties of the Bayesian computation methods used to perform inference w.r.t. $p(x|y)$: convexity allows for faster convergence rates, tighter non-asymptotic bounds, and stronger guarantees on the delivered solutions. Choices of g that respect convexity include, e.g., the ℓ_1 -norm on some sparsifying basis, the total-variation pseudo-norm, as well as a number of recently proposed data-driven priors such as the convex ridge neural network regularizer [13], the convex variant of the variational network [19], or the input convex neural network [26]. Moreover, with regards to the Lipschitz condition on ∇f_y , even though this condition is not satisfied in the case of some important non-Gaussian likelihoods (e.g. Poisson with $\beta = 0$, Binomial, and Geometric), an appropriate approximation of the likelihood resolves this issue (see, e.g., [24]).

Drawing inferences about x from y by operating directly with $p(x|y)$ is usually not possible in imaging sciences because of the high dimensionality involved. As a result, imaging methods compute summaries of $p(x|y)$, such as Bayesian point estimators, posterior probabilities, moments, and expectations of interest. In particular, for the class of log-concave models considered here, a common choice is to compute the maximum-a-posteriori (MAP) estimator

$$\hat{x}_{\text{MAP}} = \arg \max_{x \in \mathcal{X}} \pi(x) = \arg \min_{x \in \mathcal{X}} \{f_y(x) + g(x)\}, \quad (5)$$

which is the optimal Bayesian estimator of x w.r.t. to the $(f_y + g)$ -Bregman loss [29] and can be efficiently computed by proximal convex optimisation techniques [7, 8]. Other Bayesian point estimators and more complicated Bayesian analyses such as uncertainty quantification, calibration of model parameters, model selection and hypothesis testing, are generally beyond the scope of optimisation algorithms as they require calculating probabilities and moments with respect to π . However, as mentioned previously, posterior probabilities and moments of interest can be computed efficiently by using proximal MCMC methods derived from the Langevin diffusion process [12]. These algorithms can be applied directly to perform inference w.r.t. π (see, e.g., [6, 24]), or embedded in more complex inference machinery to calibrate unknown model parameters [38] or perform model selection [5].

2.2 Sampling using the overdamped Langevin equation

The fastest provably accurate Bayesian computation algorithms to perform inference w.r.t. log-concave models such as π are derived from the overdamped Langevin SDE (OLSDE). Let $\bar{\pi}$ be a smooth density that is log-concave on \mathcal{X} and consider the OLSDE

$$dX_t = \nabla \log \bar{\pi}(X_t) dt + \sqrt{2} dW_t, \quad (6)$$

where $\{W_t\}_{t \geq 0}$ is a d -dimensional Brownian motion. Under mild conditions on $\bar{\pi}$, this SDE is well-posed and converges provably quickly to $\bar{\pi}$ (its unique invariant distribution), providing a foundation for fast and rigorous MCMC algorithms to perform computations for $\bar{\pi}$.

The OLSDE does not have an analytic solution in general so we consider numerical approximations of this SDE. The most widely used approximation is the Unadjusted Langevin Algorithm (ULA) given by

$$X_{n+1} = X_n + \delta \nabla \log \bar{\pi}(X_n) + \sqrt{2\delta} \xi_n, \quad (7)$$

where $\delta > 0$ and $\{\xi_n\}_{n \geq 1}$ is a sequence of d -dimensional standard Gaussian random variables. Despite its simplicity, ULA can be remarkably fast and scale efficiently to very large problems [11]. However, a necessary condition to establish the convergence of ULA is that $\delta < L^{-1}$ where L is the Lipschitz constant of $\nabla \log \bar{\pi}(x)$ [11]. Unfortunately, in problems that are ill-conditioned, as commonly encountered in imaging sciences, this leads to δ being very small in comparison to the time scales required for convergence to π , which can be shown to be of order $\kappa = L/m$ when π is m -strongly log-concave. The slow convergence leads to highly correlated samples and to an inefficient exploration of the solution space. This issue can be addressed by designing specialized stochastic integrators such as the SKROCK algorithm [1, 31], which behave in an accelerated manner in the sense that the iterates converge to π with a time scale of order $\sqrt{\kappa}$. Because often in Bayesian imaging problems $\sqrt{\kappa} \ll \kappa$, SKROCK can lead to very significant reductions in computing time w.r.t. ULA (e.g., [31, 32] report speed-up factors of the order of 20x - 30x).

Despite this remarkable empirical performance, very little is known about the theoretical convergence properties of SKROCK, with existing results only covering the case where $\bar{\pi}$ is Gaussian. As a result, whenever SKROCK is applied to a new imaging problem it is necessary to benchmark its performance against a more reliable method (e.g., ULA) on some test problems. In addition, the lack of theory for SKROCK makes it difficult to embed it within more complex inference schemes such as [38], or to provide practitioners with principled strategies to set its internal parameters. The method that we present in Section 3 seeks to address these issues by providing a new accelerated Langevin MCMC scheme that is empirically as computationally efficient as SKROCK, but has a much stronger theoretical footing. Before concluding this section, we briefly discuss the case where π is not smooth.

2.3 Non-smooth distributions

The Bayesian models encountered in imaging sciences are often not smooth (e.g., because of the use of non-smooth norms or because of constraints on the solution space), and this presents some challenges to the application of gradient-based strategies such as ULA and SKROCK that require smoothness. For models with an underlying convex geometry, proximal MCMC methods [27, 12] deal with the lack of smoothness by replacing π with a smooth approximation $\pi^\lambda(x)$ such as

$$\pi^\lambda(x) \propto \exp(-f_y(x) - g^\lambda(x)) \quad (8)$$

where g^λ is the Moreau-Yosida envelope of g defined by

$$g^\lambda(x) = \min_{u \in \mathbb{R}^d} \left\{ g(u) + \frac{1}{2\lambda} \|x - u\|^2 \right\}.$$

Approximations based on other splittings can be considered too (see, e.g., [38]). The density π^λ is log-concave and Lipschitz smooth, with gradient

$$\nabla \log \pi^\lambda(x) = -\nabla f_y(x) - \nabla g^\lambda(x) \quad (9)$$

$$= -\nabla f_y(x) - \frac{1}{\lambda} (x - \text{prox}_{\theta g}^\lambda(x)), \quad (10)$$

with Lipschitz constant $L \leq L_f + \lambda^{-1}$. Here prox_g^λ is the proximal operator of g given by

$$\text{prox}_g^\lambda(x) = \arg \min_{u \in \mathbb{R}^d} \left\{ g(u) + \frac{1}{2\lambda} \|x - u\|^2 \right\}. \quad (11)$$

Since the OLSDE with π replaced by π^λ is well-defined, we can consider sampling algorithms derived from discretisations of this SDE. In particular, applying an OLSDE-based algorithm such as ULA or SKROCK to π^λ is considered in [12, 31] as a means of performing approximate computation w.r.t. π . Indeed, since π^λ converges to π in Total Variation as $\lambda \rightarrow 0$ and can be made arbitrarily close to π by reducing λ , see [12, Proposition 1], the samples generated by targeting π^λ reliably approximate π .

It is worth mentioning at this point that both ULA and SKROCK produce biased Monte Carlo samples because of the error stemming from the discretisation of the OLSDE. Of course, there is also non-asymptotic bias stemming from the fact that the algorithms run for a finite number of iterations. In principle, the asymptotic bias related to the discretisation error could be removed by using a Metropolis-Hastings (MH) correction step, as in the Metropolis Adjusted Langevin Algorithm (MALA) [30]. However, in high-dimensional problems that are ill-conditioned or ill-posed, introducing an accept-reject mechanism can lead to drastically slower convergence and to a large increase in the non-asymptotic error (see, e.g., [12]). The development of unadjusted proximal Langevin methods with smaller bias is therefore valuable, particularly in imaging problems that require computing tail probabilities or high-order moments that are more sensitive to the estimation bias than the posterior mean.

3 Proposed method: IMLA

We are now ready to present our proposed accelerated Bayesian computation method to perform inference w.r.t. π , for which we are able to provide detailed convergence results including non-asymptotic convergence bounds and an optimal choice of the discretisation time step to maximize the convergence speed. The method is based on the following recursion:

$$\text{IMLA : } X_{n+1} = \left(1 - \frac{1}{\theta}\right) X_n + \frac{1}{\theta} \text{prox}_U^{\delta\theta}(X_n + \theta\sqrt{2\delta}\xi_n), \quad (12)$$

for $\theta \in [0, 1]$, $\delta > 0$ and $\{\xi_n\}_{n \geq 1}$ a sequence of independent d -dimensional standard Gaussian random variables. Recursion (12) is a stochastic modification of a relaxed proximal point algorithm to compute the MAP solution \hat{x}_{MAP} (definition see (5)), or equivalently to minimize the potential $U(\cdot) = -\log \pi(\cdot)$ [8].

In order to analyze (12) it is useful to express it explicitly as a minimisation problem as follows²:

$$\begin{aligned} X_{n+1} &= \text{argmin}_{x \in \mathbb{R}^d} F(x; X_n; \xi_{n+1}), \\ F(x; u, z) &:= \theta^{-1}U(\theta x + (1 - \theta)u) + \frac{1}{2\delta} \|x - u - \sqrt{2\delta}z\|^2. \end{aligned} \quad (13)$$

Note that the scheme (13) is well defined when U is convex or under the weaker condition $U \in \mathcal{C}^2$ with the second order derivative $\delta^{-1}I + \theta\nabla^2 U \succeq 0$.

When $U \in \mathcal{C}^1$, $\nabla F(X_{n+1}) = 0$ and therefore (12) and (13) are also equivalent to the following implicit (midpoint) approximation of an OLSDE (6) targeting the posterior density of interest $\pi(x) \propto \exp\{-U(x)\}$:

$$X_{n+1} = X_n - \delta\nabla U(\theta X_{n+1} + (1 - \theta)X_n) + \sqrt{2\delta}\xi_{n+1}. \quad (14)$$

The recursion (14) belongs to the class of numerical integration methods commonly known as “ θ -methods”, which combine explicit and implicit first-order gradient steps [14]. We are particularly interested in the case $\theta = 1/2$, where (14) is a stochastic modification of the implicit midpoint method, motivating the name of our proposed algorithm: Implicit Midpoint Langevin Algorithm (IMLA). Observe that when $\theta = 0$, (14) reduces to ULA, and when $\theta = 1$, this corresponds to the implicit Euler scheme applied to OLSDE [15](ILA).

Moreover, by using the convexity of U , the identity (9), and setting $\theta = 1/2$, we find that the sequence of iterates $Y_n = X_n + \theta\sqrt{2\delta}\xi_n$ is described by the following non-Markovian scheme:

$$Y_{n+1} = Y_n - \delta\nabla U^{\delta/2}(Y_n) + \sqrt{2\delta}\left(\frac{\xi_{n+1}}{2} + \frac{\xi_{n+2}}{2}\right), \quad (15)$$

which corresponds to a Leimkuhler–Matthews (LM) discretisation [21] of a Langevin diffusion targeting the Moreau–Yosida approximation $\pi^{\delta/2}$ of π given by (8), which is well-posed even when U is not smooth. Accordingly, the post-processed [39] iterates $X_n = Y_n - \sqrt{\delta/2}\xi_n$ converge to an approximation of π that

²Recall that U is defined in (4).

stems from correcting $\pi^{\delta/2}$ by performing a deconvolution of $\pi^{\delta/2}$ with the inverse of a Gaussian smoothing kernel of bandwidth $\delta/2$.

With regards to the convergence properties of the method, both the midpoint approximation (14) with $\theta = 1/2$ and the LM approximation (15) are known to be exact for Gaussian target distributions. That is, despite the discretisation of the OLSDE, the asymptotic bias vanishes completely as n increases [45, 21]. For more general log-concave distributions that are not Gaussian, the methods exhibit some bias, but are usually considerably more accurate than the conventional ULA schemes based on an Euler-Maruyama approximation of the OLSDE, such as MYULA [11].

Moreover, while a small bias is desirable, a primary concern in Bayesian computation for imaging problems is to achieve a fast convergence rate without incurring excessive bias. And while the dimension of the problem can affect the rate of convergence, the main challenge is to mitigate the impact of the high anisotropy of the target density π , which is inherent to imaging problems that are ill-conditioned or ill-posed. More precisely, as mentioned previously, the convergence rate of ULA strategies deteriorates as κ for target densities that are κ -strongly log-concave [11]. The accelerated algorithm SKROCK [31] empirically achieves a significantly better rate that only deteriorates as $\sqrt{\kappa}$, but this can only be verified theoretically in the case of a Gaussian target distribution.

Sections 3.1 and 3.2 below present a detailed non-asymptotic convergence analysis of (12) which shows that the method achieves acceleration for all smooth distributions tested that are strongly log-concave. The analysis also allows for identifying the optimal time step δ that maximizes the convergence speed of the method. It is worth mentioning at this point that our proposed method has similarities to the implicit schemes proposed in [15, 42], as they all coincide in the Gaussian case. We have chosen to use (12) as opposed to the schemes described in [15, 42] because (12) is well-posed for models that are not smooth, and because our analysis shows that (12) has better non-asymptotic convergence bounds.

It is also worth mentioning at this point that if U is L -smooth and m -strongly convex, then F is strongly convex with conditioning number $\frac{1+\theta L\delta}{1+\theta m\delta}$ meaning that from a computational point of view, for any δ , one step of the algorithm is at most as expensive as minimizing U directly to compute a MAP solution. In practice, a step of the algorithm is much cheaper because it is possible to use past iterations to warm-start computations. This, combined with an efficient solver, allows implementing (12) with a computational cost that is comparable to explicit accelerated methods from the state of the art such as SKROCK [31].

3.1 Analysis in the Gaussian case

We start our analysis for our proposed method from the representation (14) in the case where

$$\pi(x) \propto \exp\left(-\frac{1}{2}x^T \Sigma^{-1}x\right), \quad \Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2), \quad (16)$$

and study its convergence in the 2-Wasserstein distance, as a function of the parameter θ and the condition number $\kappa = \sigma_{\max}^2/\sigma_{\min}^2$. In this case, because $\log \pi$ is quadratic, we can solve directly for X_{n+1} in (14). As the covariance matrix is diagonal, each co-ordinate in (6) becomes independent from the rest. Note that the same analysis could be done for any covariance matrix Σ using singular value decomposition.

