Efficient Bayesian computation by proximal Markov chain Monte Carlo: when Langevin meets Moreau

Alain Durmus*, Éric Moulines†, and Marcelo Pereyra‡

Abstract. Modern imaging methods rely strongly on Bayesian inference techniques to solve challenging imaging problems. Currently, the predominant Bayesian computation approach is convex optimisation, which scales very efficiently to high dimensional image models and delivers accurate point estimation results. However, in order to perform more complex analyses, for example image uncertainty quantification or model selection, it is necessary to use more computationally intensive Bayesian computation techniques such as Markov chain Monte Carlo methods. This paper presents a new and highly efficient Markov chain Monte Carlo methodology to perform Bayesian computation for high dimensional models that are log-concave and non-smooth, a class of models that is central in imaging sciences. The methodology is based on a regularised unadjusted Langevin algorithm that exploits tools from convex analysis, namely Moreau-Yoshida envelopes and proximal operators, to construct Markov chains with favourable convergence properties. In addition to scaling efficiently to high dimensions, the method is straightforward to apply to models that are currently solved by using proximal optimisation algorithms. We provide a detailed theoretical analysis of the proposed methodology, including asymptotic and non-asymptotic convergence results with easily verifiable conditions, and explicit bounds on the convergence rates. The proposed methodology is demonstrated with four experiments related to image deconvolution and tomographic reconstruction with total-variation and ℓ_1 priors, where we conduct a range of challenging Bayesian analyses related to uncertainty quantification, hypothesis testing, and model selection in the absence of ground truth.

Key words. Mathematical imaging; inverse problems; Bayesian inference; Markov chain Monte Carlo methods; convex optimisation; uncertainty quantification; model selection.

AMS subject classifications. primary 65C40, 68U10, 62F15; secondary 65C60, 65J22.

1. Introduction.

Image estimation problems are ubiquitous in science and engineering. For example, problems related to image denoising [24], deconvolution [5], compressive sensing reconstruction [12], super-resolution [30], tomographic reconstruction [27], inpainting [8], source separation [48], fusion [21], and phase retrieval [6]. The development of new theory, methodology, and algorithms for imaging problems is a focus of significant research efforts. Particularly, convex imaging problems have received a lot of attention lately, leading to major developments in this area.

Most recent works in the imaging literature adopt formal mathematical approaches to analyse problems, derive solutions, and study the underpinning algorithms. There are several mathematical frameworks available to solve imaging problems [22]. In particular, many modern methods are formulated in the Bayesian statistical framework, which relies on statistical models to represent the data observation process and the prior knowledge available, and then derives solutions by using inference techniques rooted in Bayesian decision theory [22].

*LTCI, Telecom ParisTech (alain.durmus@telecom-paristech.fr).

†Centre de Mathématiques Appliquées, UMR 7641, Ecole Polytechnique (eric.moulines@polytechnique.edu).

‡Maxwell Institute for Mathematical Sciences and School of Mathematical and Computer Sciences, Heriot-Watt University (m.pereyra@hw.ac.uk).

There are currently two main approaches in Bayesian imaging methodology. The predominant approach is to use a convex formulation of the estimation problem and postulate a prior distribution that is log-concave. This leads to a posterior distribution that is also log-concave, and where maximum-a-posteriori (MAP) estimation can be computed efficiently by using high dimensional convex optimisation algorithms [17]. In addition to scaling well to large settings, convex optimisation algorithms have two additional advantages that are important for practical Bayesian computation: they are well understood theoretically and their conditions for convergence are clear and simple to check; and the main algorithms are general and can be applied similarly to wide range of problems. However, convex optimisation on its own cannot deliver basic aspects of the Bayesian paradigm and struggles to support the complex statistical analyses that are inherent to modern scientific reasoning and decision-making.

The second main approach in Bayesian imaging methodology is based on stochastic simulation algorithms, namely Markov chain Monte Carlo (MCMC) algorithms. Such methods, which were already actively studied over two decades ago, have regained significant attention lately because of their capacity to address very challenging imaging problems that are beyond the scope of optimisation-based techniques [39]. Additionally to complex models such as hierarchical or empirical Bayesian models, MCMC methods also enable advanced analyses such as hypotheses test and model selection. Unfortunately, despite great progress in high dimensional MCMC methodology, solving imaging problems by stochastic simulation remains too expensive for applications involving moderate or large datasets. Another drawback of existing MCMC methods is that the conditions for their convergence are often significantly more difficult to check than those of optimisation schemes. As a result, most practitioners only assess convergence empirically. It is worth mentioning that some of these limitations can be partially mitigated by resorting to variational Bayes or message passing approximations, which are generally significantly more computationally efficient than stochastic simulation. Unfortunately, such approximations are available only for specific models, and we currently have little theory to analyse the approximation error involved. Similarly, it is generally difficult to provide convergence guarantees for the related algorithms, which often suffer from local convergence issues. Observe that this is in sharp contrast with the convex optimisation approach, which despite its clear limitations, is general and well understood theoretically.

In summary, convex optimisation and MCMC methods have complementary strengths and weaknesses related to their computational efficiency, theoretical underpinning, and the inferences they can support. As a result, it is increasingly acknowledged that the two methodologies should be used together. In this view, the future imaging methodological toolbox should provide a flexible framework where it is possible to perform very efficiently a first analysis of a full dataset by using convex optimisation algorithms, followed by in-depth analyses by MCMC simulation for specific data (e.g., particular data that will be used as evidence to support a hypothesis or a decision). Also, in this framework practitioners should be able to use MCMC algorithms to perform preliminary analyses, which then set the basis for a full scale analysis with convex optimisation techniques. These could be, for example, exploratory analyses with selected data aimed at calibrating the model or performing Bayesian model selection, and benchmarking analyses to assess efficient approximations (e.g., optimisation-based approximate confidence intervals [35]). Unfortunately, it is currently difficult to use optimisation and MCMC methodologies in this complementary manner because optimisation

methods use predominantly non-conjugate priors that are not smooth, such as priors involving the ℓ_1 or the total-variation noms, whereas MCMC methods are mainly restricted to models with priors that are either conjugate to the likelihood function, or that are smooth with Lipchitz gradients (the latter enables efficient high dimensional MCMC algorithms such as the Metropolis-adjusted Langevin algorithm or Hamiltonian Monte Carlo [39]).

Proximal MCMC algorithms, proposed recently in [36], are an important first step towards bridging this methodological gap between convex optimisation and stochastic simulation. Unlike conventional high dimensional MCMC algorithms that use gradient mappings and require Lipchitz differentiability, proximal MCMC algorithms draw their efficiency from convex analysis, namely proximal mappings and Moreau-Yoshida envelopes. This allows MCMC-based Bayesian computation for precisely the type of models that are solved by convex optimisation (i.e., high dimensional models that are log-concave but not smooth), which in turn enables advanced Bayesian analyses for these models (e.g., see [35; 2] for applications of proximal MCMC to Bayesian uncertainty quantification and sparse regression). However, the proximal MCMC algorithms presented in [36] have three shortcomings that limit their impact in imaging sciences, and which this paper seeks to address. First, the conditions that guarantee the convergence of the algorithms are difficult to check in practice. Second, the algorithms assume that it is possible to compute the proximal mapping of the log-posterior distribution; in practice however this mapping is often approximated by using a forward-backward splitting scheme. Third, the algorithms rely on a Metropolis-Hastings correction step to remove the asymptotic bias introduced by the approximations and to guarantee that the Markov chains target the desired posterior distribution. Unfortunately, this correction step can degrade significantly the efficiency of the algorithms (i.e., the asymptotic bias is removed at the expense of a potentially significant increase in estimation variance and some additional bias from the Markov chain's transient or burn-in regime).

This paper presents a new and significantly better proximal MCMC methodology that address all the issues of the original proximal algorithms discussed above. This new methodology is highly computationally efficient and general, in that it can be applied straightforwardly to most models currently addressed by convex optimisation (in particular, to any model that can be solved by forward-backward splitting). Moreover, we provide simple theoretical conditions to guarantee the convergence of the Markov chains, as well as bounds on its convergence rate. To conclude, we emphasise again that our aim is to provide a Bayesian computation methodology that complements rather than competes with modern convex optimisation, particularly by enabling advanced Bayesian analyses for high-dimensional models that are log-concave.

The remainder of the paper is organised as follows: Section 2 defines notation, introduces the class of models considered, and recalls the Langevin MCMC approach that is the basis of our method. In Section 3 we present the proposed MCMC method, analyse its theoretical properties in detail, provide practical implementation guidelines, and discuss connections with the original proximal MCMC algorithms described in [36]. Section 4 illustrates the methodology on four experiments related to image deconvolution and tomographic reconstruction with total-variation and ℓ_1 sparse priors, where we conduct a range of challenging Bayesian analyses related to model comparison and uncertainty quantification. Conclusions and perspectives for future work are reported in Section 5. Proofs are finally reported in Appendices A and C.

sec:bac

2. Bayesian analysis and computation.

2.1. Notations and Conventions. Denote by $\mathcal{B}(\mathbb{R}^d)$ the Borel σ-field of $\mathbb{R}^d$. For all $\mathsf{A} \in \mathcal{B}(\mathbb{R}^d)$, denote by $\mathrm{Vol}(\mathsf{A})$ its Lebesgue measure. Denote by $\mathbb{M}(\mathbb{R}^d)$ the set of all Borel measurable functions on $\mathbb{R}^d$ and for $f \in \mathbb{M}(\mathbb{R}^d)$, $\|f\|_\infty = \sup_{x\in\mathbb{R}^d} |f(x)|$. For μ a probability measure on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ and $f \in \mathbb{M}(\mathbb{R}^d)$ a μ-integrable function, denote by $\mu(f)$ the integral of f w.r.t. μ. For two probability measures μ and ν on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$, the total variation norm of μ and ν is defined as

$$\|\mu - \nu\|_{\mathrm{TV}} = \sup_{f\in\mathbb{M}(\mathbb{R}^d), \|f\|_\infty \le 1} \left| \int_{\mathbb{R}^d} f(x)\mathrm{d}\mu(x) - \int_{\mathbb{R}^d} f(x)\mathrm{d}\nu(x) \right|$$

Let $f : \mathbb{R}^d \to (-\infty, +\infty]$. If f is a Lipschitz function, namely there exists $C \ge 0$ such that for all $x, y \in \mathbb{R}^d$, $|f(x) - f(y)| \le C \|x - y\|$, then denote

$$\|f\|_{\mathrm{Lip}} = \inf\{|f(x) - f(y)| \, \|x - y\|^{-1} \mid x, y \in \mathbb{R}^d, x \ne y\} .$$

f is said to be proper if there exists $x_0 \in \mathbb{R}^d$ such that $f(x_0) < +\infty$. Denote for all $M \in \mathbb{R}$, $\{f \le M\} = \{z \in \mathbb{R}^d \mid f(z) \le M\}$. f is said to be lower semicontinuous if for all $M \in \mathbb{R}$, $\{f \le M\}$ is a closed subset of $\mathbb{R}^d$. For $k \ge 0$, denote by $C^k(\mathbb{R}^d)$, the set of k-times continuously differentiable functions. For $f \in C^1(\mathbb{R}^d)$, denote by ∇f the gradient of f. Denote for all $\mathrm{q} \ge 1$, the ℓ_q norm $\|\cdot\|_\mathrm{q}$ on $\mathbb{R}^d$ by for all $x \in \mathbb{R}^d$, $\|x\|_\mathrm{q} = (\sum_{i=1}^d |x_i|^\mathrm{q})^{1/\mathrm{q}}$. Denote by $\|\cdot\|$ the Euclidian norm on $\mathbb{R}^d$. For all $x \in \mathbb{R}^d$ and $M > 0$, denote by $\mathrm{B}(x, M)$, the ball centered at x of radius M. For a closed convex $\mathcal{K} \subset \mathbb{R}^d$, denote by $\mathrm{proj}_\mathcal{K}(\cdot)$, the projection onto $\mathcal{K}$, and $\iota_\mathcal{K}$ the convex indicator of $\mathcal{K}$ defined by $\iota_\mathcal{K}(x) = 0$ if $x \in \mathcal{K}$, and $\iota_\mathcal{K}(x) = +\infty$ otherwise. In the sequel, we take the convention that $\inf \emptyset = \infty$, $1/\infty = 0$ and for $n, p \in \mathbb{N}$, $n < p$ then $\sum_p^n = 0$ and $\prod_p^n = 1$.

2.2. Imaging inverse problems. We consider inverse problems where we seek to estimate an unknown quantity $x \in \mathbb{R}^d$ from an observation y, related to x by a forward statistical model with likelihood function $p(y|x)$. Following a Bayesian approach, we use prior knowledge about x to reduce the uncertainty and deliver accurate estimation results [22]. Precisely, we specify a prior distribution $p(x)$ promoting expected properties (e.g., sparsity, piecewise regularity, or smoothness), and combine observed and prior information by using Bayes' theorem, leading to the posterior distribution [40]

$$\pi(x) \triangleq p(x|y) = \frac{p(y|x)p(x)}{\int_{\mathbb{R}^d} p(y|x)p(x)\mathrm{d}x} ,$$

that we henceforth denote as π, and which models our knowledge about x after observing y. In this paper we focus on inverse problems that are convex. We assume that π is log-concave, i.e.

posterior (1) $$\pi(x) = \frac{\mathrm{e}^{-U(x)}}{\int_{\mathbb{R}^d} \mathrm{e}^{-U(s)}\mathrm{d}s},$$

for some measurable function $U : \mathbb{R}^d \to (-\infty, +\infty]$ satisfying the following condition.

H1. *$U = f + g$, where $f : \mathbb{R}^d \to \mathbb{R}$ and $g : \mathbb{R}^d \to (-\infty, +\infty]$ are two lower bounded functions satisfying:*

(i) f is convex, continuously differentiable, and gradient Lipschitz with Lipschitz constant L_f, i.e. for all $x, y \in \mathbb{R}^d$

$$\|\nabla f(x) - \nabla f(y)\| \leq L_f \|x - y\| \ . \tag{2}$$

(ii) g is proper, convex and lower semi continuous (l.s.c).

Notice that the class (1) comprises many important models that are used extensively in modern imaging sciences. Particularly, models of the form $U(x) = \|y - Ax\|^2/2\sigma^2 + \phi(Bx)$ for some linear operators A, B, and convex regulariser ϕ that is typically non-smooth, and which may also encode convex constraints on the parameter space. In such cases $f(x) = \|y - Ax\|^2/2\sigma^2$ and $g(x) = \phi(Bx)$ for instance.

When x is high-dimensional, drawing inferences from π directly is generally not possible. Instead we use summaries, particularly point estimators, that capture some of the information about π that is relevant for the application considered [40]. In particular, modern statistical imaging methodology relies strongly on the maximum-a-posteriori (MAP) estimator defined by:

$$\hat{x}_{\text{MAP}} = \arg\max_{x \in \mathbb{R}^d} \pi(x) = \arg\min_{x \in \mathbb{R}^d} U(x) \ , \tag{3}$$

which can often be computed efficiently, even in very large problems, by using proximal convex optimisation algorithms [10; 32]. From the practitioner's viewpoint, this is a main advantage w.r.t. most other summaries that require high-dimensional integration w.r.t. π, which is generally significantly more computationally expensive [39].

However, in its raw form, mathematical imaging based on optimisation struggles to support complex statistical analyses. For example, such methods are typically unable to assess the uncertainty in the solutions delivered and to support uncertainty quantification and decision-making procedures (e.g. hypothesis tests). Similarly, they have difficulty checking and comparing alternative mathematical models intrinsically (i.e., without ground truth available). To perform such advanced (often Bayesian) analyses and deliver the full richness of the statistical paradigm it is necessary to use Monte Carlo stochastic simulation algorithms [17].

As mentioned previously, the high-dimensionality and the lack of smoothness of π pose important challenges from a Bayesian computation viewpoint. This paper presents a new MCMC methodology to tackle this problem. The proposed methodology is general, robust, theoretically sound, and computationally efficient, and can be applied straightforwardly to any model satisfying (1) that can be addressed by using proximal convex optimisation (particularly by using the gradient of f and the proximal operator of g, similarly to forward-backward splitting algorithms).

Finally, we mention at this point some recent works that also consider new MCMC methods to sample from non-smooth posterior distributions with ℓ_1 priors, which is a specific subclass of (1). Most of these works consider Gibbs sampling strategies based either on auxiliary variables [34] or on direct simulation from the univariate conditional densities involved [25; 26]. An alternative strategy is to use a non-linear transformations to change the ℓ_1 prior into a

Gaussian distribution, which then enables using the randomize-then-optimise (RTO) method of [4] to generate samples (see [47] for details). Similarly to our methodology, RTO combines optimisation and sampling steps, albeit in a completely different way (precisely, RTO simulates high-dimensional Gaussian vectors by minimising a loss function with random parameters).