One step of IMLA for the i -th coordinate in the case of (16) is given by

$$X_{n+1}^i = R_1(z_i)X_n^i + \sqrt{2\delta}R_2(z_i)\xi_n^i, \quad \xi_n^i \sim \mathcal{N}(0, 1), \quad (17)$$

where $z_i = -\Delta t/\sigma_i^2$ and $X_0 = (X_0^1, \dots, X_0^d)$ is a deterministic initial condition, while

$$R_1(z) = \frac{1 + (1 - \theta)z}{1 - \theta z}, \quad R_2(z) = \frac{1}{1 - \theta z}. \quad (18)$$

Using the fact that Gaussian distributions are closed under linear transformations, and assuming that the initial condition X_0 is deterministic, we derive the distribution of X_n for any $\delta > 0$. In fact, we have the following proposition [31]

Proposition 3.1. Let $\pi(x) \propto \exp(-\frac{1}{2}x^T \Sigma^{-1}x)$ with $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$, and let Q_n be the probability measure associated with n iterations of the generic Markov kernel (17). Then the 2-Wasserstein distance between π and Q_n is given by

$$W_2(\pi; Q_n)^2 = \sum_{i=1}^d (D_n(z_i, x_0^i) + B_n(z_i, \sigma_i)) \quad (19)$$

where

$$D_n(z, x) = (R_1(z))^{2n} x^2, \quad B_n(z, \sigma) = \left[\sigma - \sqrt{2\delta} R_2(z) \left(\frac{1 - (R_1(z))^{2n}}{1 - (R_1(z))^2} \right)^{1/2} \right]^2.$$

In addition the following bound holds

$$W_2(\pi; Q_n) \leq W_2(\pi; \tilde{\pi}) + C^m W_2(\tilde{\pi}, Q_0) \quad (20)$$

where

$$\tilde{\pi} = \mathcal{N} \left(0, 2\delta (R_2(z))^2 \left[\frac{1}{1 - R_1^2(z)} \right] \right) \quad (21)$$

is the numerical invariant measure and

$$C = \sqrt{\max_{1 \leq i \leq d} R_1(z_i)^2}. \quad (22)$$

A closer look at the bound (20) reveals that the constant C controls how fast the numerical algorithm converges to its equilibrium behaviour and the term $W_2(\pi; \tilde{\pi})$ characterizes the bias of the numerical algorithm at equilibrium. The first step in analyzing the convergence of the algorithm is thus to understand the behaviour of the constant C for different values of the parameter δ . Surprisingly, this analysis will also be relevant in the general (non-Gaussian) log-concave case, since the convergence to equilibrium of this numerical scheme is controlled by the same constant. The analysis of the behaviour of the constant C as a function of δ for different values of θ is presented in detail in the Appendix, and can be used to determine the optimal number of steps n required such that $W_2(\pi, Q_n)^2 < \epsilon^2$ according to the following proposition:

Proposition 3.2. Let Q_n be the probability measure associated with the n -th iteration of the generic Markov kernel (17) starting at X_0 . Then the number of steps n required such that $W_2(\pi, Q_n)^2 < \epsilon^2$ is given by

$$n \approx \begin{cases} \frac{\sqrt{\kappa}}{2} [\log(W_2(\pi, Q_0)) - \log(\epsilon)], & \theta = \frac{1}{2} \\ \min \left(\frac{d\sigma_{\max}^2}{\epsilon^2}, \frac{\sqrt{d\kappa}\sigma_{\max}}{2\epsilon} \right) [\log(W_2(\pi, Q_0)) - \log(2^{-1}\epsilon)], & \theta = 1 \end{cases} \quad (23)$$

with δ given by

$$\delta = \begin{cases} \delta_*, & \theta = \frac{1}{2} \\ \max \left(\frac{\epsilon^2}{d}, \frac{2\epsilon\sigma_{\min}}{\sqrt{d}} \right) & \theta = 1. \end{cases}$$

where $\delta_* = \frac{2}{\sqrt{Lm}} = 2\sigma_{\min}\sigma_{\max}$.

Proof. The proof of this proposition can be found in the Appendix A.1. \square

In principle, we are able to repeat the analysis in Proposition 3.2 for all the different values of θ , but the two cases of interest are the ones where $\theta = 1/2$ or $\theta = 1$. In the case $\theta = 1/2$, the expression for n does not depend on d , while in the case of $\theta = 1$ it does depend on d . For $\theta = 1$, the dependence on ϵ is worse than logarithmic and the expression of n remains bounded as $\kappa \uparrow \infty$. These bounds show that $\theta = 1/2$ will outperform $\theta = 1$ unless $d\sigma_{\max}^2/(\sqrt{\kappa}\epsilon^2)$ is small.

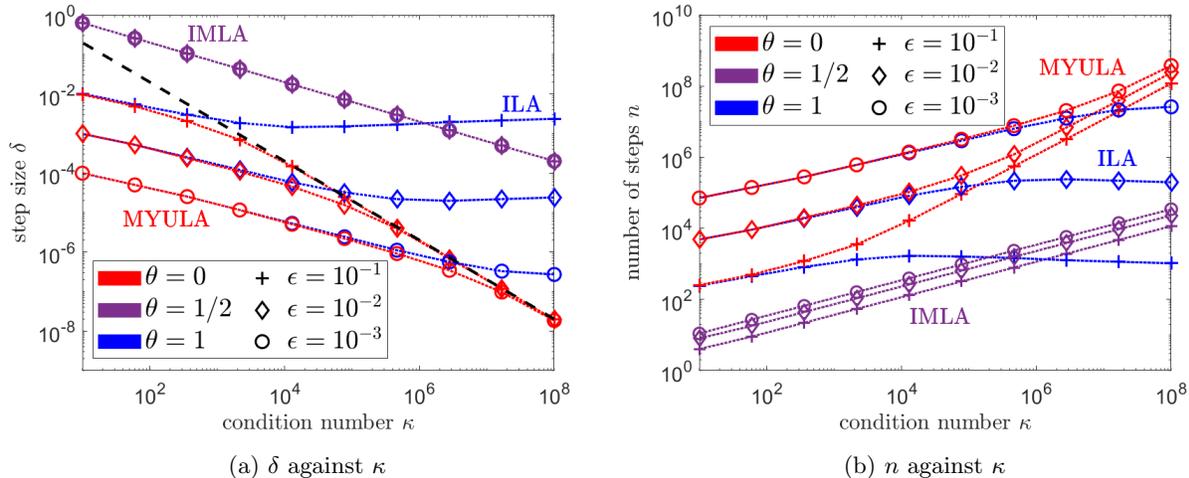


Figure 1: (a) step-size δ as a function of the condition number κ for different choices of θ (the black dotted line corresponds to $\kappa\delta = 2$, the stability limit of the explicit integrator) and level of accuracy ϵ . (b) Number of steps n such that $W_2(\pi, Q_n)^2 < \epsilon^2$ as a function of the condition number κ for different choices of θ and ϵ .

In order to visually illustrate the results of Proposition 3.2 we now proceed with the following numerical experiment: We choose $X_0 = 1/\sqrt{d}$ (so that $\|X_0\| = 1$) and σ_i geometrically spaced between 1 and $1/\sqrt{\kappa}$ to mimic an ill-conditioned posterior distribution. We also plot results for $\theta = 0$ to compare with an explicit method in addition to the cases $\theta = 1/2$ and $\theta = 1$. In the case of $\theta = 1/2$ since $W_2(\tilde{\pi}, \pi) = 0$ we have performed all our calculations for $\delta = \delta_*$, because this choice maximizes the convergence speed to equilibrium. For $\theta = 1$, we numerically determined the largest δ for which the bias $W_2(\pi, \tilde{\pi})$ (which has a closed-form expression) is $\leq \epsilon/2$. An explicit calculation is presented in the Appendix and shows the following dependence

$$\delta = \max\left(\frac{\epsilon^2}{d}, \frac{2\epsilon}{\sqrt{d\kappa}}\right). \quad (24)$$

Similarly, for $\theta = 0$, we numerically computed the largest δ such that $C < 1$ and $W_2(\pi, \tilde{\pi}) \leq \epsilon/2$, and subsequently determined the smallest n for which $W_2(\pi, Q_n)$ given by (19) is $\leq \epsilon$.

We summarise the results of these experiments³ in Figures 1(a) and 1(b). In Figure 1(a) we plot the step size δ as a function of the conditioning number κ . We observe that the behavior is as predicted by (24) in Proposition 3.2: For $\theta = 1$ and fixed ϵ , the step size δ decreases as $1/\sqrt{\kappa}$ initially, and eventually stops decreasing once δ becomes the order of $1/\kappa$ for increasing κ . The behaviour of δ as ϵ changes is as follows: On the left-hand side of the figure, dividing ϵ by 10 results in a division by 10 of δ , whereas on the right-hand side of the figure, dividing ϵ by 10 leads to a decrease in δ by a factor of 100. The figure shows that the value of δ for $\theta = 1$ is smaller than the optimal value for $\theta = 1/2$ except for $\epsilon = 0.1$ and very large κ . In the case where $\theta = 0$, ϵ is fixed and κ is varying there are two regimes: For small κ , the value of δ is chosen to ensure that the bias is sufficiently small, and the results are very similar⁴ to those corresponding to $\theta = 1$. However, for κ large, the value of δ is really determined by the stability limit.

Figure 1(b) gives n as a function of κ for the runs in Figure 1(a). We again observe that the behavior is as predicted by (24) in Proposition 3.2: For $\theta = 1/2$, increasing κ and fixed ϵ , n increases in the order of $\sqrt{\kappa}$ and changes insignificantly depending on ϵ . This is not the case for $\theta = 1$ or $\theta = 0$; in these cases, a reduction in ϵ implies a reduction in δ in order to decrease the bias, which in turn increases the contractivity constant $C \approx 1 - \delta$. As a consequence, in the case of $\theta = 1$, n initially increases like $\sqrt{\kappa}$ and then becomes

³We have used $d = 100$ and $\epsilon = 0.1, 0.01, 0.001$.

⁴For $\theta = 0$ the bias can be shown similarly to the proof of the proposition (24) to have a bound proportional to $\sqrt{d\kappa}\delta$ provided that δ is bounded away from the stability limit $2/\kappa$.

constant as predicted. Only when $\epsilon = 0.1$ and κ is very large, $\theta = 1$ improves on $\theta = 1/2$. In the case $\theta = 0$ and κ large, the number of steps required behaves as $1/\kappa$ and is up to four orders of magnitude larger than the number of steps for $\theta = 1/2$. Therefore, for large κ , the implicit integrator with $\theta = 1/2$ is to be preferred to the explicit algorithm even if its computational complexity per step is substantially higher.

3.2 Analysis in the strongly log-concave case

We extend the Gaussian non-asymptotic analysis to the strongly log-concave case. The following theorem is the analogue of Proposition 3.1.

Theorem 3.3. *Let $U = -\log \pi$ and suppose that $U \in C^2$, m -strongly convex and has gradient which is L -Lipschitz. For any probability distribution, Q_0 , let Q_n denote the probability distribution of X_n where $X_0 \sim Q_0$ and X_k is given according to the scheme (13). We assume that each step of the scheme (13) is solved to a tolerance of ϵ , i.e. that $\|\nabla F(X_{n+1}; X_n, \xi_n)\| \leq \epsilon$ for every n . Then for any initial probability distribution Q_0 with finite second moments we have*

$$W_2(Q_n, \pi) \leq C^n W_2(Q_0, \pi) + \frac{1 - C^{n+1}}{1 - C} \frac{\frac{1}{2}\delta^2 L^{\frac{3}{2}} \sqrt{d} + \frac{2}{3} L \delta^{\frac{3}{2}} \sqrt{2d} + \epsilon \delta}{1 + \theta \delta m} \quad (25)$$

and

$$C = \sqrt{\max_{z \in [m\delta, L\delta]} R_1(-z)^2}. \quad (26)$$

Proof. The proof is an adaptation of [15, Theorem 2] for the numerical scheme (13) and is deferred to the Appendix A.2. \square

The bound we have derived in (25) has two terms, the numerical contraction rate and the bias. The rate of contraction C we obtain in Theorem 3.3 is sharp. The constant C here is the same as the contraction rate for the Gaussian problem. This is a significant improvement upon the conclusions found in [15]. Indeed, [15, Theorem 2] derives a contraction rate of $\kappa_\delta C^n$ for the numerical scheme

$$X_{k+1} = X_k - \theta \delta \nabla U(X_{k+1}) + (1 - \theta) \nabla U(X_k) + \sqrt{2\delta} \xi_k. \quad (27)$$

Here, κ_δ is given by $\kappa_\delta = \frac{2+\theta\delta L}{2+\theta\delta m}$. On the other hand, the bias term in (25) is not sharp and is consistent with the bias found in [15, Theorem 2]. As in the Gaussian case, the following proposition uses the bound in (25) to determine the optimal number of steps to achieve accuracy ϵ .

Proposition 3.4. *Let Q_n be the probability measure associated with the n -th iteration of the Markov kernel (14) with $\theta = 1/2$ starting at X_0 . Then for any $\epsilon > 0$ there exists $\delta > 0$ such that for*

$$n \approx \max \left\{ \frac{\sqrt{\kappa}}{4}, \frac{2\kappa\sqrt{\kappa d}}{2\sqrt{m}\epsilon}, \frac{64}{9} \frac{d\kappa}{m\epsilon^2} \right\} [\log(\epsilon/2) - \log(W_2(\pi, Q_0))]$$

we have $W_2(Q_n, \pi) \leq \epsilon$.

As noted above, the bias in (25) is not sharp and as a result the description of the small ϵ regime is suboptimal. However, if ϵ is sufficiently large then the number of steps grows with $\sqrt{\kappa}$ and hence in this regime the result is consistent with the Gaussian setting.

3.3 Illustrative experiment: a toy denoising model based on a Gaussian mixture prior

We now consider a Gaussian mixture model for denoising an image x from observations $y = x + \xi$ with $\xi \sim \mathcal{N}(0, \sigma^2 I_d)$ with noise variance σ^2 . This choice allows us to demonstrate our approach with a posterior

| | | | | | |
|-------|-------|--------------|--------------|------------|------------------|
| m_0 | m_1 | σ_0^2 | σ_1^2 | σ^2 | $\tilde{\omega}$ |
| 0 | 0 | 0.0025 | 0.0809 | 0.0016 | 0.9 |

Table 1: Parameters for the Gaussian mixture model for denoising.

| Algorithm | Mean $\mathcal{W}_2(\pi, \tilde{\pi})$ | Std. Dev. |
|-----------|---|------------|
| EXACT | 1.4123e-07 | 1.4448e-07 |
| IMLA | 1.4095e-07 | 1.5540e-07 |
| ILA | 6.2550e-04 | 1.7596e-06 |
| ULA | 6.6183e-05 | 7.7698e-07 |
| SKROCK | 2.0804e-04 | 1.3499e-06 |

Table 2: $\mathcal{W}_2(\pi, \tilde{\pi})$ errors for one arbitrary pixel.

distribution that deviates from a Gaussian density but for which we can compute quantities exactly. More precisely, the prior is the product of d 1-dimensional identical Gaussian mixtures of two Gaussian random variables \tilde{Z}_k with means $m_k \in \mathbb{R}$, variances $\sigma_k^2 \in \mathbb{R}$ for $k \in \{0, 1\}$ and weights $\tilde{\omega}$ and $1 - \tilde{\omega}$ for some $\tilde{\omega} \in [0, 1]$. It is possible to show that the posterior distribution is now a product of d 1-dimensional Gaussian mixtures of two independent Gaussian random variables Z_k with means $\mu_k(y_i)$, variances δ_k^2 and weights $w(y_i)$. Note that the means and the weights of each of the d Gaussian mixtures are now depending on the pixel values y_i (the precise relationship can be found in the Appendix A.4).

In Section 3.1 we showed that IMLA samples are exact for Gaussian distributions. Here, we investigate the bias for a non-Gaussian problem. Note that the posterior distribution associated with this mixture prior is not usually strongly log-concave. However, for the range of value of pixels y_i that we have used in our experiments, as well as our choices of prior parameters $(m_0, m_1, \sigma_0^2, \sigma_1^2, \tilde{\omega})$ and noise variance σ^2 that can be seen in Table 1, the corresponding posterior distribution is indeed strongly log-concave and fits the assumptions of our theory.

Furthermore, in our setting the posterior distribution is separable into one-dimensional distributions over the pixel domain, which significantly simplifies the analysis of the estimation error. In particular, for any pixel it is possible to reliably compute the Wasserstein-2 error w.r.t. the true pixel’s distribution by using [33, Remark 2.30]. We conducted this toy experiment by using a 60×60 pixel region of the `lizard` test image from the BSDS300 data set [23].