2.3. Bayesian computation: unadjusted and Metropolis-adjusted Langevin algorithms.

The MCMC method proposed in this paper is derived from the discretization of overdamped Langevin diffusions. Let $\bar{U} : \mathbb{R}^d \to \mathbb{R}$ be a continuously differentiable function and consider the Langevin stochastic differential equations (SDE) given by

eq:langevin-1

$$\mathrm{d}\mathbf{X}_t = -\nabla \bar{U}(\mathbf{X}_t)\mathrm{d}t + \sqrt{2}\mathrm{d}B_t^d \,, \tag{4}$$

where $(B_t^d)_{t\geq 0}$ is a d-dimensional Brownian motion. Under additional mild assumptions, this equation has a unique strong solution. In addition if $\int_{\mathbb{R}} \mathrm{e}^{-\bar{U}(x)}\mathrm{d}x < \infty$, then $\bar{\pi}(x) \propto \mathrm{e}^{-\bar{U}(x)}$ is the unique invariant distribution of the semi-group associated with the Langevin SDE, see [23]. Consequently, if we could solve (4) and let $t \to \infty$, this would provide samples from $\bar{\pi}$ useful for Bayesian computation. Since it is possible to analytically solve (4) only in very specific cases, we consider a discrete-time Euler-Maruyama approximation and obtain the following Markov chain $(X_k)_{k\geq 0}$: for all $k \geq 0$

definition-Euler

$$\text{ULA} : \; X_{k+1} = X_k - \gamma \nabla \bar{U}(X_k) + \sqrt{2\gamma} Z_{k+1} \,, \tag{5}$$

where $\gamma > 0$ is a given step size and $(Z_k)_{k\geq 1}$ is a sequence of i.i.d. d-dimensional standard Gaussian random variables. This scheme has been first introduced in molecular dynamics by [15] and [33], and then popularized in artificial intelligence by [18], [19] and in computational statistics by [31] and [42]. Following [42], this algorithm is referred to as the Unadjusted Langevin Algorithm (ULA).

In Bayesian computation, the samples $(X_k)_{k\geq 0}$ generated by ULA (5) are used to estimate probabilities and expectations w.r.t. $\bar{\pi}$. This scheme has attracted significant attention recently, in particular for high-dimensional problems were most Monte Carlo methods struggle. Theory for ULA advanced significantly recently with the development of non-asymptotic bounds in total variation distance between $\bar{\pi}$ and the marginal laws of the Markov chain $(X_k)_{k\geq 0}$ defined by ULA [11; 13], with explicit dependence on the stepsize γ and the dimension d (see Subsection 3.2). These new theoretical results are important because they provide estimation accuracy guarantees for ULA, as well as valuable new insights into the convergence properties of the algorithm. In particular, they establish that if $\bar{U}$ is convex and gradient Lipchitz, then ULA's convergence properties deteriorate at most polynomially as d increases. Remarkably, if in addition $\bar{U}$ is strongly convex, then it deteriorates at most linearly with d, confirming the empirical evidence that ULA is a highly computationally efficient method to sample in high-dimensional settings.

It is worth emphasising at this point that this deep understanding of ULA is very recent. Indeed, without a proper theoretical underpinning, ULA has been traditionally regarded as unreliable and rarely applied directly in statistics or statistical image processing. Instead, most applications reported in the literature adopt a safe approach and complement ULA with a Metropolis-Hastings correction step targeting $\bar{\pi}$, as recommended by [44] and [42]. This

correction guarantees that the resulting Metropolis Adjusted Langevin Algorithm (MALA) generates a reversible Markov chain with respect to $\bar{\pi}$, and therefore eliminates the asymptotic bias. And perhaps more importantly, it places ULA within the sound theoretical framework of Metropolis-Hasting algorithms. For sufficiently smooth densities MALA inherits the good convergence properties of ULA and scales efficiently to high-dimensional settings [42].

Unfortunately, neither ULA nor MALA are well defined for non-smooth target densities, which strongly limits their application to modern mathematical imaging problems. In fact, both theory and experimental evidence show that ULA and MALA often run into difficulties if π is not sufficiently regular. For example, when $\nabla \log \pi$ is not Lipchitz continuous ULA is generally explosive and MALA is not geometrically ergodic (see [42; 36, Figure 2]). Similarly, when $\nabla \log \pi$ is subdifferentiable and therefore, at least from a purely algorithmic viewpoint, the algorithms could still be applied, the theory underpinning the ULA and MALA collapses and even the convergence of the time-continuous Langevin diffusion driving the algorithms becomes unclear. Moreover, many applications involve constraints on the parameter space and then π is supported only on a bounded convex set $\mathcal{K}$. In such case, $\nabla \log \pi$ is bounded on $\mathcal{K}$ and infinite or not defined outside $\mathcal{K}$. Then it is not possible to use ULA, and MALA typically behaves very poorly (the algorithm gets "stuck" whenever the proposal drives the Markov chain outside $\mathcal{K}$). Following a proximal MCMC approach [36], in the following section we present a new ULA that exploits tools from convex calculus and proximal optimisation to address these issues, and sample efficiently from high-dimensional log-concave densities of the form H1 that are beyond the scope of conventional ULAs and MALAs.

nore-yosida-regul

3. Proximal MCMC: Moreau-Yosida regularised Unadjusted Langevin Algorithm.

ssec:MYULAl

3.1. Proposed method. A central idea in this work is to replace the non-smooth potential U with a carefully designed smooth approximation U^λ which, by construction, has the following two key properties: 1) its Euler-Maruyama discrete-time approximations are always stable and have favourable convergence properties, and 2) we can make $\pi^\lambda \propto \mathrm{e}^{-U^\lambda}$ arbitrarily close to π by adjusting an approximation parameter $\lambda > 0$.

In a manner akin to [36], we define such approximations by using Moreau-Yosida envelopes [9] which we recall below. Let $\mathrm{g} : \mathbb{R}^d \to (-\infty, +\infty]$ be a l.s.c convex function and $\lambda > 0$. The λ-Moreau-Yosida envelope of g is a carefully regularised approximation of g given by

–eq:id-MY-env

$$\mathrm{g}^\lambda(x) = \min_{y \in \mathbb{R}^d} \left\{ \mathrm{g}(y) + (2\lambda)^{-1} \|x - y\|^2 \right\} , \tag{6}$$

where λ is a regularisation parameter that controls a trade-off between the regularity properties of g^λ and the approximation error involved. Remarkably, by [43, Example 10.32, Theorem 9.18], the approximation g^λ inherits the convexity of g and is always continuously differentiable, even if g is not. In fact, g^λ is gradient Lipshitz [43, Proposition 12.19]: for all $x, y \in \mathbb{R}^d$,

'moreau'yosida

$$\left\| \nabla \mathrm{g}^\lambda(x) - \nabla \mathrm{g}^\lambda(y) \right\| \leq \lambda^{-1} \|x - y\| . \tag{7}$$

The gradient is given by for all $x \in \mathbb{R}^d$

ition-grad-prox

$$\nabla \mathrm{g}^\lambda(x) = \lambda^{-1} \left(x - \mathrm{prox}_{\mathrm{g}}^\lambda(x) \right) , \tag{8}$$

where

$$\text{prox}_{\text{g}}^{\lambda}(x) = \arg\min_{y\in\mathbb{R}^d} \left\{ \text{g}(y) + (2\lambda)^{-1} \|x-y\|^2 \right\} , \tag{9}$$

is the proximal operator of g [9]. This operator is used extensively in imaging methods based on convex optimisation, where it is generally computed efficiently by using a specialised algorithm [10; 32]. Indeed, similarly to gradient mappings, $\text{prox}_{\text{g}}^{\lambda}$ also moves points in the direction of the minimum of g (by an amount related to the value of λ), and has many properties that are useful for devising fixed-point methods [9].

In addition, g^{λ} envelopes g from below: for all $x \in \mathbb{R}^d$, $\text{g}^{\lambda}(x) \leq \text{g}(x)$, and since for $0 < \lambda < \lambda'$ and $x, y \in \mathbb{R}^d$, $\text{g}(y) + (2\lambda')^{-1}\|x-y\|^2 \leq \text{g}(y) + (2\lambda)^{-1}\|x-y\|^2$, we get that for all $x \in \mathbb{R}^d$ $\text{g}^{\lambda'}(x) \leq \text{g}^{\lambda}(x)$. By [43, Theorem 1.25], g^{λ} converges pointwise to g as λ goes to 0, i.e. for all $x \in \mathbb{R}^d$,

$$\lim_{\lambda\to 0} \text{g}^{\lambda}(x) = \text{g}(x) . \tag{10}$$

Hence, g^{λ} provides a convex and smooth approximation to g that we can make arbitrarily close to g by adjusting the value of λ.

So under **H** 1, if g is not continuously differentiable, but the proximity operator associated with g is available, we can consider sampling algorithms that use the λ-Moreau-Yosida envelope g^{λ} instead of g. Here we propose to replace the potential U with the approximation $U^{\lambda} : \mathbb{R}^d \to \mathbb{R}$ defined for all $x \in \mathbb{R}^d$ by

$$U^{\lambda}(x) = g^{\lambda}(x) + f(x) ,$$

which we will use to define a surrogate target density $\pi^{\lambda} \propto \text{e}^{-U^{\lambda}}$. We will see that such approximation is endowed with very useful regularity and approximation accuracy properties.

Proposition 1 below implies that the probability measure π^{λ} on $\mathbb{R}^d$, with density with respect to the Lebesgue measure, also denoted by π^{λ} and given for all $x \in \mathbb{R}^d$ by

$$\pi^{\lambda}(x) = \frac{\text{e}^{-U^{\lambda}(x)}}{\int_{\mathbb{R}^d} \text{e}^{-U^{\lambda}(s)}\text{d}s} ,$$

is well defined, log-concave, Lipschitz continuously differentiable, and as close to π as required.