We quantify the bias of each method in the following two complementary ways: (i) by computing the Wasserstein-2 error for the marginal distribution of a randomly selected image pixel; and (ii) by computing the distribution of the scalar summary statistic $\log \pi(x)$, which defines the typical set of the posterior distribution and plays a central role in uncertainty quantification analyses [28]. The Wasserstein-2 error $\mathcal{W}_2(\pi, \tilde{\pi})$ was computed using 50 repetitions of 10^6 realisations for each of the algorithms under scrutiny; for completeness, we also included a fully implicit ($\theta = 1$) Langevin algorithm (ILA) in the comparison. Table 2 reports the mean and the standard deviation of the Wasserstein-2 error for each method w.r.t. the truth. To provide an indication of the baseline error stemming from the finite Monte Carlo sample size, we also include the results obtained for a perfect sampler that draws exact samples from π . We observe that the error of IMLA is almost exactly the same as the baseline error of the exact samples, confirming the previous analysis and the fact that this algorithm has a remarkably low bias. This result is confirmed in Table 3, showing the summed Wasserstein errors over all pixels and IMLA having the least error compared to all other methods.

Moreover, Figure 2 depicts the histogram of the scalar statistic $T(x) = \log \pi(x)$, as computed from 5×10^5 iterations of each algorithm. For IMLA ($\theta = 1/2$), SKROCK, and ILA ($\theta = 1$) we have used the optimal time step δ_* associated with IMLA, while for ULA we used $\delta = 1/L$ as recommended in [11]. We observe that the histogram obtained by using IMLA is in close agreement with the truth, while all the other methods

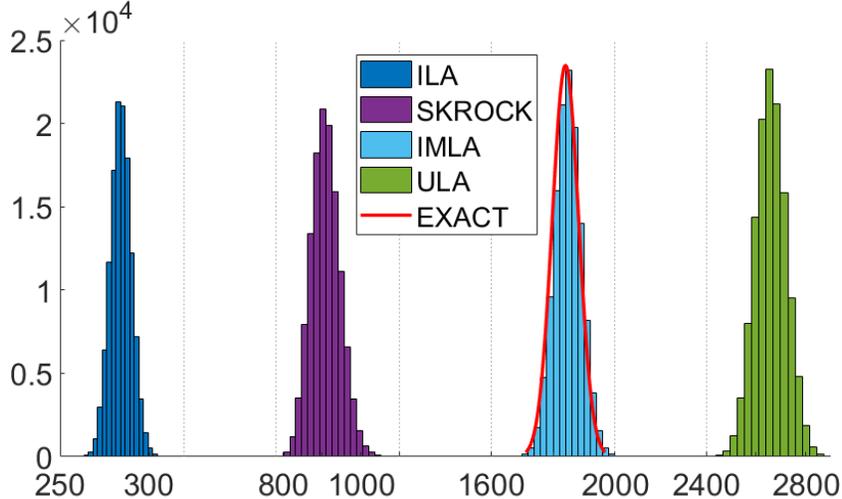


Figure 2: Comparison of histograms of the scalar statistic $\log \pi(x)$ for the Gaussian mixture model. This figure shows the different biases induced by the sampling algorithms, and significant agreement between the exact and the IMLA statistic.

| Algorithm | $\mathcal{W}_2(\pi, \tilde{\pi})$ |
|-----------|-----------------------------------|
| EXACT | 0.2490 |
| IMLA | 0.3991 |
| ILA | 3.3662 |
| ULA | 1.1514 |
| SKROCK | 2.0769 |

Table 3: Summed $\mathcal{W}_2(\pi, \tilde{\pi})$ errors over all pixels computed over 15000 samples.

exhibit some significant bias, with ILA exhibiting the largest amount of bias, and SKROCK and ULA having similar moderate biases.

3.4 Illustrative experiment: One dimensional distributions

We continue our illustrative experiments by comparing the algorithms MYULA ($\theta = 0$), IMLA ($\theta = 1/2$), and ILA ($\theta = 1$) used to sample from three representative one-dimensional distributions, namely

- Laplace distribution: $\pi(x) \propto e^{-|x|}$,
- Uniform distribution: $\pi(x) = e^{-\iota_{[0,1]}(x)}$,
- Light-tailed distribution: $\pi(x) \propto e^{-x^4}$.

We use $\iota_{[0,1]}$ to denote the indicator function which takes the value 0 on $[0, 1]$ and $+\infty$ otherwise (such that $\pi(x) = 1$ in $[0, 1]$ and 0 outside that interval). We have chosen these distributions because they exhibit different tail behaviors and are representative situations where ULA can not be applied directly without smoothing the distribution (the distributions are not Lipschitz-differentiable). In contrast, using the formulation (12) of IMLA we can apply IMLA directly for these distributions without smoothing. Moreover, since the proximal operator has an analytic expression in each of these cases, the implementation of IMLA and ILA is straightforward and has the same cost as MYULA.

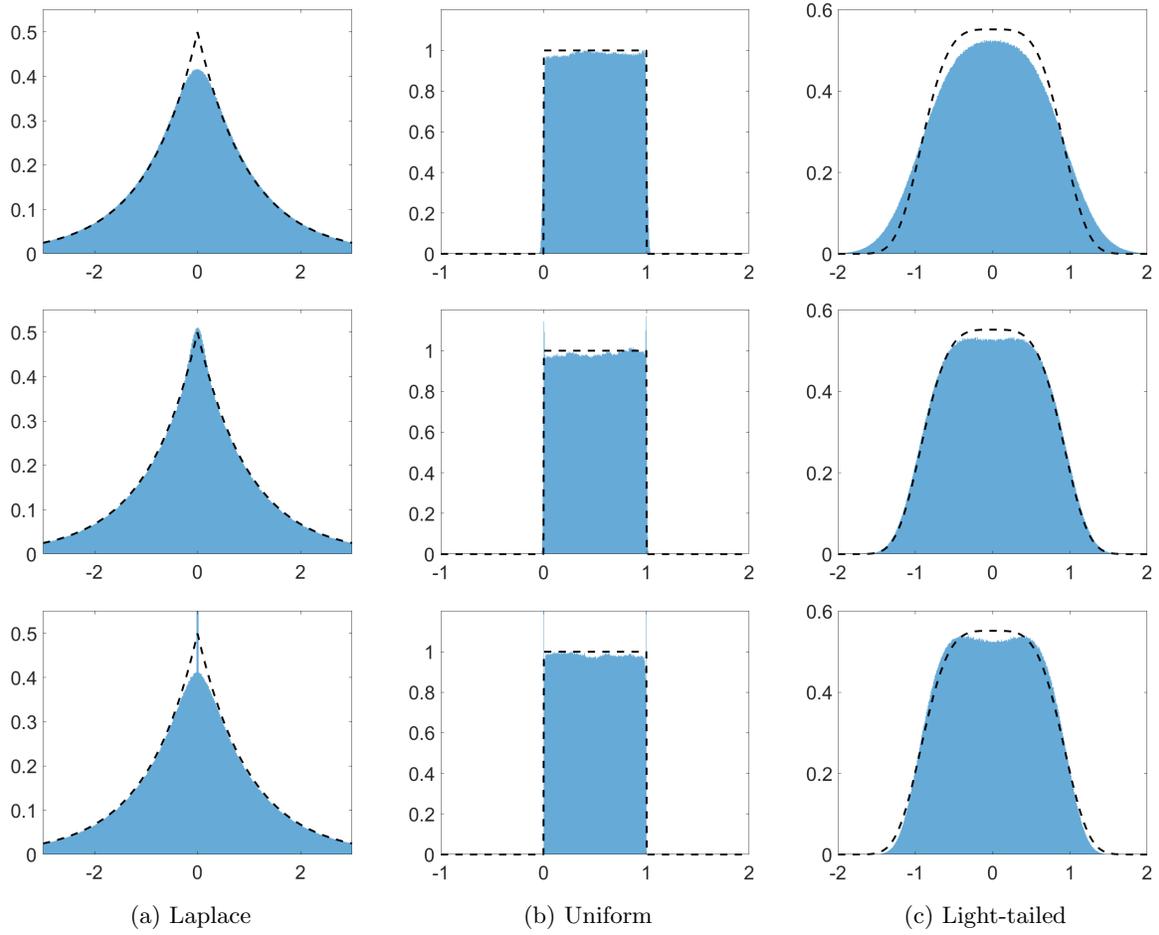


Figure 3: Histograms of the Laplace distribution (left), the uniform distribution (middle) and a light-tailed distribution with $\pi(x) \propto e^{-x^4}$ (right) computed with 15×10^6 iterations generated by MYULA (top) and IMLA (middle) and ILA with $\theta = 1$ (bottom), for $\delta = 0.05$ (left, right) and $\delta = 10^{-4}$ (middle). The dashed black curve represents the probability density function of the respective distribution.

| Distribution | SD IMLA | SD ILA | SD MYULA | SD EXACT |
|--------------|---------|--------|----------|----------|
| Laplace | 1.4046 | 1.4005 | 1.4356 | 1.4142 |
| Uniform | 0.2923 | 0.2936 | 0.2949 | 0.2887 |
| $\exp(-x^4)$ | 0.5964 | 0.5777 | 0.6590 | 0.5813 |

Table 4: Estimated vs. the exact standard deviations (SD) for each method (IMLA, ILA and MYULA) for the experiment shown in Figure 3.

Algorithm 1 IMLA

Require: $N \geq 0$, $\delta > 0$ and $X_0 \in \mathbb{R}^d$.

for $n=0 : N-1$ **do**

Draw

$$\xi_n \sim \mathcal{N}(0, I_d)$$

Set

$$X_{n+1} \leftarrow \arg \min_{x \in \mathbb{R}^d} 2U \left(\frac{1}{2}x + \frac{1}{2}X_n \right) + \frac{1}{2\delta} \|x - X_n - \sqrt{2\delta}\xi_n\|^2$$

end for

Figure 3 depicts the histograms associated with the Markov chains generated by using MYULA, IMLA, and ILA for the three test distributions considered. In the case of the Laplace and the light-tailed distributions, we have used $\delta = 5 \times 10^{-2}$, while for the uniform distribution, we have used $\delta = 10^{-4}$. We have run the algorithms for 15×10^6 iterations and set the smoothing parameter of MYULA to $\lambda = \delta$, as we found that this gives the best approximation of π in terms of the convergence speed and asymptotic bias trade-off.

As can be observed in Figure 3, IMLA and ILA achieve a visibly smaller bias than MYULA for all the distributions, which is in agreement with our analysis in Section 3.1. For a quantitative comparison, we report the standard deviation (SD) estimated by each method in Table 4. It can be observed that MYULA overestimates the SD in all cases, IMLA and ILA are significantly closer to the true value, and IMLA is the most accurate for the Laplace and the Uniform distributions.

3.5 Algorithmic recommendations: IMLA

The analysis in Sections 3.1 indicates that to converge as quickly as possible, one should choose $\theta = 1$ with arbitrary large step size δ . However, the bias in this case would be very large as seen in Figures 1(a) and 1(b). In contrast, using $\theta = 1/2$ in the Gaussian case is bias-free. Moreover, it can be observed in Figure 1(b) that it also provides a better non-asymptotic convergence behavior than $\theta = 1$. The illustrative numerical experiments in Section 3.3 reveal that this behavior persists outside the Gaussian regime. Furthermore, the numerical experiments in Section 3.4, even though not covered by our theoretical analysis, reveal no substantial benefit of using $\theta = 1$ over $\theta = 1/2$ for the values of δ considered. Hence, from now on we concentrate on the choice $\theta = 1/2$, which we summarise in Algorithm 1. In the case when U is strongly convex, we set $\delta = \delta_* = 2/\sqrt{Lm}$ which gives the optimal contraction rate. Algorithm 1 does not incorporate hard constraints, but it can be modified if needed, as will be discussed in Section 4.2.

4 Imaging experiments

In this section, we demonstrate the proposed Bayesian methodology on a range of experiments related to image deconvolution. We consider problems with an underlying convex geometry (i.e., with a posterior distribution that is log-concave). In Section 4.1 we consider a model with a data-driven convex prior and a Gaussian noise characteristic, and in Section 4.2 we consider a model with an assumption-driven convex prior and a Poisson noise characteristic.

| Experiment | L | m | h_{IMLA} | stages | h_{ULA} | α | μ |
|------------|-------|-------|-------------------|--------|------------------|----------|-------|
| castle | 43521 | 17.39 | 0.002 | 8 | 2.3e-05 | 2828.43 | 8.00 |
| lizard | 50885 | 17.25 | 0.002 | 8 | 2.0e-05 | 3363.59 | 8.00 |
| person | 41896 | 2.41 | 0.006 | 12 | 2.4e-05 | 2828.43 | 11.31 |

Table 5: Algorithm parameters for motion deconvolution experiment.

4.1 Image deconvolution using a CRR-NN prior

We consider a non-blind motion deblurring problem where we seek to recover $x \in \mathbb{R}^d$ from a blurry and noisy observation $y = Ax + \epsilon$, where the linear operator A represents a known motion blur kernel and $\epsilon \sim \mathcal{N}(0, \sigma^2 I_d)$. We conduct our experiments with a range of motion blurs from [22]; in all cases, $A^\top A$ is highly ill-conditioned but full rank. To regularise the estimation problem, we leverage a state-of-the-art data-driven prior [13], which is by construction log-concave. Combining this prior with the likelihood gives us a posterior distribution with the following density

$$\pi(x) \propto \exp\left(-\frac{\|Ax - y\|^2}{2\sigma^2} - \frac{\alpha}{\mu} R_{\Theta}(\mu x)\right), \quad (28)$$

where $R_{\Theta}(\cdot)$ is the convex-ridge regularizer neural network (CRR-NN) [13], with $\mu > 0$ and $\alpha > 0$ scale parameters. In our numerical experiments, we use three different images from the BSD300 data set [23] (**castle**, **lizard**, and **person**) depicted in Figure 4. To degrade each image, we use a different blur kernel⁵ and generate the observation y by adding Gaussian noise with variance σ^2 chosen to achieve a blurred signal-to-noise ratio of 30dB. The corresponding blurred and noisy images are shown in Figure 4, while the values of the Lipschitz and strong convexity constants (L, m) for each blur kernel are reported in Table 5. In addition, as recommended in [13], the parameters α and μ of the neural network regularizer are adjusted for each experiment by using a grid search method⁶, which optimizes the parameters to yield the MAP solution with optimal peak-signal-to-noise ratio (PSNR).

For each image y , we use IMLA, SKROCK and ULA to draw Monte Carlo samples from the posterior distribution π . In the case of ULA, the step size is set to half the stability barrier, $h_{\text{ULA}} = 1/L$ as recommended in [11], which is of the order of 10^{-5} . In contrast, for IMLA we use the optimal step size $h_{\text{IMLA}} = 2/\sqrt{Lm}$ as suggested by the analysis in Section 3.1 (note that h_{IMLA} is two orders of magnitude larger than h_{ULA} , see Table 5). We use the same step size for SKROCK, and automatically derive the appropriate number of internal stages, s , for each experiment (see Table 5). To make the comparison fair between SKROCK and ULA we use 4×10^4 iterations for SKROCK and $4s \times 10^4$ iterations for ULA. For IMLA, we use the same number of iterations as for SKROCK (4×10^4), as we are initially interested in comparing the performance of the two methods without taking into account the average computational cost per iteration (since this is highly implementation dependent). We defer the analysis of the computational cost to Section 4.2.2.

To demonstrate the effectiveness of our method in an uncertainty quantification context, we summarise π by computing two different Bayesian estimators that summarise π optimally w.r.t. to complementary loss functions and by computing uncertainty visualisation plots based on (marginal) second-order moments of π . More precisely, we compute: (i) the minimum mean square estimator (MMSE) solution given by the posterior mean, (ii) the MAP solution by solving (28) as an optimisation problem⁷ (see [29] for details about the optimality of these estimators), (iii) the marginal standard deviations of x at the pixel-wise scale, and (iv) the joint marginal standard deviations for groups of pixels of size 2×2 , 4×4 , 8×8 , and 16×16 pixels, which display the uncertainty about structures of different size and in different regions of x . The MMSE and MAP solutions are depicted in Figure 5, whereas the uncertainty visualisation plots are displayed in Figure 6 and Figure 7. For comparison, we also include the results obtained with a Bayesian model using a non-convex neural network prior [37] used within the PnP-ULA [20] method to perform computations.

⁵We use the kernels from [22] with indexes {2, 3, 4}.

⁶https://github.com/axgoujon/convex_ridge_regularizers

⁷The MAP solution is computed by using the gradient method [13].



Figure 4: Motion deconvolution experiment: Ground truth images x (left) of size 320×320 with observations y (right) for castle, person and lizard images (top to bottom) and their respective blur kernels in the top left corner of y .

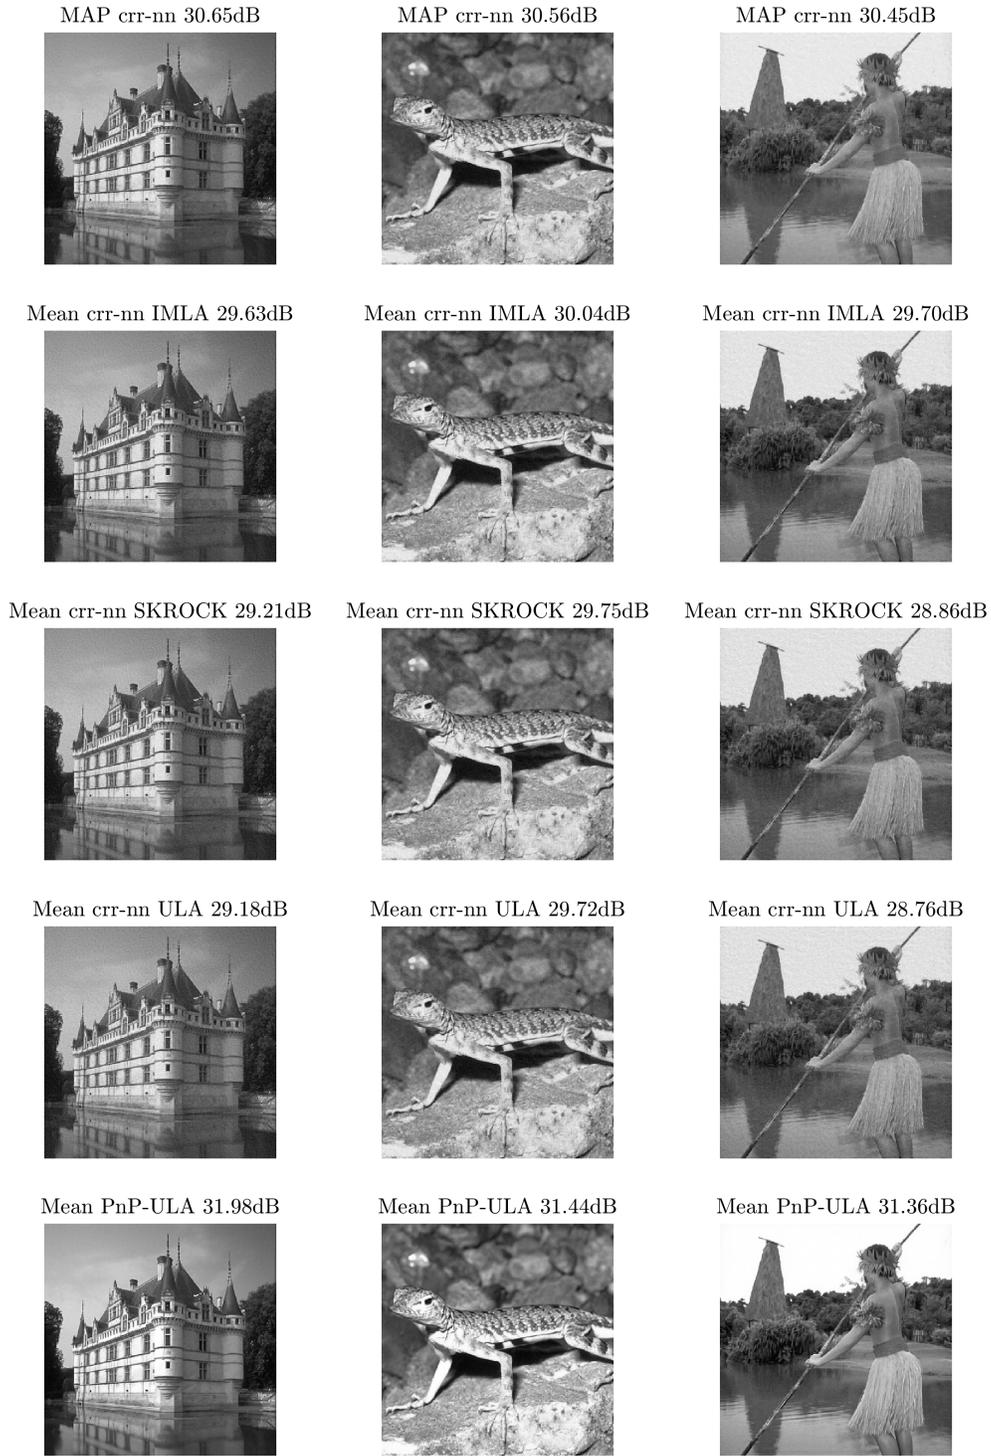


Figure 5: MAP solution for the motion deconvolution problem using a gradient method [13] for the CRR-NN model (first line) and posterior means for IMLA, SKROCK, ULA and PnP-ULA (line 2-5) for **castle**, **lizard** and **person** images (left to right) including the corresponding PSNR values in *dB*.

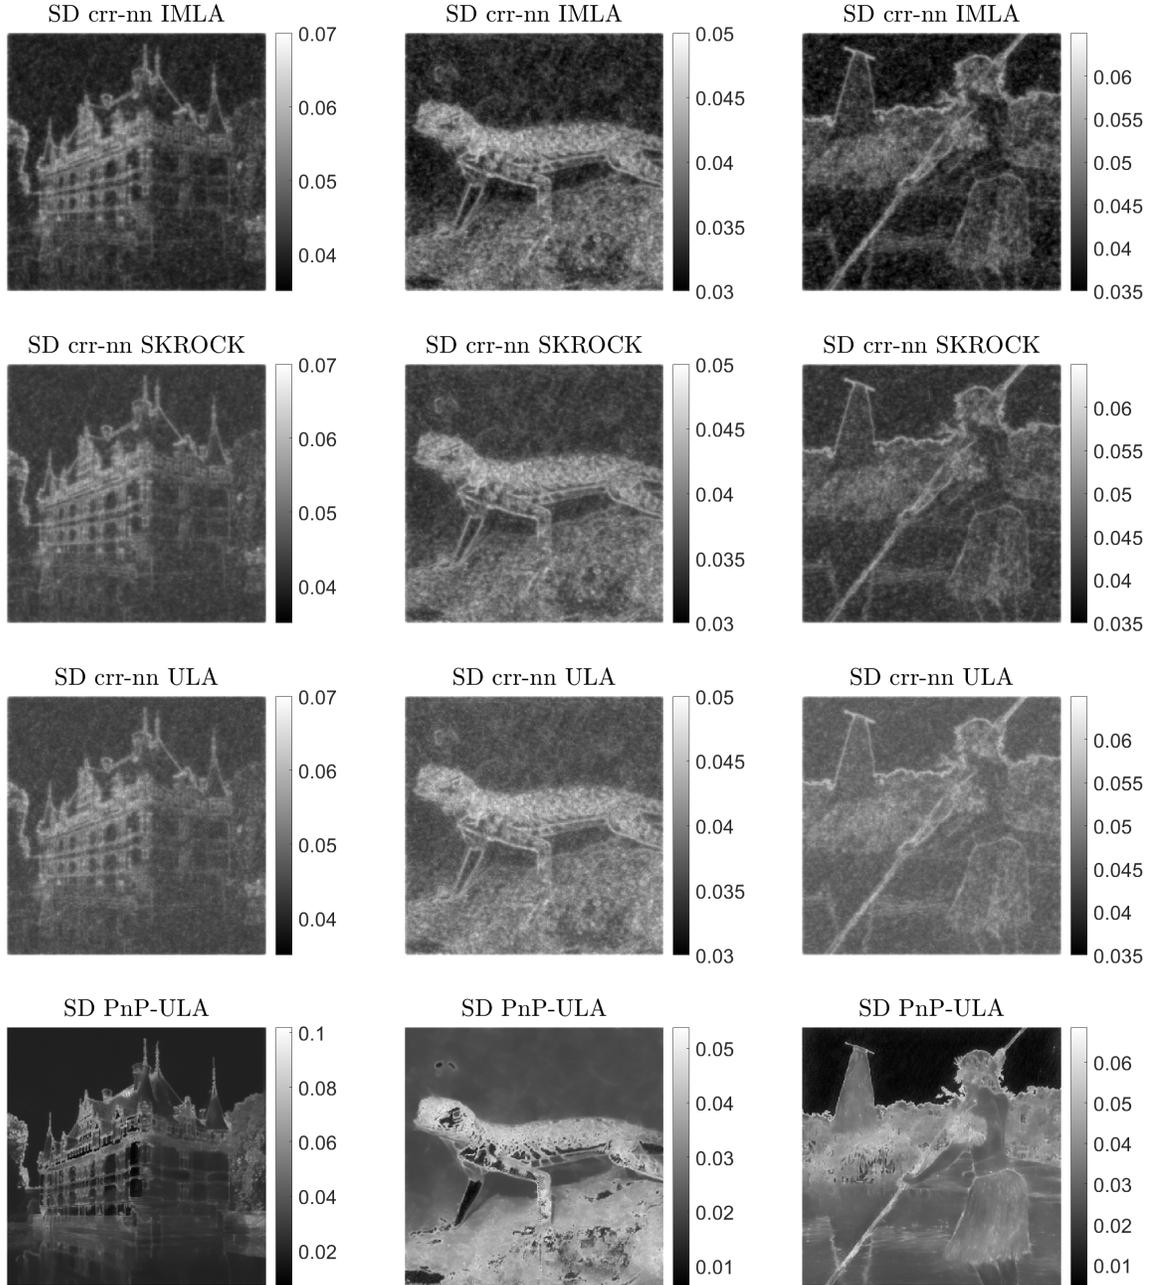


Figure 6: Marginal posterior standard deviation (SD) for the motion deconvolution problem for IMLA, SKROCK, ULA and PnP-ULA for `castle`, `lizard` and `person` images (left to right).

From Figures 5-7, we observe that IMLA achieves a higher PSNR than ULA and SKROCK (see Figure 5) and that it also delivers more stable estimates of the second-order moments. In particular, ULA and SKROCK have noticeably higher Monte Carlo error noise in the background of Figure 6, which indicates that IMLA is exploring the solution space more quickly. The MAP estimator delivers better point estimates in terms of PSNR, which is common in models that are log-concave, especially when the hyperparameters

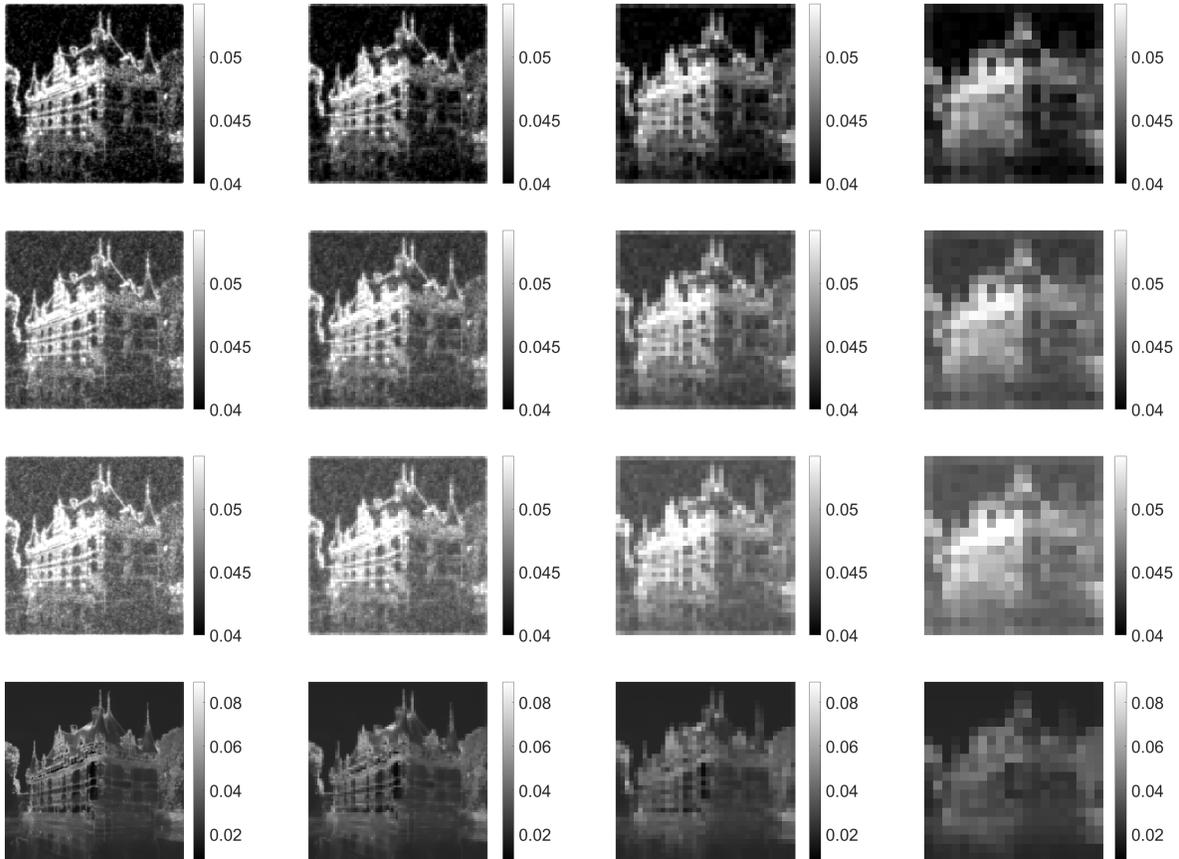


Figure 7: Visualisation of uncertainty by displaying the marginal standard deviations at scales 2×2 , 4×4 , 8×8 , and 16×16 pixels (left to right), for IMLA, SKROCK, ULA, and PnP-ULA (top to bottom) for the `castle` image.

of (28) are tuned to optimise the PSNR of the MAP solution. In addition, the MMSE calculated by PnP-ULA outperforms the MAP and MMSE solutions provided by the log-concave model (28) due to the higher accuracy of the non-convex prior. This additional accuracy comes at a potentially significantly higher computational cost, as navigating non-convex landscapes requires additional iterations and evaluating the neural network [37] is significantly more expensive than evaluating [13].