H2. *Assume that one of these two conditions holds:*

(i) e^{-g} *is integrable with respect to the Lebesgue measure.*

(ii) g *is Lipschitz.*

Proposition 1. *Assume* ***H***1 *and* ***H***2.

a) For all $\lambda > 0$, π^{λ} *defines a proper density of a probability measure on* $\mathbb{R}^d$, *i.e.*

$$0 < \int_{\mathbb{R}^d} \text{e}^{-U^{\lambda}(y)}\text{d}y < +\infty .$$

b) For all $\lambda > 0$, π^{λ} *is log-concave and continuously differentiable with*

$$\nabla U^{\lambda}(x) = -\nabla \log \pi^{\lambda}(x) = \nabla f(x) + \lambda^{-1}(x - \text{prox}_{g}^{\lambda}(x)) . \tag{11}$$

In addition, ∇U^{λ} *is Lipschitz with constant* $L \leq L_f + \lambda^{-1}$.

c) *The approximation π^λ converges to π as $\lambda \downarrow 0$ in total variation norm, i.e.*

$$\lim_{\lambda \to 0} \|\pi^\lambda - \pi\|_{\mathrm{TV}} = 0 \, .$$

d) *If **H**2-(ii) then for all $\lambda > 0$,*

$$\|\pi^\lambda - \pi\|_{\mathrm{TV}} \leq \lambda \, \|g\|^2_{\mathrm{Lip}} \, .$$

Proof. The proof is postponed to Appendix A. ∎

Figure 1 shows the approximations of two non-smooth densities that satisfy **H**1:

1. the Laplace density $\pi(x) = (1/2)\exp(|x|)$, for which

$$\pi^\lambda(x) = \frac{\exp\left\{(\lambda/2 - |x|)\mathbb{1}_{\{|x| \geq \lambda\}} - (x^2/(2\lambda))\mathbb{1}_{\{|x|<\lambda\}}\right\}}{2\left\{\mathrm{e}^{-\lambda/2} + (2\pi/\lambda)^{1/2}(\mathbf{\Phi}(\lambda^{1/2}) - 1/2)\right\}} \, ,$$

   where $\mathbf{\Phi}$ is the cumulative function of the standard normal distribution.

2. the uniform density $\pi(x) = (1/2)\exp(-\iota_{[-1,1]}(x))$, for which

$$\pi^\lambda(x) = \left\{2 + \sqrt{2\pi\lambda}\right\}^{-1} \exp\left[\left\{-\max(|x| - 1, 0)\right\}^2/(2\lambda)\right] \, .$$

We observe that the approximations are smooth and converge to π as λ decreases, as described by Proposition 1. Also for these two examples, analytic expressions for $\|\pi - \pi^\lambda\|_{\mathrm{TV}}$ can be found, and Figure 2 shows $\|\pi - \pi^\lambda\|_{\mathrm{TV}}$ as a function of $\lambda > 0$. Notice that in the case of the Laplace density $\|\pi - \pi^\lambda\|_{\mathrm{TV}}$ goes to 0 quadratically in λ as λ goes to 0, which is faster than the linear bound given in Proposition 1-d). Also note that this bound does not apply to the uniform density, and in this case $\|\pi - \pi^\lambda\|_{\mathrm{TV}}$ vanishes at rate $\sqrt{\lambda}$.

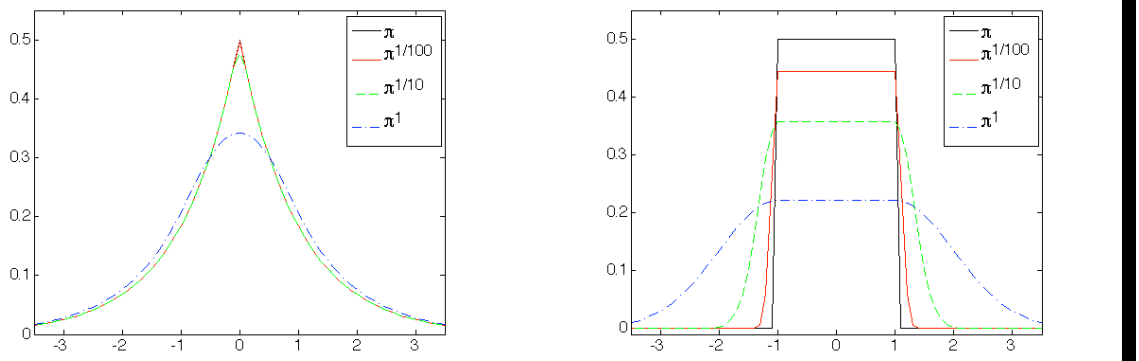

(a) $\pi(x) = \frac{1}{2}\mathrm{e}^{-|x|}$ (b) $\pi(x) = \frac{1}{2}\iota_{[-1,1]}(x)$

Figure 1. *Density plots for the Laplace (a) and uniform (b) distributions (solid black), and their smooth approximations π^λ for $\lambda = 1, 0.1, 0.01$ (dashed blue and green, and solid red).*

We now make two key observations. First, Proposition 1 shows that ∇U^λ is gradient Lipschitz and therefore it guarantees that the Langevin SDE constructed with U^λ converges

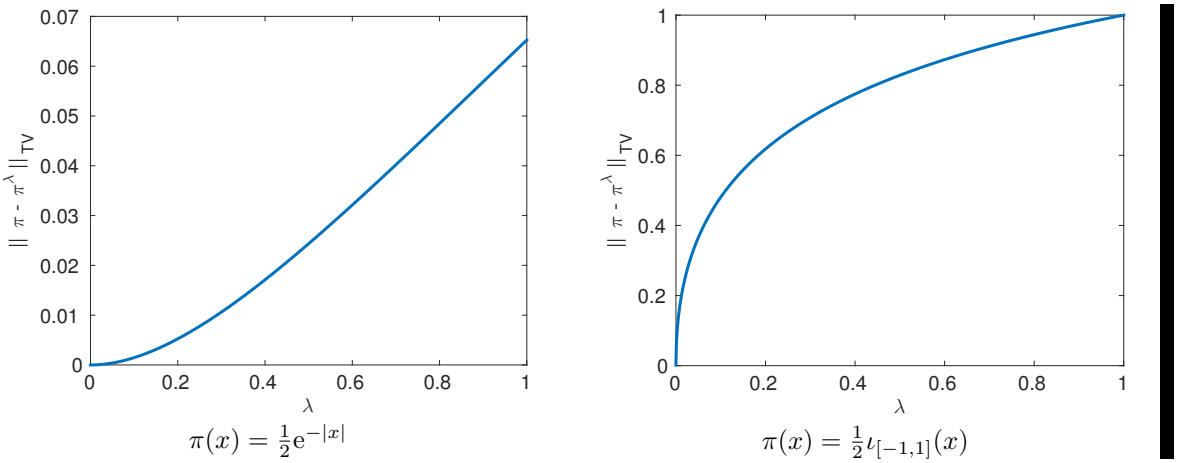

Figure 2. *Total variation norm between π and its smooth approximation π^λ as function of λ.* FigMoreauApprox-TV

to π^λ as $t \to \infty$ (formally, it guarantees that the Langevin SDE associated with π^λ admits a unique strong solution $(\mathbf{X}_t^\lambda)_{t\geq 0}$ and π^λ is the unique stationary distribution of the semigroup). More importantly, as it will be seen below, it implies that the ULA chain derived from a Euler-Maruyama discretisation of this Langevin diffusion will be, by construction, well behaved and useful for Monte Carlo integration with respect to π^λ.

Second, Proposition 1 also establishes that λ controls the estimation bias involved in performing estimations with π^λ as a substitute of π. This approximation error can be made arbitrarily small, and is bounded explicitly by $\lambda \|g\|_{\mathrm{Lip}}^2$ when g is Lipschitz.

We are now in a position to present the new MCMC methodology proposed in this work, which is essentially an application of ULA to π^λ. Precisely, given $\lambda > 0$ and a stepsize $\gamma > 0$, we use an Euler-Maruyama approximation of $(\mathbf{X}_t^\lambda)_{t\geq 0}$, and obtain the following Markov chain $(X_k^{\mathrm{M}})_{k\geq 0}$: for all $k \geq 0$

q:def-MYRULA

$$\text{MYULA}: \; X_{k+1}^{\mathrm{M}} = (1 - \tfrac{\gamma}{\lambda}) X_k^{\mathrm{M}} - \gamma \nabla f(X_k^{\mathrm{M}}) + \tfrac{\gamma}{\lambda} \operatorname{prox}_g^\lambda(X_k^{\mathrm{M}}) + \sqrt{2\gamma} Z_{k+1} \; , \tag{12}$$

where $\{Z_k, \; k \in \mathbb{N}^*\}$ is a sequence of i.i.d. d dimensional standard Gaussian random variables. This algorithm will be referred to as the *Moreau-Yosida Unadjusted Langevin Algorithm* (MYULA), and is summarised in Algorithm 1 below (see Subsection 3.3 for guidelines for setting the values of γ and λ). Note that the stationary distribution of the MYULA sequence $\{X_k^{\mathrm{M}}, \; k \in \mathbb{N}\}$ is different from the target distribution π^λ, and depends on the stepsize $\gamma > 0$. Nevertheless, we show in Subsection 3.2 that, choosing λ and γ appropriately, the samples are very close to π.

Besides, to compute the expectation of a function $h : \mathbb{R}^d \to \mathbb{R}$ under π from $\{X_k^{\mathrm{M}} \; ; \; 0 \leq k \leq n\}$, an optional importance sampling step might be used to correct the regularization. This step amounts to approximate $\int_{\mathbb{R}^d} h(x)\pi(x)\mathrm{d}x$ by the weighted sum

rtance˙sampling

$$\mathrm{S}_n(h) = \sum_{k=0}^{n} \omega_{k,n} h(X_k) \; , \; \text{with } \omega_{k,n} = \left\{ \sum_{\ell=0}^{n} \mathrm{e}^{\bar{g}^\lambda(X_\ell^{\mathrm{M}})} \right\}^{-1} \mathrm{e}^{\bar{g}^\lambda(X_k^{\mathrm{M}})} \; , \tag{13}$$

where for all $x \in \mathbb{R}^d$

$$\bar{g}^\lambda(x) = g^\lambda(x) - g(x) = g(\mathrm{prox}_g^\lambda(x)) - g(x) + (2\lambda)^{-1} \left\| x - \mathrm{prox}_g^\lambda(x) \right\|^2 .$$

To remove this asymptotic bias, we can add an Hastings-Metropolis step, which will produce a Markov chain $\{\tilde{X}_k^\lambda,\ k \in \mathbb{N}\}$ which is reversible this time with respect to π^λ and use similarly an importance sampling step to correct for the bias introduced by smoothing. This algorithm will be called the *Moreau-Yosida Regularized Metropolis-adjusted Langevin Algorithm* (MYMALA).

The focus of this work is on MYULA without importance sampling or Metropolis-Hastings correction. A study of MYMALA is currently in progress and will be reported separately.

Algo:MYULA

Algorithm 1 Moreau-Yoshida unadjusted Langevin algorithm (MYULA)

set $X_0^{\mathrm{M}} \in \mathbb{R}^d$, $\lambda > 0$, $\gamma \in (0, \lambda/(\lambda L_f + 1)]$, $n \in \mathbb{N}$
for $k = 0 : n$ **do**
  $Z_{k+1} \sim \mathcal{N}(0, \mathbb{I}_d)$
  $X_{k+1}^{\mathrm{M}} = (1 - \frac{\gamma}{\lambda}) X_k^{\mathrm{M}} - \gamma \nabla f(X_k^{\mathrm{M}}) + \frac{\gamma}{\lambda} \mathrm{prox}_g^\lambda(X_k^{\mathrm{M}}) + \sqrt{2\gamma} Z_{k+1}$
end for

For illustration, Figure 3 shows the sample approximations of the univariate Laplace and uniform distributions of Figure 1 (the true densities are depicted in solid blue for comparison). The histograms were generated using 10^4 iterations of MYULA with parameters $\lambda = 10^{-3}$ and $\delta = 2\lambda$. Observe that the samples provide a good approximation of the desired target densities, particularly of the uniform distribution which is beyond the scope of the conventional ULA. Also observe that the approximation of the uniform distribution has Gaussian tails, as per the Lipchitz differentiability of π^λ approximation (see Proposition 1 and Figure 1). From Proposition 1, this error can be made arbitrarily low by adjusting the value of λ.

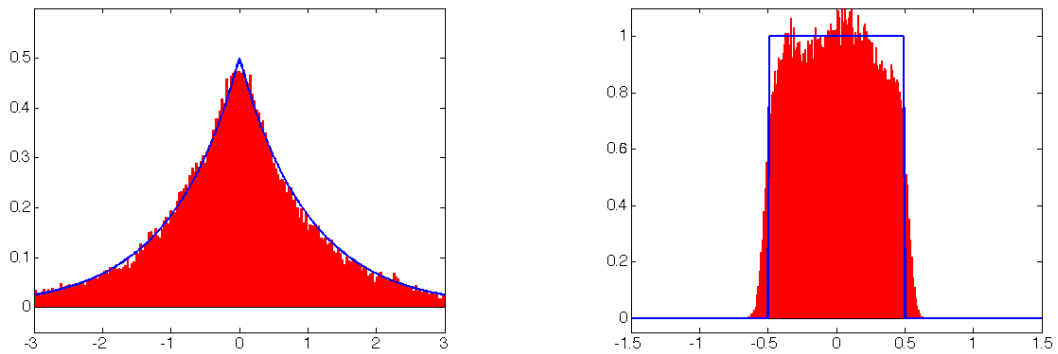

(a) $\pi(x) = \frac{1}{2}\mathrm{e}^{-|x|}$ (b) $\pi(x) = \frac{1}{2}\iota_{[-1,1]}(x)$

Figure 3. *MYULA sample approximations of the Laplace (a) and uniform (b) distributions (true density in solid blue). Histograms computed with* 10^4 *samples generated using* $\lambda = 10^{-3}$ *and* $\delta = 2\lambda$.

FigMYULAApprox

Finally, similarly to ULA, it is possible to adapt MYULA to use a stochastic gradient

strategy based on an unbiased estimator of ∇f (see [46] for details). This can be useful in applications that involve very large datasets for which computing the exact gradient would be too computationally expensive, for example machine learning applications. The specialisation of MYULA to such problems is currently under investigation and will be reported separately.

vergence˙analysis

3.2. Theoretical convergence analysis of MYULA. In this section we present a detailed theoretical analysis of MYULA implemented with fixed regularization parameter $\lambda > 0$ and step-size $\gamma > 0$. We first establish that the chains generated by MYULA converge geometrically fast to an approximation of π that is controlled by λ and γ, and which can be made arbitrarily close to π. More importantly, we also establish non-asymptotic bounds for the estimation error of MYULA with a finite number of iterations. This enables an analysis of the behaviour of MYULA as the dimensionality of the model increases, as well as deriving practical guidelines for setting λ and γ for specific models.

First, under **H** 1, it has been observed that g^λ is λ^{-1}-gradient Lipschitz, which implies that U^λ is gradient Lipschitz as well: there exists $L \geq 0$ such that for all $x, y \in \mathbb{R}^d$, $\left\|\nabla U^\lambda(x) - \nabla U^\lambda(y)\right\| \leq L \left\|x - y\right\|$ and

st˙lip˙u˙lambda

$$L \leq L_f + \lambda^{-1} \,. \tag{14}$$

Of course, this bound strongly depends on the decomposition of U in a smooth and a non-smooth part, which is arbitrary and therefore can be pessimistic (for instance, if U is continuously differentiable, g can be chosen to be 0 which implies $U^\lambda = U$ and $L = L_f$).

We assume first the following assumption on the potential U^λ.

ssum:potentialU2

H3. *There exist a minimizer $x^\star$ of U^λ, $\eta_{\mathrm{c}} > 0$ and $\mathrm{R}_{\mathrm{c}} \geq 0$ such that for all $x \in \mathbb{R}^d$, $\|x - x^\star\| \geq \mathrm{R}_{\mathrm{c}}$,*

rexpo˙potential

$$U^\lambda(x) - U^\lambda(x^\star) \geq \eta_{\mathrm{c}} \left\|x - x^\star\right\| \,. \tag{15}$$

Note that in fact **H**3 always holds under **H**1 and **H**2, since by Lemma 4 and Proposition 1 there exist $C_1, C_2 > 0$ such that $U^\lambda(x) \geq C_1 \left\|x\right\| - C_2$. Therefore, since U^λ is continuous on $\mathbb{R}^d$, there exists a minimizer $x^\star$ of U^λ and (15) holds with $\eta_{\mathrm{c}} \leftarrow C_1/2$ and $\mathrm{R}_{\mathrm{c}} \leftarrow 2(C_2 + \|x^\star\| + U^\lambda(x^\star))/C_1$. However, these constants are non quantitative, and that is why we introduce **H** 3 to derive quantitative bounds.