In order to further illustrate the convergence properties of the IMLA, ULA and SKROCK algorithms, Figure 8 depicts the PSNR of the running mean, the trace of the statistic $\log \pi$, and the autocorrelation function (ACF) of the slowest and fastest mixing components of the Markov chains for the `castle` image (the results for the other two images were quantitatively equivalent). Note that the slow/fast components have been identified by performing an approximate singular value decomposition of the posterior covariance and then computing the ACF for the projections of the chain on the eigenvectors with largest (fastest) and smallest (slowest) eigenvalues.

In this figure, one can observe that IMLA and SKROCK behave very similarly both in terms of their transient regimes (evolution of PSNR and of $\log \pi$), and in terms of their stationary regimes (described by the ACFs), whereas ULA converges quickly to the typical set of π (as illustrated by the evolution of $\log \pi$) but explores the space more slowly (see the ACF plots for ULA). This is consistent with the theory for ULA and IMLA, and in agreement with our understanding of SKROCK based on its analysis for Gaussian models.

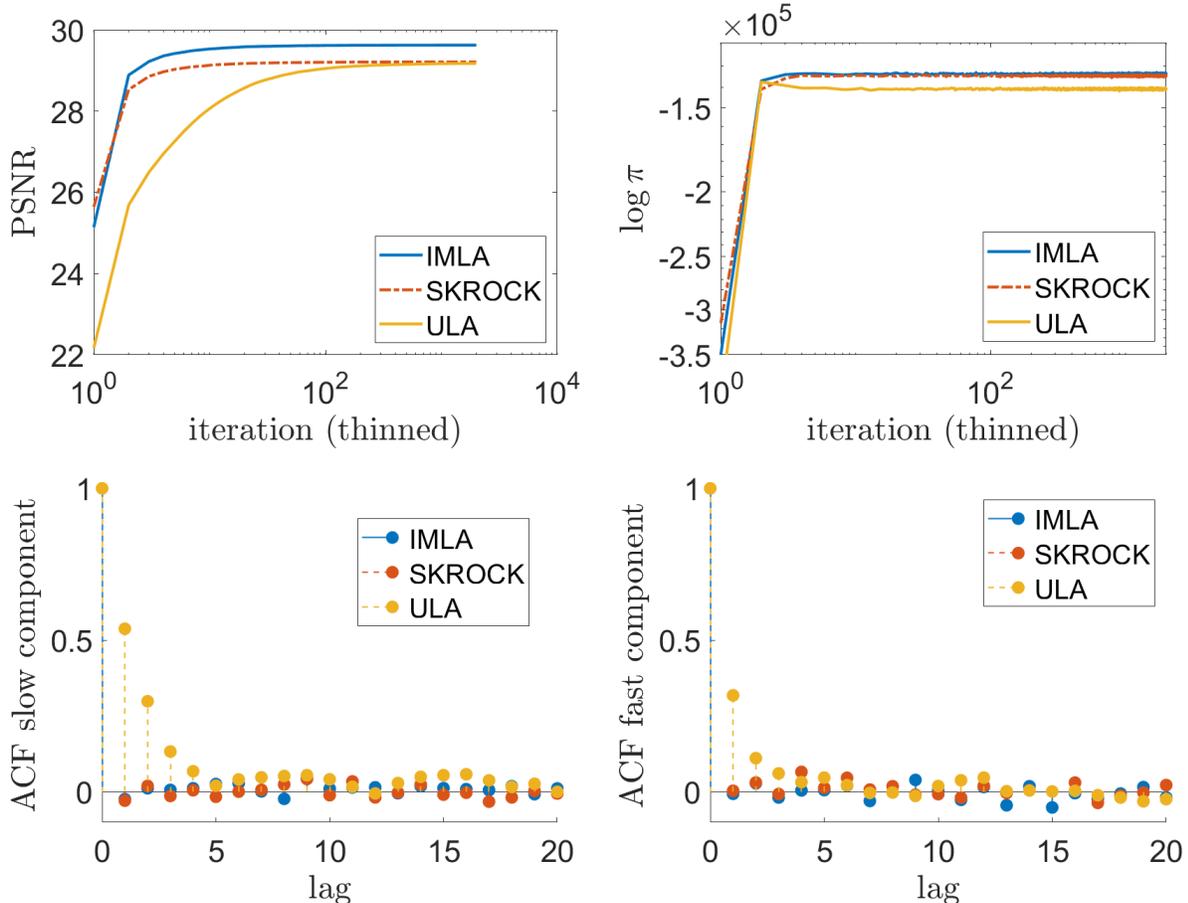


Figure 8: Motion deconvolution experiment. PSNR of the running mean (top left) and $\log \pi(x)$ (top right) for N iterations, thinned fairly to obtain 2000 data points, and ACF of the slowest and fastest component (bottom left and right) for IMLA, SKROCK and ULA for the `castle` image.

4.2 Deconvolution with Poisson noise

We now consider a total-variation (non-blind) image deconvolution problem with Poisson noise [24], which allows us to demonstrate our methodology in a strongly non-Gaussian and non-smooth situation. The posterior distribution that we consider is given by

$$\pi(x) \propto \exp \left(- \sum_{i=1}^d [(Ax)_i + \beta - y_i \log((Ax)_i + \beta) + \iota_{(Ax)_i \geq 0}] - \theta_{\text{TV}} \text{TV}(x) \right), \quad (29)$$

where $\beta > 0$ is an arbitrary constant that can be interpreted as a known background level.

This model has two forms of non-smoothness that represent a challenge to the application of strategies derived from the OLSDE (6): (i) the model contains a (hard) positivity constraint; (ii) the total-variation pseudo-norm $\text{TV}(\cdot)$ is not Lipschitz smooth. Also note that (29) is weakly log-concave, hence this allows us to illustrate our approach in a situation where the analysis of Sections 3.1 and 3.2 does not apply directly. In particular, because of the Poisson likelihood and the positivity constraint, π is very different from a Gaussian model, so one could expect IMLA to exhibit a noticeable bias when the step size is sufficiently large.

Following [11], to address the lack of smoothness of the total-variation prior we replace $\text{TV}(\cdot)$ by its



Figure 9: cameraman image x (left) and observed image y (right) of size 256×256 for the Poisson deconvolution experiment.

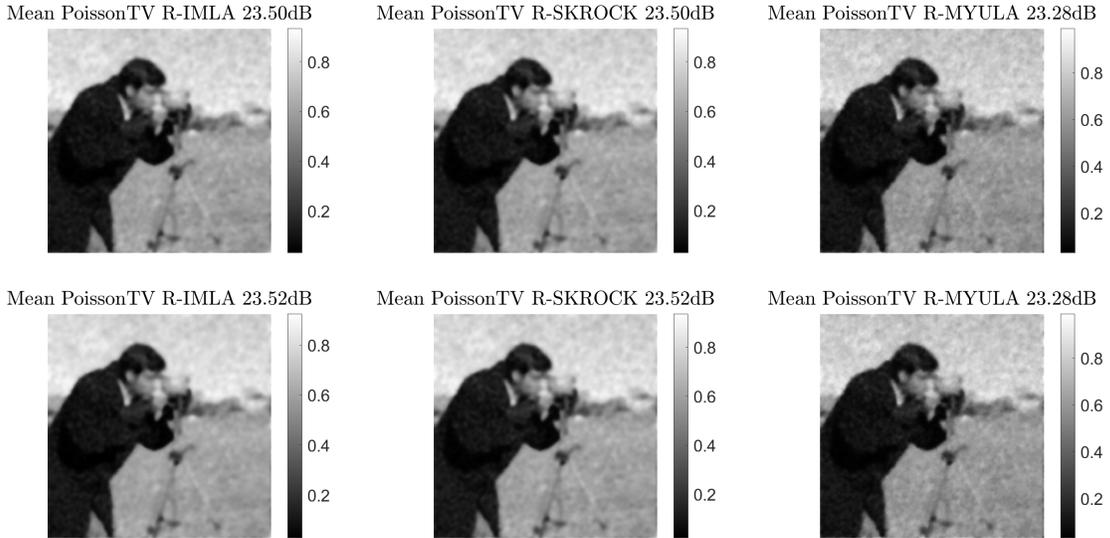


Figure 10: Posterior mean for Poisson deconvolution experiment. R-IMLA was run using $h = 6.65 \times 10^{-5}$ (top), and $h = 1.16 \times 10^{-3}$ (bottom) which is equivalent to the effective step size of R-SKROCK when $s = 10$ (top) or $s = 40$ (bottom); R-MYULA was run using $h = 1/L$.

Moreau-Yosida envelope

$$\text{TV}^\lambda(x) = \min_{u \in \mathbb{R}^d} \left\{ \text{TV}(u) + \frac{1}{2\lambda} \|x - u\|^2 \right\},$$

and construct the regularised posterior density

$$\pi^\lambda(x) \propto \exp \left(- \sum_{i=1}^d [(Ax)_i + \beta - y_i \log((Ax)_i + \beta) + \iota_{(Ax)_i \geq 0}] - \theta_{\text{TV}^\lambda} \text{TV}^\lambda(x) \right).$$

Applying gradient-based sampling strategies to π^λ remains challenging due to the positivity constraint. Following [24], we address this issue by using a reflected version of the overdamped Langevin diffusion, which

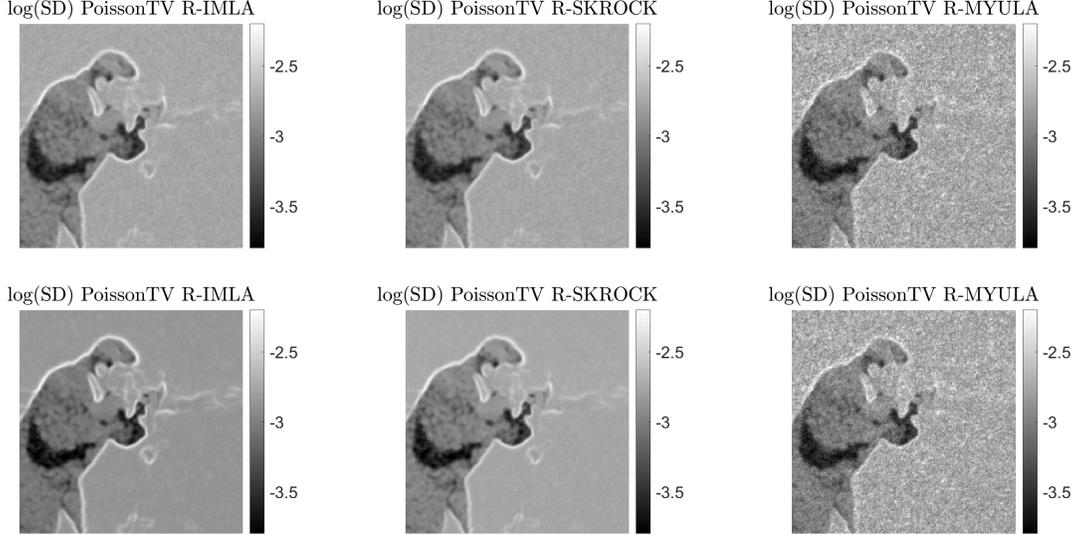


Figure 11: Marginal posterior standard deviation for the Poisson deconvolution experiment, in log-scale. R-IMLA was run using $h = 6.65 \times 10^{-5}$ (top), and $h = 1.16 \times 10^{-3}$ (bottom) which is equivalent to the effective step size of R-SKROCK when $s = 10$ (top) or $s = 40$ (bottom); R-MYULA was run using $h = 1/L$.

naturally incorporates the positivity constraint. This diffusion underpins the R-MYULA and R-SKROCK methods recently introduced in [24]. Analogously, we construct a reflected version of IMLA (R-IMLA) to approximate a reflected Langevin diffusion targeting π^λ . The resulting scheme is summarised in Algorithm 2⁸.

For this experiment, A is a 5×5 box blur operator, while we use the `cameraman` image scaled to have a mean intensity value (MIV) of 10 (ground truth and the observed data can be observed in Figure 9). Similarly to [24], the background noise level β is set to 1% of the MIV. Using the guidelines from [24, Section 4.4] we set $\lambda = L_{f_y}^{-1}$ where L_{f_y} is the Lipschitz constant of $x \mapsto \nabla \log p(y|x)$, and set the θ_{TV^λ} automatically by marginal likelihood estimation by using the SAPG algorithm [38].

We apply the R-IMLA, R-SKROCK, and R-MYULA algorithms to generate Monte Carlo samples from π^λ . As in the previous experiment, we set the time-step for R-MYULA to half the stability barrier $h_{R-MYULA} = 1/L$, with $L = L_{f_y} + 1/\lambda$. For R-SKROCK, [24] recommends setting $s \in \{5, \dots, 15\}$ and conducts experiments with $s = 10$, so we also set $s = 10$. For completeness, we also report results with $s = 20$ and $s = 40$ in order to assess the effect of the step size on the convergence speed and the estimation bias - the corresponding step size values can be seen in Table 6. Again to make the comparison fair between R-MYULA and R-SKROCK we use 8×10^5 iterations for R-MYULA and $8 \times 10^5/s$ iterations for R-SKROCK.

With regard to IMLA, unlike most examples seen before, the posterior π^λ is not strongly log-concave and hence we do not have an optimal choice for the time-step. Thus, we consider comparisons with R-IMLA in two different ways. In the first comparison, we choose to ignore the computational cost of solving the optimisation problem (30) at each iteration and use the same time-step and the same number of iterations for IMLA and SKROCK. These comparisons are reported in Section 4.2.1. In contrast, in the second comparison, we keep the time steps equal but take into account the cost per iteration of each algorithm by modifying the number of iterations for IMLA so that IMLA and SKROCK have the same clock time. These experiments are reported in Section 4.2.2.

⁸One can also apply IMLA without replacing g by g^λ , but this would make the comparisons with MYULA and SKROCK less clear as the algorithms would be targeting different models.