Consider the Markov kernel R_γ associated to the Euler-Maruyama discretization (12) given, for all $\mathsf{A} \in \mathcal{B}(\mathbb{R}^d)$ and $x \in \mathbb{R}^d$ by

finition˙R˙kernel

$$R_\gamma(x, \mathsf{A}) = (4\pi\gamma)^{-d/2} \int_{\mathsf{A}} \exp\left(-(4\gamma)^{-1} \left\|y - x + \gamma \nabla U^\lambda(x)\right\|^2\right) \mathrm{d}y \,. \tag{16}$$

The sequence $(X_n^{\mathrm{M}})_{n \geq 0}$ defined by (12) is a homogeneous Markov chain associated with the Markov kernel R_γ. Therefore for all $n \in \mathbb{N}$, $n \geq 1$, and $x \in \mathbb{R}^d$, the distribution of X_n^{M} started at x is $R_\gamma^n(x, \cdot)$ defined by induction for all $\mathsf{A} \in \mathcal{B}(\mathbb{R}^d)$ by

$$R_\gamma^n(x, \mathsf{A}) = \int_{\mathsf{A}} R_\gamma^{n-1}(x, \mathrm{d}y) R_\gamma(y, \mathsf{A}) \,.$$

It is easily seen that under **H**1, since U^λ is continuously differentiable, R_γ is irreducible with respect to the Lebesgue measure, all compact sets are 1-small and the kernel is strongly aperiodic. In addition under **H**3, since U is also convex then [13, Proposition 13] shows that R_γ satisfies a Foster-Lyapunov drift condition, i.e. for all $\bar\gamma \in (0, L]$, $\gamma \in (0, \bar\gamma]$ and for all $x \in \mathbb{R}^d$,

$$R_\gamma V_{\mathrm{c}}(x) \leq \varrho_{\mathrm{c}}^\gamma V_{\mathrm{c}}(x) + b_{\mathrm{c}}\gamma \;,$$

where

$$V_{\mathrm{c}}(x) = \exp\left\{(\eta_{\mathrm{c}}/4)\left(\|x - x^\star\|^2 + 1\right)^{1/2}\right\} \tag{17a}$$

$$\varrho_{\mathrm{c}} = \mathrm{e}^{-2^{-4}\eta_{\mathrm{c}}^2(2^{1/2}-1)} \;,\; \mathrm{a}_{\mathrm{c}} = \max(1, 2d/\eta_{\mathrm{c}}, \mathrm{R}_{\mathrm{c}}) \tag{17b}$$

$$b_{\mathrm{c}} = \{(\eta_{\mathrm{c}}/4)(d + (\eta_{\mathrm{c}}\bar\gamma/4)) - \log(\varrho_{\mathrm{c}})\}\, \mathrm{e}^{\eta_{\mathrm{c}}(\mathrm{a}_{\mathrm{c}}^2+1)^{1/2}/4 + (\eta_{\mathrm{c}}\bar\gamma/4)(d+(\eta_{\mathrm{c}}\bar\gamma/4))} \;. \tag{17c}$$

eq:convex_drift

By [29, Theorem 16.0.1], R_γ has a unique invariant distribution π_γ^λ and is V_{c}-uniformly geometrically ergodic: there exists $\kappa_{\mathrm{c}} \in (0,1)$ and $C_{\mathrm{c}} \geq 0$ such that all $n \geq 0$ and $x \in \mathbb{R}^d$,

$$\|R_\gamma^n(x,\cdot) - \pi_\gamma^\lambda\|_{\mathrm{TV}} \leq C_{\mathrm{c}} V_{\mathrm{c}}(x) \kappa_{\mathrm{c}}^n \;.$$

Note π_γ^λ is different from π^λ, nevertheless the following result shows that choosing γ small enough, the ULA generates samples very close to the distribution π^λ.

We are now ready to present our main theoretical result: a non-asymptotic bound of the total-variation distance between π and the marginal laws of the samples generated by MYULA. Denote in the following by $\omega : \mathbb{R}_+ \to \mathbb{R}_+$ the function given for all $r \geq 0$ by

–eq:Fsmall

$$\omega(r) = r^2 / \left\{2\boldsymbol{\Phi}^{-1}(3/4)\right\}^2 \;. \tag{18}$$

c-stepsize-convV

Theorem 2 ([13, Corollary 19]). *Assume* ***H***1 *and* ***H***3. *Let* $\bar\gamma \in \left(0, L^{-1}\right]$. *For all* $\varepsilon > 0$ *and* $x \in \mathbb{R}^d$, *we have*

$$\|R_\gamma^n(x,\cdot) - \pi\|_{\mathrm{TV}} \leq \varepsilon \;,$$

provided that $n > T\gamma^{-1}$ *with*

$$T = \max\left\{32\,\eta_{\mathrm{c}}^{-2}\log\left(8\varepsilon^{-1}A_1(x)\right), \log(16\varepsilon^{-1})\Big/(-\log(\kappa))\right\}$$

$$\gamma \leq \frac{-d + \sqrt{d^2 + (2/3)A_2(x)\varepsilon^2(L^2T)^{-1}}}{2A_2(x)/3} \wedge \bar\gamma \;,$$

where $\alpha_{\mathrm{c}} = \max(1, 4d/\eta_{\mathrm{c}}, \mathrm{R}_{\mathrm{c}})$

$$\beta_{\mathrm{c}} = (\eta_{\mathrm{c}}/4)\left[\eta_{\mathrm{c}}\alpha_{\mathrm{c}}/4 + d\right] \max\left\{1, (\alpha_{\mathrm{c}}^2+1)^{-1/2} \exp(\eta_{\mathrm{c}}(\alpha_{\mathrm{c}}^2+1)^{1/2}/4)\right\}$$

$$A_1(x) = (1/2)(V_{\mathrm{c}}(x) + b_{\mathrm{c}}(-\varrho_{\mathrm{c}}^{\gamma}\log(\varrho_{\mathrm{c}}))^{-1} + 8\eta_{\mathrm{c}}^{-2}\beta_{\mathrm{c}}) + 16\eta_{\mathrm{c}}^{-2}\beta_{\mathrm{c}}\mathrm{e}^{32^{-1}\eta_{\mathrm{c}}^2\omega\left\{(8/\eta_{\mathrm{c}})\log(32\eta_{\mathrm{c}}^{-2}\beta_{\mathrm{c}})\right\}}$$

$$A_2(x) = L^2\left(4\eta_{\mathrm{c}}^{-1}\left[1 + \log\left\{V_{\mathrm{c}}(x) + b_{\mathrm{c}}(-\varrho_{\mathrm{c}}^{\gamma}\log(\varrho_{\mathrm{c}}))^{-1}\right\}\right]\right)^2$$

$$\log(\kappa) = -\log(2)(\eta_{\mathrm{c}}^2/32)\left[\log\left\{8\eta_{\mathrm{c}}^{-2}\beta_{\mathrm{c}}\left(3 + 4\eta_{\mathrm{c}}^{-2}\mathrm{e}^{32^{-1}\eta_{\mathrm{c}}^2\omega\left\{(8/\eta_{\mathrm{c}})\log(32\eta_{\mathrm{c}}^{-2}\beta_{\mathrm{c}})\right\}}\right)\right\} + \log(2)\right]^{-1} ,$$

$\mathrm{a}_{\mathrm{c}}, \varrho_{\mathrm{c}}, b_{\mathrm{c}}, V_{\mathrm{c}}$ *are defined in* (17) *and* ω *in* (18).

Proof. The proof follows from combining [13, Lemma 4, Theorem 14, Theorem 16]. ∎

This result implies that the number of iteration to reach a precision target ε is, at worse, of order $d^5\log^2(\varepsilon^{-1})\varepsilon^{-2}$ for this class of models. Significantly more precise bounds can be obtained under more stringent assumption on U^{λ}. In particular, we consider the case where U^{λ} is strongly convex outside some ball; see [14].

utsideBallDrift

H4. *There exist* $\mathrm{R}_{\mathrm{s}} \geq 1$ *and* $m > 0$, *such that for all* $x, y \in \mathbb{R}^d$, $\|x - y\| \geq \mathrm{R}_{\mathrm{s}}$,

$$\left\langle \nabla U^{\lambda}(x) - \nabla U^{\lambda}(y), x - y \right\rangle \geq m\|x - y\|^2 .$$

Of course, in the case where f is strongly convex then this assumption holds.

dec-stepsize-St

Theorem 3 ([13, Lemma 4, Theorem 21]). *Assume* ***H***1 *and* ***H***4. *Let* $\bar{\gamma} \in \left(0, L^{-1}\right]$. *Then for all* $\varepsilon > 0$, *we get* $\|R_{\gamma}^n(x, \cdot) - \pi\|_{\mathrm{TV}} \leq \varepsilon$ *provided that* $n > T\gamma^{-1}$ *with*

$$T = \left(\log\{A_1(x)\} - \log(\varepsilon/2)\right) \Big/ \left(-\log(\kappa)\right)$$

$$\gamma \leq \frac{-d + \sqrt{d^2 + (2/3)A_2(x)\varepsilon^2(L^2T)^{-1}}}{2A_2(x)/3} \wedge \bar{\gamma} ,$$

where

$$A_1(x) = 5 + \left(d/m + \mathrm{R}_{\mathrm{s}}^2\right)^{1/2} + (A_1(x)/L^2)^{1/2}$$

$$A_2(x) = L^2\left(\|x - x^{\star}\|^2 + 2(d + m\mathrm{R}_{\mathrm{s}}^2)(\mathrm{e}^{-\gamma(2m+\bar{\gamma}L^2)}/(2m + \bar{\gamma}L^2))^{-1}\right)$$

$$\log(\kappa) = -(\log(2)m/2)\left[\log\left\{\left(1 + \mathrm{e}^{m\omega\{\max(1,\mathrm{R}_{\mathrm{s}})\}/4}\right)(1 + \max(1, \mathrm{R}_{\mathrm{s}}))\right\} + \log(2)\right]^{-1} ,$$

and ω *is given in* (18).

This result implies that the worst minimal number of iterations to achieve a precision level $\varepsilon > 0$ is this time of order $d\log(d)\log^2(\varepsilon^{-1})\varepsilon^{-2}$.

guidelines

3.3. Selection of λ and γ. We now discuss practical guidelines for setting the values for λ and for γ. As mentioned previously, our aim is to provide an efficient computation methodology that can be applied straightforwardly to any model satisfying **H** 1. Hence,

rather than seeking optimal values for specific models, we focus on general rules that are simple, robust, and which only involve tractable quantities such as Lipschitz constants.

First, by Theorem 2, γ should take its value in the range $\gamma \in (0, \lambda/(L_f\lambda+1)]$ to guarantee the stability of the Euler-Maruyama discretisation, and where we recall that L_f is the Lipschitz constant of ∇f. The values of γ within this range are subject to the a bias-variance trade-off. Precisely, large values of γ produce a fast-moving chain that convergences quickly and has low estimation variance, but potentially relatively high asymptotic bias. Conversely, small values of γ lead to low asymptotic bias, but produce a Markov chain that moves slowly and requires a large number of iterations to produce a stable estimate (such chains often also suffer from some additional bias from the transient or burn-in period). Because applications in imaging sciences involve high dimensionality and require moderately low computing times, as a general rule we recommend setting γ to a relatively large value. For example, in our experiments we use

$$\gamma \in [\lambda/5(L_f\lambda + 1), \lambda/2(L_f\lambda + 1)] \ .$$

Observe that this range depends on the value of λ, which is also subject to a bias-variance tradeoff. Letting $\lambda \to 0$ to bring π^λ close to π reduces asymptotic bias, but forces $\gamma \to 0$ and consequently reduces significantly the efficiency of the chain. Conversely, increasing the value of λ accelerates the chain at the expense of some asymptotic bias. Based on our experience, and again with an emphasis on efficiency in high dimensional settings, we recommend using values of λ in the order of L_f^{-1} (there is no benefit in using larger values of λ because γ saturates at L_f^{-1}). In all our experiments we use $\lambda = 1/L_f$ and $\gamma \in [L_f^{-1}/10, L_f^{-1}/4]$ and obtain estimation errors of the order of 1%.

3.4. Connections to the proximal Metropolis-adjusted Langevin algorithm. We conclude this section with a discussion of the connections between the proposed MYULA method and the original proximal Metropolis-adjusted Langevin algorithm (Px-MALA) [36]. That algorithm is also based on a Euler-Maruyama approximation of a Langevin SDE targeting a Moreau-Yoshide-type regularised approximation of π. However, unlike MYULA, that algorithm uses this approximation as proposal mechanism to drive a Metropolis-Hastings (MH) algorithm targeting π (not the regularised approximation). The role of the MH is two-fold: it removes the asymptotic bias related to the approximations involved, and it provides a theoretical framework for Px-MALA by placing the scheme within the framework of MH algorithms (recall that many theoretical results regarding ULAs are very recent). However, as mentioned previously, the introduction of the MH step often slows down the algorithm, thus leading to higher estimation variance and longer chains (and potentially some bias from the chain's initial transient regime). Of course, it also introduces a significant computational overhead related to the computation of the MH acceptance ratio [36]. Another importance difference between MYULA and Px-MALA is that the latter uses the proximal operator of U, which is often unavailable and has to be approximated by using a forward-backward scheme based on the decomposition $U = f + g$ that we also use in this paper. This approximation error is corrected in practice by the MH step, but it is not considered in the theoretical analysis of the algorithm. Conversely, in MYULA this decomposition is explicit, both in the computational aspects of the method as well as in its theoretical analysis. Furthermore, the theory for MYULA presented in this paper is significantly more complete than that currently

available for Px-MALA and other MALAs. Finally, MYULA is also more robust and simple to implement than Px-MALA. For example, identifying suitable values of γ for MYULA is straightforward by using the guidelines described above, whereas setting γ for Px-MALA can be challenging and often requires using an adaptive MCMC approach based on a stochastic approximation scheme [36; 17].

sec:experiments

4. Experimental results.

In this section we illustrate the proposed methodology with four canonical imaging inverse problems related to image deconvolution and tomographic reconstruction with total-variation and ℓ_1 sparse priors. In the Bayesian setting these problems are typically solved by MAP estimation, which delivers accurate solutions and can be computed very efficiently by using proximal convex optimisation algorithm. Here we demonstrate MYULA by performing some advanced and challenging Bayesian analyses that are beyond the scope of optimisation-based mathematical imaging methodologies. For example, in Section 4.1 we report two experiments where we use MYULA to perform Bayesian model choice for image deconvolution models, and where a novelty is that comparisons are performed intrinsically (i.e., without ground truth available) by computing the posterior probability of each model given the observed data. Following on from this, in Section 4.2 we report the two additional experiments where we use MYULA to explore the posterior uncertainty about x and analyse specific aspects about the solutions delivered, particularly by computing simultaneous credible sets (joint Bayesian confidence sets).

Moreover, to assess the computational efficiency and the accuracy of MYULA we benchmark our estimations against the results of Px-MALA [36] targeting the exact posterior $\pi(x) = p(x|y)$ (recall that this algorithm has no asymptotic estimation bias). We emphasise at this point that we do not seek to compare explicitly and quantitatively the methods because: 1) MYULA and Px-MALA do not target the exact same stationary distribution; 2) high-dimensional quantitative efficiency comparisons may depend strongly on the summary statistics used to define the efficiency metrics; and 3) results can often be marginally improved by fine tuning the algorithm parameters (e.g., step sizes, burn-in periods, etc.). What our comparisons seek to demonstrate is that MYULA can deliver reliable approximate inferences with a computational cost that is often significantly lower than Px-MALA, and more importantly, that it provides a general, robust, and theoretically sound computational framework for performing advanced Bayesian analyses for imaging problems. In all experiments the Lipchitz constant L_f required to set the value γ was computed explicitely. Experiments were conducted on a Apple Macbook Pro computer running MATLAB 2015.

exp:BMS

4.1. Bayesian model selection.

4.1.1. Bayesian analysis and computation.

Most mathematical imaging problems can be solved with a range of alternative models. Currently, the predominant approach to select the best model for a specific problem is to compare their estimations against ground truth. For example, given K alternative Bayesian models $\mathcal{M}_1, \ldots, \mathcal{M}_K$, practitioners often benchmark models by artificially degrading a set of test images, computing the MAP estimator for each model and image, and then measuring estimation error with respect to the truth. The model with the best overall performance is then used in applications to analyse real data. Of course this approach to model selection has some limitations: 1) it relies strongly on test data that

may not be representative of the unknown, and 2) conclusions can depend on the estimation error metrics used.

An advantage of formulating inverse problems within the Bayesian framework is that, in addition to strategies to perform point estimation, this formalism also provides theory to compare models objectively and intrinsically, and hence perform model selection in the absence of ground truth. Precisely, K alternative Bayesian models are compared through their marginal posterior probabilities

–margPost

$$p(\mathcal{M}_j|y) = \frac{p(y|\mathcal{M}_j)K^{-1}}{\sum_{k=1}^{K} p(y|\mathcal{M}_k)K^{-1}}, \quad j = \{1, \dots, K\}, \tag{19}$$

where for objectiveness here we use an uniform prior on the auxiliary variable j indexing the models, $p(y|\mathcal{M}_j)$ is the marginal likelihood