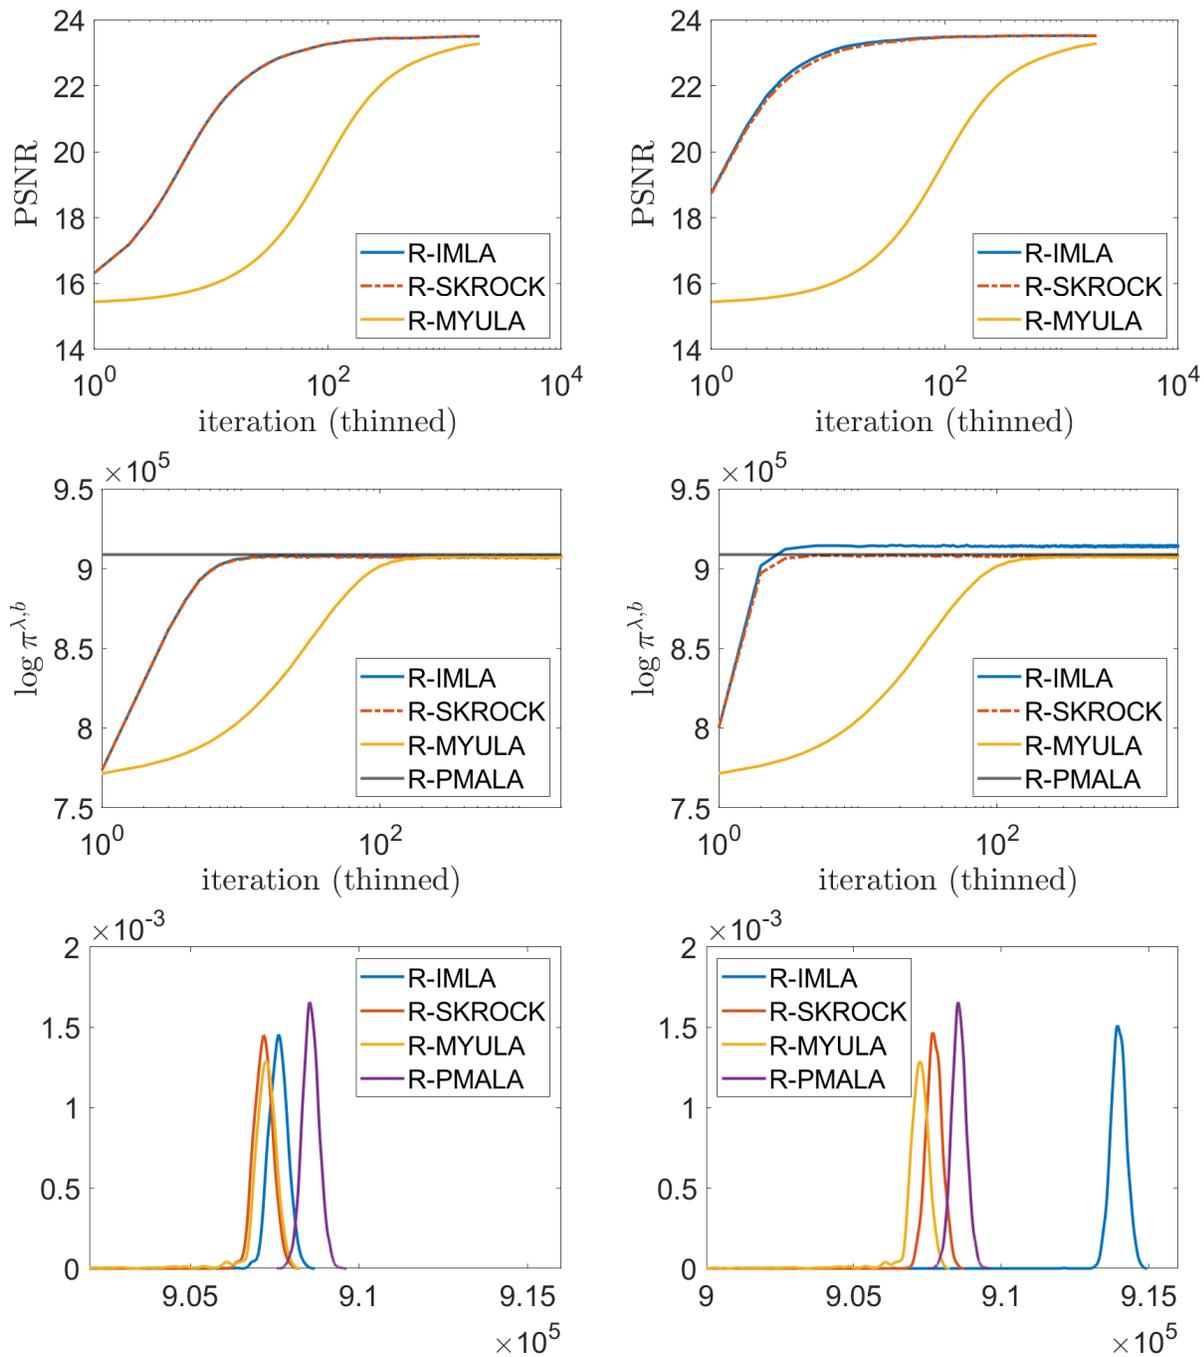


Figure 12: Poisson deconvolution experiment. PSNR of the running mean (top), $\log \pi^{\lambda,b}(x)$ (middle), and non-parametric density estimates of $\log \pi^{\lambda,b}(x)$ at stationarity (bottom). R-IMLA was run using $h = 6.65 \times 10^{-5}$ (left), and $h = 1.16 \times 10^{-3}$ (right) which is equivalent to the effective step size of R-SKROCK when $s = 10$ (left) or $s = 40$ (right), R-MYULA was run using $h = 1/L$, R-PMALA is shown as an unbiased reference.

Algorithm 2 R-IMLA

Require: $N \geq 0$, $\delta > 0$ and $X_0 \in \mathbb{R}^d$.**for** $n=0 : N-1$ **do****Draw**

$$\xi_n \sim \mathcal{N}(0, I_d)$$

Set

$$X_{n+1}^i = |\tilde{X}_{n+1}^i|, \quad \text{for all } i = 1, \dots, d,$$

$$\tilde{X}_{n+1} \leftarrow \arg \min_{x \in \mathbb{R}^d} 2U^\lambda \left(\frac{1}{2}x + \frac{1}{2}X_n \right) + \frac{1}{2\delta} \|x - X_n - \sqrt{2\delta}\xi_n\|^2 \quad (30)$$

end for

| stages | step size | cost/iter(s) | | factor | iterations | |
|--------|-----------------------|--------------|--------|---|------------|--------|
| | | R-SKROCK | R-IMLA | $\frac{t(\text{R-IMLA})}{t(\text{R-SKROCK})}$ | R-SKROCK | R-IMLA |
| 1 | 6.65×10^{-6} | | | | 800000 | |
| 10 | 6.65×10^{-5} | 0.22 | 0.22 | 0.99 | 80000 | 80000 |
| 20 | 2.82×10^{-4} | 0.42 | 0.28 | 0.67 | 40000 | 62000 |
| 40 | 1.16×10^{-3} | 0.84 | 0.45 | 0.54 | 20000 | 36000 |

Table 6: Cost in seconds per iteration for different step sizes used to run R-IMLA and the equivalent number stages for R-SKROCK. Listing the factor showing the cost improvement of R-IMLA and the number of iterations used to run the experiments. R-SKROCK with 1 stage is equivalent to R-MYULA.

4.2.1 Comparison according to integration time

For each of the algorithms R-MYULA, R-IMLA and R-SKROCK the estimated posterior means are reported in Figure 10 and the pixel-wise standard deviations are reported in Figure 11. In both figures, the results for R-SKROCK and R-IMLA are almost identical, which is consistent with the PSNRs obtained with the posterior means. In contrast, the reconstruction with R-MYULA is significantly worse which illustrates cases where R-MYULA has not converged.

In Figure 12 we compare the PSNR of the running mean, the trace of the statistic $\log \pi^\lambda(x)$, and non-parametric density estimates of the statistic $\log \pi^\lambda(x)$ at stationarity. In each case, R-IMLA and R-SKROCK show similar convergence speeds while being significantly faster than R-MYULA. To assess the true posterior distribution of the statistics $\log \pi^\lambda$, Figure 12 also reports the results obtained by running the MH-corrected R-PMALA method, which produces asymptotically unbiased samples from π^λ (we calculated the density plots displayed in Figure 12 by using 9.8×10^6 iterations of R-PMALA). We observe that all the algorithms produce accurate estimates of the posterior distribution of the statistics $\log \pi^\lambda$, with errors of the order of 0.5% or smaller. In this specific experiment IMLA achieves a slightly smaller bias than the other methods in the case of a smaller time step ($s = 10$), and the other methods are more accurate than IMLA when the step size is large ($s = 40$), but the differences are very small and this ordering could be modified by considering a marginally different MIV value or a different choice of θ .

We conclude that IMLA and SKROCK perform similarly in terms of integration time, and they significantly outperform MYULA, both in situations that are strongly log-concave as well also in settings involving weak log-concavity and constraints, as illustrated in this last set of experiments.

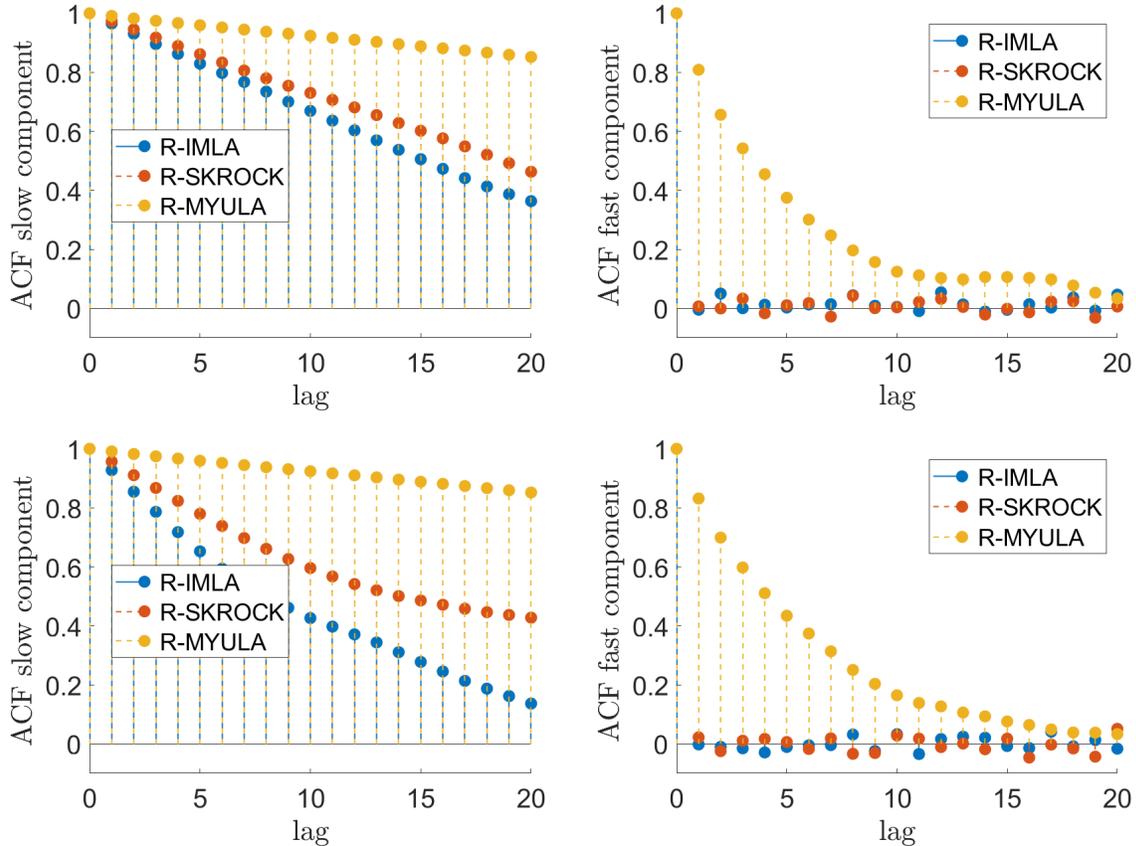


Figure 13: Comparison of ACF for the slowest (left) and the fastest component (right) for the Poisson deconvolution experiment 4.2.2. R-IMLA was run using $h = 2.82 \times 10^{-4}$ (top), and $h = 1.16 \times 10^{-3}$ (bottom) which is equivalent to the effective time step of R-SKROCK when $s = 20$ (top) or $s = 40$ (bottom), R-MYULA was run using $h = 1/L$.

4.2.2 Comparison according to computational cost

The experiments that we have conducted so far suggest that IMLA and SKROCK perform similarly when compared without taking into account the computational cost per iteration. The computational cost per iteration of SKROCK depends directly on the number of internal stages s used, whereas the computational cost of IMLA depends on the choice of solver used. To demonstrate that IMLA can provide a competitive alternative to SKROCK despite requiring implicit steps, and to highlight the importance of choosing a suitable solver, we now perform a comparison where IMLA is implemented by using a limited memory quasi-Newton solver⁹ [4]. As in the previous section, we use the same time step for R-IMLA and R-SKROCK, but we now run R-IMLA for the same clock time as for R-SKROCK (see Table 6). We observe that for $s = 10$, IMLA and SKROCK complete the same number of iterations in the same clock time, whereas IMLA is 50% more computationally efficient than SKROCK when $s = 20$, and almost twice as efficient in the large step regime ($s = 40$). Figure 13 compares the convergence speed of IMLA and SKROCK with time-normalised ACF plots, for the time steps related to $s = 20$ and $s = 40$. Observe that IMLA exhibits a moderately higher convergence speed than SKROCK for the moderately large step ($s = 20$) and a significantly better speed when a large step is used ($s = 40$).

⁹we use the FORTRAN implementation <https://github.com/stephenbeckr/L-BFGS-B-C> with a MATLAB wrapper, and set the tolerance level to 10^{-4} .

5 Discussion and conclusions

This paper presented IMLA, a new proximal MCMC methodology for Bayesian inference in imaging inverse problems with an underlying convex geometry. Similarly to previous proximal MCMC methods such as MYULA [11] and SKROCK [31], IMLA is derived from a discrete-time approximation of an overdamped Langevin diffusion that targets the posterior distribution of interest (or a Moreau-Yosida approximation when the original target is not smooth). IMLA takes the form of a stochastic relaxed proximal-point iteration that admits two complementary interpretations. For models that are smooth or regularised by Moreau-Yosida smoothing, the approximation constructed by IMLA is equivalent to an implicit midpoint approximation targeting the posterior distribution of interest. This discretisation is asymptotically unbiased for Gaussian targets and shown to converge in an accelerated manner for any target that is κ -strongly log-concave, (i.e., requiring in the order of $\sqrt{\kappa}$ iterations to converge, similarly to accelerated optimisation schemes, with κ the condition number), comparing favourably to previous proximal MCMC methods which are biased and either not provably accelerated or only provably accelerated for Gaussian targets. For models that are not smooth, the algorithm is equivalent to a Leimkuhler–Matthews discretisation of a Langevin diffusion targeting a Moreau-Yosida approximation of the posterior distribution of interest. For targets that are κ -strongly log-concave, the provided non-asymptotic convergence analysis also identifies the optimal time step which maximizes the convergence speed. We demonstrated the proposed methodology with a range of numerical experiments, including some tractable toy models designed to illustrate different classes of models, as well as two challenging experiments related to image deconvolution with Gaussian and Poisson noise.

From a methodological perspective, future works could focus on embedding IMLA within more Bayesian inference complex schemes, such as the stochastic approximation proximal gradient schemes for empirical Bayesian estimation [38] and proximal nested sampling [5] for Bayesian model selection. Extending IMLA to mildly non-convex problems is also an important perspective for future work. Moreover, with regard to applications, we believe that IMLA has great potential for low-photon scientific imaging applications, where uncertainty quantification is critical and where the combination of challenging noise statistics (e.g., Binomial and Geometric noise) and the lack of ground truth data make log-concave priors a very convenient option [24].

A Appendix

The appendix contains collected proofs of the results from Section 3.2 and details of the model used Section 3.3.

A.1 Proof of Proposition 3.2

Before giving the proof of Proposition 3.4 we give two auxiliary lemmas which summarise the properties of the constant C as a function of δ and the bias for $\theta = 1$. As C has the same form for both the Gaussian and the strongly log-concave setting the following Lemma holds for both cases.

Lemma A.1. *Let C be given by (22) or (26), then we have*

$$C = \begin{cases} \frac{(1-\theta)L\delta-1}{\theta L\delta+1} & \delta > \delta_* \\ \frac{1-(1-\theta)m\delta}{1+\theta m\delta} & \delta \leq \delta_* \end{cases}$$

where δ_* is the value of δ which minimises C and is given by

$$\delta_* = \frac{(2\theta - 1)(L + m) + \sqrt{(1 - 2\theta)^2(L + m)^2 + 16(1 - \theta)\theta Lm}}{4(1 - \theta)\theta Lm}$$

The constant C here corresponds to constant ρ in [15, Theorem 2] with $\delta = h/2$, we include the proof of the properties of C here for completeness.

Proof of Lemma A.1. We split the range of z in the maximum of C based on when the term $R_1(-z)$ is positive, for $z \leq (1-\theta)^{-1}$, or negative, for $z \geq (1-\theta)^{-1}$:

$$C = \left(\max_{z \in [m\delta, (1-\theta)^{-1} \vee L\delta]} R_1(-z) \right) \wedge \left(\max_{z \in [(1-\theta)^{-1} \wedge m\delta, L\delta]} -R_1(-z)^2 \right)$$

Here we use the notation \wedge to denote the maximum and \vee to denote the minimum. Since $z \mapsto R_1(-z)$ is a decreasing function we can write C as

$$C = \begin{cases} R_1(-m\delta) & \text{if } L\delta \leq (1-\theta)^{-1}, \\ R_1(-m\delta) \wedge (-R_1(-L\delta)) & \text{if } m\delta \leq (1-\theta)^{-1} \leq L\delta, \\ -R_1(-L\delta) & \text{if } (1-\theta)^{-1} \leq m\delta. \end{cases}$$

As $R_1(-m\delta)$ is decreasing in δ and $-R_1(-L\delta)$ is increasing in δ to determine the maximum we only need to determine for what value of δ they are equal. Observe that $R_1(-m\delta) = -R_1(-L\delta)$ if $\delta = \delta_*$ with δ_* as in the statement of Lemma A.1. Therefore we have that $C = R_1(-m\delta)$ for $\delta \leq \delta_*$ and $C = -R_1(-L\delta)$ for $\delta > \delta_*$, using the definition of R_1 we obtain the claim of the lemma. \square

Proposition A.2. *Set $\theta = 1$ then the bias may be bounded as*

$$W_2(\pi, \tilde{\pi}) \leq \min \left(\frac{1}{2} \sqrt{d\delta}, \frac{\sqrt{d\delta}}{4\sigma_{\min}} \right). \quad (31)$$

Proof of Proposition A.2. Taking the limit as $n \rightarrow \infty$ in (19) we have

$$W_2(\pi, \tilde{\pi})^2 = \sum_{i=1}^d \phi(\sigma_i, \delta)^2$$

where, from (21),

$$\phi(\sigma, \delta) = \sigma \left(1 - \frac{1}{\sqrt{1 + \frac{\delta}{2\sigma^2}}} \right).$$

It is easy to prove the bounds

$$\phi \leq \frac{\delta}{4\sigma}, \quad \phi \leq \sigma. \quad (32)$$

Then

$$\begin{aligned} W_2(\pi, \tilde{\pi})^2 &= \sum_{\delta\sigma_i^{-2} \leq 4} \phi(\sigma_i, \delta)^2 + \sum_{\delta\sigma_i^{-2} > 4} \phi(\sigma_i, \delta)^2 \\ &\leq \sum_{\delta\sigma_i^{-2} \leq 4} \frac{\delta^2}{16\sigma_i^2} + \sum_{\delta\sigma_i^{-2} > 4} \sigma_i^2 \\ &\leq \sum_{\delta\sigma_i^{-2} \leq 4} \frac{\delta}{4} + \sum_{\delta\sigma_i^{-2} > 4} \frac{\delta}{4} \end{aligned}$$

so that $W_2(\pi, \tilde{\pi}) \leq \sqrt{d\delta}/2$. Furthermore, from the first bound in (32) we readily obtain $W_2(\pi, \tilde{\pi}) \leq \frac{\sqrt{d\delta}}{4\sigma_{\min}}$. \square

Proof of Proposition 3.2. We first consider the case $\theta = 1/2$. Formula (21) shows that in this case there is no bias: $\tilde{\pi} = \pi$. From (20), in order to ensure that $W_2(\pi, Q_n) \leq \epsilon$ we have to take n such that

$$n |\log(C)| \geq \log(W_2(\pi, Q_0)) - \log(\epsilon).$$

The best performance will correspond to the value of δ that minimizes C in (22). By Lemma A.1, C is minimized when $\delta = \delta_* = 2\sigma_{\max}\sigma_{\min}$, and

$$C_{\theta=1/2}^{\text{opt}} = \frac{1 - \sigma_{\max}^{-2}\delta_*/2}{1 + \sigma_{\max}^{-2}\delta_*/2}.$$

Since we are interested in case where σ_d is small, leading to small δ_* , $C_{\theta=1/2}^{\text{opt}} \approx 1 - \sigma_{\max}^{-2}\delta_* = 1 - 2/\sqrt{\kappa}$ and $|\log C| \approx 2/\sqrt{\kappa}$. This yields

$$n \approx \frac{\sqrt{\kappa}}{2} [\log(W_2(\pi, Q_0)) - \log(\epsilon)] \quad (33)$$

where we note that in the righthand side there is *no dependence on d , and the dependence on ϵ is logarithmic.*

Now consider the case $\theta = 1$. By (20) in order to ensure that $W_2(\pi, Q_n) \leq \epsilon$ we first take δ small enough to get $W_2(\pi, \tilde{\pi}) \leq \epsilon/2$ and then n large enough to get $C^n W_2(\pi, Q_0) \leq \epsilon/2$, i.e.

$$n|\log(C)| \geq \log(W_2(\pi, Q_0)) - \log(2^{-1}\epsilon).$$

From (31), the bias has the upper bound $\sqrt{d\delta}/2$, which is independent of κ . This independence is achieved at the prize of the bound being proportional to $\sqrt{\delta}$ rather than to δ , as one may naively expect for a first-order integrator. For δ small, more precisely for $\delta \leq 4\sigma_{\min}^2$, (31) yields the upper bound $\sqrt{d\delta}/(4\sigma_{\min})$, with a first-order, i.e. $\mathcal{O}(\delta)$, behaviour. However the regime $\delta \leq 4\sigma_{\min}^2$ is without any practical interest, because for values of δ of the order of σ_{\min}^2 , one may use the explicit, $\theta = 0$, integrator, with a much lower computational complexity per time-step.

According to (31), we choose the value of δ

$$\delta^* = \max\left(\frac{\epsilon^2}{d}, \frac{2\epsilon\sigma_{\min}}{\sqrt{d}}\right)$$

Unlike the optimal δ for $\theta = 1/2$ we see that δ^* remains bounded away from zero as $\sigma_{\min} \downarrow 0$. On the other hand δ^* depends on ϵ and d . Since $R_1 = 1/(1 + \delta\lambda)$ is a positive, monotonically decreasing function of $\lambda > 0$, the contractivity constant is then

$$C_{\theta=1} = R_1(z_1) = 1/(1 + \sigma_{\max}^{-2}\delta^*).$$

and, when δ^* is small, $C_{\theta=1} \approx 1 - \sigma_{\max}^{-2}\delta^*$ and $|\log C| \approx \sigma_{\max}^{-2}\delta^*$ which leads to

$$n \approx \min\left(\frac{d\sigma_{\max}^2}{\epsilon^2}, \frac{\sqrt{d\kappa}\sigma_{\max}}{2\epsilon}\right) [\log(W_2(\pi, Q_0)) - \log(2^{-1}\epsilon)]. \quad (34)$$

Comparison with (33) shows that now there is a dependence of n on d and, in addition, the dependence on ϵ is now much worse than logarithmic. On the other hand the right hand side of (34) remains bounded as $\kappa \uparrow \infty$. The bounds show that $\theta = 1/2$ will outperform $\theta = 1$ unless $d/(\sqrt{\kappa}\epsilon^2)$ is small. \square

A.2 Proof of Theorem 3.3

Proof of Theorem 3.3. We will estimate $W_2(Q_n, \pi)$ by setting up a coupling between the numerical scheme X_k and a Langevin SDE L_t which is in stationarity with π as its invariant distribution. Let $L_0 \sim \pi$ and set

$$L_t = L_0 - \int_0^t \nabla U(L_s) ds + \sqrt{2}W_t.$$

Let (X_0, L_0) be such that $\mathbb{E}[|X_0 - L_0|^2] = W_2(Q_0, \pi)$. Define $L_s^k = L_{k\delta+s}$ and

$$D_k = X_k - L_0^k.$$

We will set up a synchronous coupling by setting $\xi_k = \delta^{-1/2}(W_{k\delta} - W_{(k-1)\delta})$ then

$$X_k = X_{k-1} - \delta \nabla U(\theta X_k + (1-\theta)X_{k-1}) + \sqrt{2}(W_{k\delta} - W_{(k-1)\delta}) + E_k.$$

Here E_k is given by

$$\begin{aligned} E_k &= X_k - X_{k-1} + \nabla f(\theta X_k + (1-\theta)X_{k-1})\delta - \sqrt{2}(W_{k\delta} - W_{(k-1)\delta}) \\ &= \delta \nabla F(X_k; X_{k-1}, \xi_k). \end{aligned}$$

By assumption we have $\|E_k\| \leq \delta\varepsilon$. Subtracting the expression for L_t from X_k we have

$$D_{k+1} = D_k - \delta \nabla U(\theta D_{k+1} + (1-\theta)D_k + \theta L_\delta^k + (1-\theta)L_0^k) + \int_0^\delta \nabla U(L_s^k) ds + E_k. \quad (35)$$

By the mean value theorem there exists a matrix $\bar{H} = \bar{H}(L_k, L_{k+1}, D_k, D_{k+1})$ with $\sigma(\bar{H}) \subseteq [m, L]$ such that

$$\nabla U(\theta D_{k+1} + (1-\theta)D_k + \theta L_\delta^k + (1-\theta)L_0^k) = \nabla U(\theta L_\delta^k + (1-\theta)L_0^k) + \theta \bar{H} D_{k+1} + (1-\theta) \bar{H} D_k. \quad (36)$$

Substituting (36) into (35) we have

$$(1 + \theta \delta \bar{H}) D_{k+1} = (1 - (1-\theta)\delta \bar{H}) D_k + \int_0^\delta [\nabla U(L_s^k) - \nabla U(\theta L_\delta^k + (1-\theta)L_0^k)] ds + E_k.$$

Observe that this now has a similar form to the Gaussian case since the D_k terms correspond to the numerical contraction of the Gaussian scheme with $\Sigma^{-1} = \bar{H}$, the integral term controls the bias and E_k corresponds to the error incurred by not solving the implicit step exactly.

Taking L^2 -norms (with respect to \mathbb{E}) we have

$$\begin{aligned} \|D_{k+1}\|_{L^2} &\leq \|R_1(-\delta \bar{H}) D_k\|_{L^2} \\ &\quad + \int_0^\delta \|R_2(-\delta \bar{H})(\nabla U(L_s^k) - \nabla U(\theta L_\delta^k + (1-\theta)L_0^k))\|_{L^2} ds + \|R_2(-\delta \bar{H}) E_k\|_{L^2}. \end{aligned}$$

Recall R_1, R_2 were defined as functions on \mathbb{R} in (18), we lift these to functions on matrices. Note that $R_1(-\delta \bar{H})$ is bounded by C , where C is given by (26), and $R_2(-\delta \bar{H}) \leq R_2(-m\delta)$. Using these bounds and that $\|E_k\| \leq \delta\varepsilon$ we have

$$\|D_{k+1}\|_{L^2} \leq C \|D_k\|_{L^2} + \frac{1}{1 + \theta \delta m} \int_0^\delta \|(\nabla U(L_s^k) - \nabla U(\theta L_\delta^k + (1-\theta)L_0^k))\|_{L^2} ds + \frac{\varepsilon \delta}{1 + \theta \delta m}. \quad (37)$$

Now consider

$$\int_0^t \|(\nabla U(L_s^k) - \nabla U(\theta L_\delta^k + (1-\theta)L_0^k))\|_{L^2} ds$$

We will now follow the procedure of [10] to bound this term. Using that U is L -Lipschitz we have

$$\|\nabla U(L_s^k) - \nabla U(\theta L_\delta^k + (1-\theta)L_0^k)\|_{L^2} \leq L\theta \|L_s^k - L_\delta^k\|_{L^2} + L(1-\theta) \|L_s^k - L_0^k\|_{L^2}. \quad (38)$$

By the definition of L_s we have

$$\begin{aligned} \|L_s - L_0\|_{L^2} &= \left\| - \int_0^s \nabla U(L_r) dr + \sqrt{2} W_s \right\|_{L^2} \\ &\leq \left\| - \int_0^s \nabla U(L_r) dr \right\|_{L^2} + \sqrt{2} \|W_s\|_{L^2} \\ &\leq \int_0^s \|\nabla U(L_r)\|_{L^2} dr + \sqrt{2} s d. \end{aligned}$$

Since L_r is stationary we have

$$\|(L_s - L_0)\|_{L^2} \leq s\|\nabla U(L_0)\|_{L^2} + \sqrt{2sd}.$$

By stationarity we also have

$$\|(L_\delta - L_s)\|_{L^2} \leq (\delta - s)\|\nabla U(L_0)\|_{L^2} + \sqrt{2(\delta - s)d}.$$

By [10, Lemma 3] we have $\|\nabla U(L_0)\|_{L^2} \leq \sqrt{Ld}$. Combining these estimates with (38)

$$\begin{aligned} & \|\nabla U(L_s) - \nabla U(\theta L_1 + (1 - \theta)L_0)\|_{L^2} \leq \\ & L\theta((\delta - s)\sqrt{Ld} + \sqrt{2(\delta - s)d}) + L(1 - \theta)(s\sqrt{Ld} + \sqrt{2sd}). \end{aligned}$$

Integrating gives

$$\begin{aligned} & \int_0^\delta \|\nabla U(L_s) - \nabla U(\theta L_1 + (1 - \theta)L_0)\|_{L^2} ds \leq L\theta\left(\frac{1}{2}\delta^2\sqrt{Ld} + \frac{2}{3}\delta^{\frac{3}{2}}\sqrt{2d}\right) \\ & + L(1 - \theta)\left(\frac{1}{2}\delta^2\sqrt{Ld} + \frac{2}{3}\delta^{\frac{3}{2}}\sqrt{2d}\right) = \frac{1}{2}\delta^2 L^{\frac{3}{2}}\sqrt{d} + \frac{2}{3}L\delta^{\frac{3}{2}}\sqrt{2d}. \end{aligned} \quad (39)$$

Combining (37) and (39) we have

$$\|D_{k+1}\|_{L^2} \leq C\|D_k\|_{L^2} + \frac{\frac{1}{2}\delta^2 L^{\frac{3}{2}}\sqrt{d} + \frac{2}{3}L\delta^{\frac{3}{2}}\sqrt{2d} + \varepsilon\delta}{1 + \theta\delta m}.$$

Therefore we have

$$\|D_{k+1}\|_{L^2} \leq C^{k+1}\|D_0\|_{L^2} + \frac{1 - C^{k+2}}{1 - C} \frac{\frac{1}{2}\delta^2 L^{\frac{3}{2}}\sqrt{d} + \frac{2}{3}L\delta^{\frac{3}{2}}\sqrt{2d} + \varepsilon\delta}{1 + \theta\delta m}.$$

Using that $W_2(Q_n, \pi) \leq D_n$ and $D_0 = W_2(Q_0, \pi)$ we have

$$W_2(Q_n, \pi) \leq C^n W_2(Q_0, \pi) + \frac{1 - C^{n+1}}{1 - C} \frac{\frac{1}{2}\delta^2 L^{\frac{3}{2}}\sqrt{d} + \frac{2}{3}L\delta^{\frac{3}{2}}\sqrt{2d} + \varepsilon\delta}{1 + \theta\delta m}.$$

□

A.3 Proof of Proposition 3.4

Proof of Proposition 3.4. Fix $\epsilon > 0$, then we ensure that $W_2(Q_n, \pi) \leq \epsilon$ holds by first finding δ such that the bias $\lim_{n \rightarrow \infty} W_2(Q_n, \pi) \leq \epsilon/2$ and then finding n such that the contraction term $C^n W_2(Q_0, \pi) \leq \epsilon/2$. If these hold we have by (25) that $W_2(Q_n, \pi) \leq \epsilon$ using that $1 - C^{n+1} \leq 1$.

Setting $\delta = \delta_*$ gives the optimal contraction rate and hence we either have that $\delta = \delta_*$ gives a bias smaller than $\epsilon/2$ or we take δ to be the largest time step such that the bias is less than $\epsilon/2$. For this reason we calculate the smallest value of ϵ such that the bias is less than $\epsilon/2$ when $\delta = \delta_*$. By taking the limit $n \rightarrow \infty$ in (25) we have the following estimate of the bias:

$$\lim_{n \rightarrow \infty} W_2(Q_n, \pi) \leq \frac{1}{1 - C} \frac{\frac{1}{2}\delta_*^2 L^{\frac{3}{2}}\sqrt{d} + \frac{2}{3}L\delta_*^{\frac{3}{2}}\sqrt{2d}}{1 + \frac{1}{2}\delta_* m}. \quad (40)$$

Setting $\delta = \delta_*$ in (40) and using Lemma A.1 to determine C we have

$$\begin{aligned} \lim_{n \rightarrow \infty} W_2(Q_n, \pi) & \leq \frac{2 + m\delta_*}{2m\delta_*} \frac{\frac{1}{2}\delta_*^2 L^{\frac{3}{2}}\sqrt{d} + \frac{2}{3}L\delta_*^{\frac{3}{2}}\sqrt{2d}}{1 + \frac{1}{2}\delta_* m} \\ & = \frac{\sqrt{d}\kappa}{\sqrt{m}} \left(1 + \frac{4}{3}\kappa^{-\frac{1}{4}}\right). \end{aligned}$$

Hence if $\varepsilon \geq \frac{\sqrt{d}\kappa}{\sqrt{m}} \left(1 + \frac{4}{3}\kappa^{-\frac{1}{4}}\right)$ then the optimal choice of time step $\delta = \delta_*$ is valid.

On the other hand if $\varepsilon < \frac{\sqrt{d}\kappa}{\sqrt{m}} \left(1 + \frac{4}{3}\kappa^{-\frac{1}{4}}\right)$ then we need to determine the largest δ for which the bias is less than $\varepsilon/2$.

Case 1: If $\delta > \frac{32}{9}L^{-1}$ then we have that

$$\lim_{n \rightarrow \infty} W_2(Q_n, \pi) \leq \frac{1}{1-C} \frac{\delta^2 L^{\frac{3}{2}} \sqrt{d}}{1 + \frac{1}{2}\delta m}.$$

Since $\delta < \delta_*$ we have by Lemma A.1

$$\lim_{n \rightarrow \infty} W_2(Q_n, \pi) \leq \sqrt{L}\kappa\delta\sqrt{d}.$$

Therefore we can ensure that the bias is less than $\varepsilon/2$ by using

$$\delta = \frac{\varepsilon}{2\sqrt{Ld}\kappa}.$$

Note this is permissible provided $\delta > \frac{32}{9}L^{-1}$ which requires

$$\varepsilon > \frac{64}{9}\sqrt{d}\kappa L^{-\frac{1}{2}}.$$

Case 2: If $\delta < \frac{32}{9}L^{-1}$, then as above we have

$$\lim_{n \rightarrow \infty} W_2(Q_n, \pi) \leq \frac{1}{m} \frac{4}{3} L \delta^{\frac{1}{2}} \sqrt{2d}.$$

This is less than $\varepsilon/2$ if we set

$$\delta = \frac{9\varepsilon^2}{128d\kappa^2}.$$

Which gives a permissible value of δ for

$$\varepsilon^2 < d\kappa^2 \frac{4096}{81} L^{-1} = \left(\frac{64}{9}\sqrt{d}\kappa L^{-\frac{1}{2}}\right)^2.$$

It remains to find the value of n such that $C^n W_2(\pi_0, \pi) \leq \varepsilon/2$, that is we require

$$n \geq \frac{\log(\varepsilon/2) - \log(W_2(\pi, Q_0))}{|\log C|}.$$

Since we always use $\delta \leq \delta_*$ we have $C = 1 - 4m\delta/(2 + m\delta)$ and we use the approximation $\log(C) \approx 2m\delta$.

From the above analysis we have that the best time step to obtain accuracy ε is

$$\delta = \min \left\{ \frac{2}{\sqrt{Lm}}, \frac{\varepsilon}{2\kappa\sqrt{Ld}}, \frac{9}{128} \frac{\varepsilon^2}{d\kappa} \right\}.$$

Therefore we have

$$n \approx \frac{1}{2m} \max \left\{ \frac{\sqrt{Lm}}{2}, \frac{2\kappa\sqrt{Ld}}{\varepsilon}, \frac{128}{9} \frac{d\kappa}{\varepsilon^2} \right\} [\log(\varepsilon/2) - \log(W_2(\pi, Q_0))].$$

□

A.4 Details of illustrative experiment 3.3

The posterior distribution $\pi(x)$ for the denoising problem using a Gaussian mixture prior model can be written for $(x)_{i=1}^d \in \mathbb{R}^d$ and data $(y)_{i=1}^d \in \mathbb{R}^d$ as follows

$$\begin{aligned} \pi(x) &= \prod_{i=1}^d [\omega(y_i)f(x_i, y_i, \mu_0, \delta_0^2) + (1 - \omega(y_i))f(x_i, y_i, \mu_1, \delta_1^2)] \\ f(x_i, y_i, \mu_k, \delta_k^2) &= \exp\left(-\frac{(x_i - \mu_k(y_i))^2}{2\delta_k^2}\right) \\ \mu_k(y_i) &= \left(\frac{y_i}{\sigma^2} + \frac{m_k}{\sigma_k^2}\right)\delta_k^2; \quad \delta_k^2 = \frac{\sigma^2\sigma_k^2}{\sigma^2 + \sigma_k^2} \\ \omega(y_i) &= \frac{C_0(y_i)\tilde{\omega}}{C}; \quad C = \tilde{\omega}C_0(y_i) + (1 - \tilde{\omega})C_1(y_i). \\ C_k(y_i) &= \frac{1}{\sqrt{2\pi(\sigma_k^2 + \sigma^2)}} \exp\left(-\frac{(m_k - y_i)^2}{2(\sigma_k^2 + \sigma^2)}\right) \end{aligned}$$

Refer to Table 1 for the parameter values $(m_0, m_1, \sigma_0^2, \sigma_1^2, \tilde{\omega})$ and noise variance σ^2 used in the experiment.

References

- [1] A. Abdulle, I. Almuslimani, and G. Vilmart. Optimal explicit stabilized integrator of weak order 1 for stiff and ergodic stochastic differential equations. *SIAM-ASA J. Uncertain. Quantif.*, 6(2):937–964, 2018.
- [2] J. Adler and O. Öktem. Deep bayesian inversion. *arXiv preprint arXiv:1811.05910*, 2018.
- [3] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb. Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174, 2019.
- [4] R. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM J Sci Comput*, 16(5):1190–1208, 1995.
- [5] X. Cai, J. D. McEwen, and M. Pereyra. Proximal nested sampling for high-dimensional bayesian model selection. *Stat Comput*, 32(5), 2022.
- [6] X. Cai, M. Pereyra, and J. D. McEwen. Uncertainty quantification for radio interferometric imaging – I. Proximal MCMC methods. *Mon. Notices Royal Astron. Soc.*, 480(3):4154–4169, 2018.
- [7] A. Chambolle and T. Pock. An introduction to continuous optimization for imaging. *Acta Numerica*, 25:161–319, 2016.
- [8] P. L. Combettes and JC. Pesquet. *Proximal Splitting Methods in Signal Processing*, pages 185–212. Springer New York, New York, NY, 2011.
- [9] M.-C. Corbineau, D. Kouamé, É. Chouzenoux, J.-Y. Tourneret, and J.-C. Pesquet. Preconditioned P-ULA for joint deconvolution-segmentation of ultrasound images. *IEEE Signal Process. Lett.*, 26(10):1456–1460, 2019.
- [10] A. S. Dalalyan and A. Karagulyan. User-friendly guarantees for the langevin monte carlo with inaccurate gradient. *Stoch. Process. Their Appl*, 129(12):5278–5311, 2019.
- [11] A. Durmus and E. Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *Ann. Appl. Probab.*, 27(3):1551–1587, 2017.

- [12] A. Durmus, E. Moulines, and M. Pereyra. Efficient bayesian computation by proximal Markov chain Monte Carlo: When Langevin meets Moreau. *SIAM J. Imaging Sci.*, 11(1):473–506, 2018.
- [13] A. Goujon, S. Neumayer, P. Bohra, S. Ducotterd, and M. Unser. A neural-network-based convex regularizer for image reconstruction. *arXiv preprint arXiv:2211.12461*, 2022.
- [14] E. Hairer, S. P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations: Vol.II. Stiff and Differential–Algebraic Problems*, volume 14 of *Springer ser. comput. math.* Springer, 1991.
- [15] L. Hodgkinson, R. Salomone, and F. Roosta. Implicit langevin algorithms for sampling from log-concave densities. *J Mach Learn Res*, 22:19–292, 2021.
- [16] B. Hosseini and N. Nigam. Well-posed bayesian inverse problems: Priors with exponential tails. *SIAM-ASA J. Uncertain.*, 5(1):436–465, 2017.
- [17] J. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems*. Springer, Dordrecht, 2005.
- [18] B. Kawar, M. Elad, S. Ermon, and J. Song. Denoising diffusion restoration models. In *ICLR Workshop on Deep Generative Models for Highly Structured Data*, 2022.
- [19] E. Kobler, T. Klatzer, K. Hammernik, and T. Pock. Variational networks: Connecting variational methods and deep learning. In *Pattern Recognit.*, pages 281–293, 2017.
- [20] R. Laumont, V. D. Bortoli, A. Almansa, J. Delon, A. Durmus, and M. Pereyra. Bayesian imaging using plug & play priors: When langevin meets tweedie. *SIAM J. Imaging Sci.*, 15(2):701–737, 2022.
- [21] B. Leimkuhler and C. Matthews. Rational Construction of Stochastic Numerical Methods for Molecular Sampling. *Appl Math Res Express*, 2013(1):34–56, 2012.
- [22] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Understanding and evaluating blind deconvolution algorithms. In *CVPR*, pages 1964–1971. IEEE Computer Society, 2009.
- [23] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int’l Conf. Computer Vision*, volume 2, pages 416–423, 2001.
- [24] S. Melidonis, P. Dobson, Y. Altmann, M. Pereyra, and K. Zygalakis. Efficient bayesian computation for low-photon imaging problems. *SIAM J. Imaging Sci.*, 16(3):1195–1234, 2023.
- [25] W. Mou, Y.-A. Ma, M. J. Wainwright, P. L. Bartlett, and M. I. Jordan. High-order langevin diffusion yields an accelerated mcmc algorithm. *J Mach Learn Res*, 22(42):1–41, 2021.
- [26] S. Mukherjee, S. Dittmer, Z. Shumaylov, S. Lunz, O. Öktem, and C.-B. Schönlieb. Learned convex regularizers for inverse problems. *arXiv preprint arXiv:2008.02839*, 2021.
- [27] M. Pereyra. Proximal Markov chain Monte Carlo algorithms. *Statistics and Computing*, 26(4):745–760, 2016.
- [28] M. Pereyra. Maximum-a-posteriori estimation with bayesian confidence regions. *SIAM J. Imaging Sci.*, 10(1):285–302, 2017.
- [29] M. Pereyra. Revisiting maximum-a-posteriori estimation in log-concave models. *SIAM J. Imaging Sci.*, 12(1):650–670, 2019.
- [30] M. Pereyra, P. Schniter, E. Chouzenoux, J.-C. Pesquet, J.-Y. Tournier, A. O. Hero, and S. McLaughlin. A Survey of Stochastic Simulation and Optimization Methods in Signal Processing. *IEEE J Sel Top Signal Process*, vol. 10(2):pp. 224–241, 2015.

- [31] M. Pereyra, L. A. Vargas-Mieles, and K. C. Zygalakis. Accelerating proximal Markov chain Monte Carlo by using an explicit stabilized method. *SIAM J. Imaging Sci.*, 13(2):905–935, 2020.
- [32] M. Pereyra, L. A. Vargas-Mieles, and K. C. Zygalakis. The split gibbs sampler revisited: improvements to its algorithmic structure and augmented target distribution. *arXiv preprint arXiv:2206.13894*, 2023.
- [33] G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. *Found. Trends Mach. Learn.*, 11(5-6):355–607, 2019.
- [34] A. Repetti, M. Pereyra, and Y. Wiaux. Scalable bayesian uncertainty quantification in imaging inverse problems via convex optimization. *SIAM J. Imaging Sci.*, 12(1):87–118, 2019.
- [35] A. Repetti, M. Pereyra, and Y. Wiaux. Scalable bayesian uncertainty quantification in imaging inverse problems via convex optimization. *SIAM J. Imaging Sci.*, 12(1):87–118, 2019.
- [36] C. P. Robert. *The Bayesian choice: from decision-theoretic foundations to computational implementation*, volume 2. Springer, 2007.
- [37] 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 *ICML*, volume 97 of *Proc. Mach. Learn. Res.*, pages 5546–5557. PMLR, 2019.
- [38] A. F. Vidal, V. De Bortoli, M. Pereyra, and A. Durmus. Maximum likelihood estimation of regularization parameters in high-dimensional inverse problems: An empirical bayesian approach part i: Methodology and experiments. *SIAM J. Imaging Sci.*, 13(4):1945–1989, 2020.
- [39] G. Vilmart. Postprocessed integrators for the high order integration of ergodic sdes. *SIAM J. Sci. Comput.*, 37(1):A201–A220, 2015.
- [40] M. Vono, N. Dobigeon, and P. Chainais. Split-and-augmented gibbs sampler—application to large-scale inference problems. *IEEE Trans. Signal Process.*, 67(6):1648–1661, 2019.
- [41] M. Vono, D. Paulin, and A. Doucet. Efficient mcmc sampling with dimension-free convergence rate using admm-type splitting. *J Mach Learn Res*, 23(25):1–69, 2022.
- [42] A. Wibisono. Sampling as optimization in the space of measures: The langevin dynamics as a composite optimization problem. volume 75 of *Proc. Mach. Learn. Res.*, pages 2093–3027, 2018.
- [43] D. Yao, S. McLaughlin, and Y. Altmann. Patch-based image restoration using expectation propagation. *SIAM J. Imaging Sci.*, 15(1):192–227, 2022.
- [44] Q. Zhou, T. Yu, X. Zhang, and J. Li. Bayesian inference and uncertainty quantification for medical image reconstruction with poisson data. *SIAM J. Imaging Sci.*, 13(1):29–52, 2020.
- [45] K. C. Zygalakis. On the existence and the applications of modified equations for stochastic differential equations. *SIAM J. Sci. Comput.*, 33(1):102–130, 2011.