–margLike

$$p(y|\mathcal{M}_j) = \int p(x, y|\mathcal{M}_j)\mathrm{d}x, \quad j = \{1, \dots, K\}, \tag{20}$$

measuring model-fit-to-data and $p(y, x|\mathcal{M}_j)$ is the joint probability density associated with $\mathcal{M}_j$ (see Appendix B for details regarding the case of improper priors). Following Bayesian decision theory, to perform model selection we simply chose the model with the highest posterior probability (this is equivalent to performing MAP estimation on the model index j):

$$\mathcal{M}^* = \underset{j\in\{1,\dots,K\}}{\arg\max}\ p(\mathcal{M}_j|y).$$

From a computation viewpoint, performing Bayesian model selection for imaging problems is challenging because it requires evaluating the likelihoods $p(y|\mathcal{M}_j)$ up to a proportionality constant, or equivalently the Bayes factors $p(y|\mathcal{M}_j)/p(y|\mathcal{M}_i)$ for $i, j \in \{1, \cdots, K\}$ (see Appendix C.2 for details regarding the case of improper priors). Here we perform this computation by Monte Carlo integration. Precisely, given n samples $X_1^M, \dots, X_n^M$ from $p(x|y, \mathcal{M}_j)$, we approximate the marginal likelihood of model $\mathcal{M}_j$ by using the truncated harmonic mean estimator [41]

monicEstimator

$$p(y|\mathcal{M}_j) \approx \left(\sum_{k=1}^{n} \frac{\mathbb{1}_{\mathsf{A}^\star}(X_k^M)}{p(X_k^M, y|\mathcal{M}_j)} \right)^{-1} \mathrm{Vol}(\mathsf{A}^\star), \quad j = \{1, 2, K\} \tag{21}$$

where for all x, y, $p(x, y|\mathcal{M}_j)$ is joint density of $\mathcal{M}_j$ and $\mathsf{A}^\star = \cup_{j=1}^{K} \mathcal{C}_{j,\alpha}^\star$ is the union of highest posterior density regions (24) of each model at level $(1-\alpha)$ (see Section 4.2 for details about HPD regions). In our experiments we use the samples to calibrate each $\mathcal{C}_{j,\alpha}^\star$ for $\alpha = 0.8$. Notice that it is not necessary to compute $\mathrm{Vol}(\mathsf{A}^\star)$ to calculate (21) because the normalisation is retrieved via $\sum_{j=1}^{K} p(\mathcal{M}_j|y) = 1$. See Appendix C for more details about this estimator and its use to compute the Bayes factors.

ssec:exp1

4.1.2. Experiment 1: Image deconvolution with total-variation prior.

Experiment setup. To illustrate the Bayesian model selection approach we consider an image deconvolution problem with three alternative models related to three different blur operators. The goal of image deconvolution is to recover a high-resolution image $x \in \mathbb{R}^n$ from a blurred and noisy observation $y = Hx + w$, where H is a circulant blurring matrix and $w \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I}_n)$. This inverse problem is ill-conditioned, a difficulty that Bayesian image deconvolution methods address by exploiting the prior knowledge available. For this first experiment we consider three alternative models involving three different blur operators H_1, H_2, and H_3. With regards to the prior, we use the popular total-variation prior that promotes regularity by using the pseudo-norm $TV(x) = \|\nabla_d x\|_{1-2}$, where $\|\cdot\|_{1-2}$ is the composite $\ell_1 - \ell_2$ norm and ∇_d is the two-dimensional discrete gradient operator. The posterior distribution $p(x|y)$ for the models is given by

deconvolutionTV

$$\mathcal{M}_j : \quad \pi(x) \propto \exp\left[-(\|y - H_j x\|^2/2\sigma^2) - \beta TV(x)\right] \tag{22}$$

with fixed hyper-parameters $\sigma > 0$ and $\beta > 0$ set manually by an expert. This density is log-concave and MAP estimation can be performed efficiently by proximal convex optimisation.

Figure 4 presents an experiment with the `Boat` test image of size $d = 256 \times 256$ pixels. Figure 4(a) shows a blurred and noisy observation y, generated by using a 5×5 uniform blur and Gaussian noise with $\sigma = 0.47$, related to a blurred signal-to-noise ratio of 40dB. Moreover, Figures 4(b)-(d) show the MAP estimates associated with three alternative instances of model (22) involving the following blur operators:

- $\mathcal{M}_1$: H_1 is the correct 5×5 uniform blur operator.
- $\mathcal{M}_2$: H_2 is a mildly misspecified 6×6 uniform blur operator.
- $\mathcal{M}_3$: H_3 is a strongly misspecified 7×7 uniform blur operator.

(All models share the same hyper-parameter values $\sigma = 0.47$ and $\beta = 0.03$ selected manually to produce good image deconvolution results.) We observe in Figure 4 that models $\mathcal{M}_1$ and $\mathcal{M}_2$ have produced sharp images with fine detail, whereas $\mathcal{M}_3$ is clearly misspecified. In terms of estimation performance with respect to the truth, as expected the estimate of Figure 4(c) corresponding to model $\mathcal{M}_1$ achieves the highest peak signal-to-noise-ratio (PSNR) of 33.8dB, $\mathcal{M}_2$ scores 33.4dB, and $\mathcal{M}_3$ scores 13.4dB. Finally, computing the estimates displayed in Figure 4 with a forward-backward optimisation algorithm [17], which is algorithmically similar to MYULA, required approximately 1 000 iterations and 30 seconds per model[1].

Model selection in the absence of ground truth. We now demonstrate the Bayesian approach to perform model selection intrinsically. Precisely, we ran 10^5 iterations of MYULA with the specific blur operators corresponding to $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$. For this experiment we implemented MYULA with $f(x) = \|y - H_j x\|^2/2\sigma^2$ and $g(x) = \beta TV(x)$, with fixed algorithm parameters $\lambda = L_f^{-1} = 0.45$ and $\gamma = L_f^{-1}/5 = 0.1$, and by using Chambolle's algorithm [7] to evaluate the proximal operator of the TV-norm. Computing these samples required approximately 30 minutes per model. Following on from this, we used the samples to calibrate the high-posterior-density regions $\mathcal{C}_j^\star$ of each model at level 20%, and then computed the Bayes factors between the models by using (21) (see C.1 for details).

[1]The computation of the MAP estimates with the SALSA convex optimisation algorithm [1], which is faster than the forward-backward splitting algorithm, required 2 seconds per model.

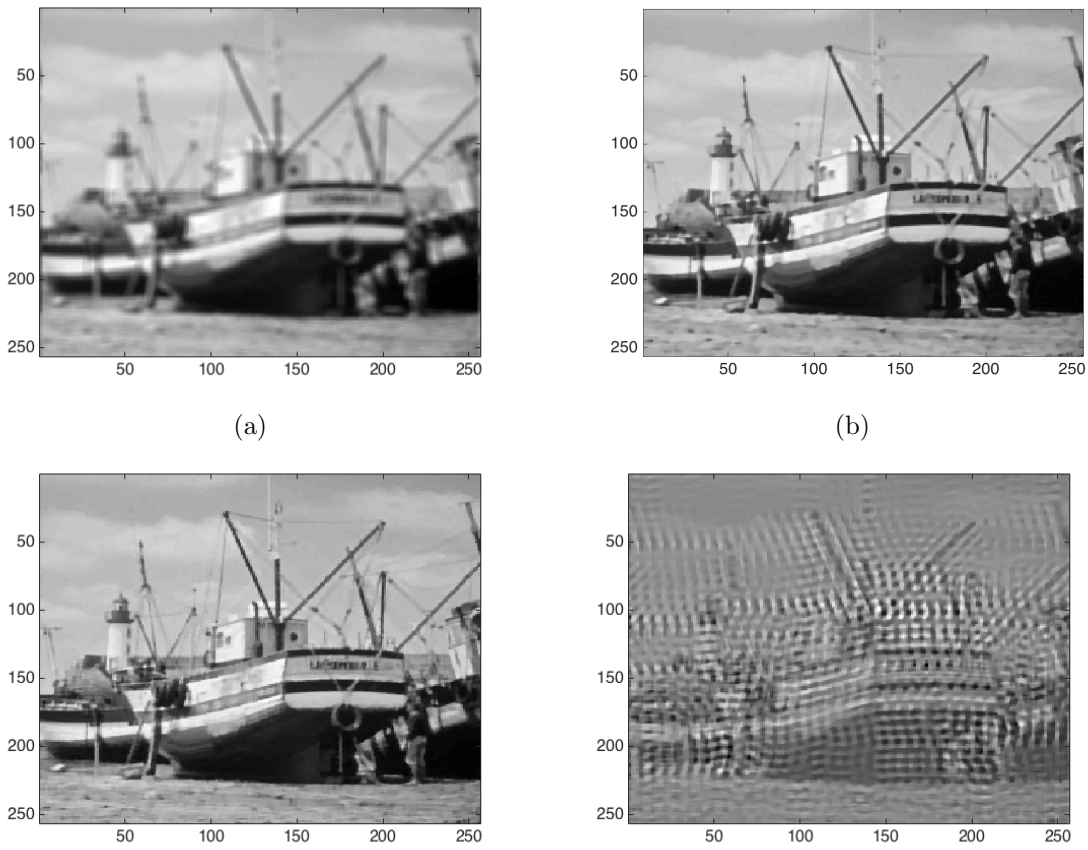

Figure 4. *Deconvolution experiment - `Boat` test image (256 × 256 pixels): (a) Blurred and noisy image y, (b)-(d) MAP estimators corresponding to models $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$.* FibBoat1

By applying this procedure we obtained that $\mathcal{M}_1$ has the highest posterior probability $p(\mathcal{M}_1|y) = 0.964$, followed by $p(\mathcal{M}_2|y) = 0.036$ and $p(\mathcal{M}_3|y) < 0.001$ (the values of the Bayes factors for this experiment are $\hat{B}_{1,2}(y) = 26.8$ and $\hat{B}_{1,3}(y) > 10^3$). These results, which have been computing without using any form of ground truth, are in agreement with the PSNR values calculated by using the true image and provide strong evidence in favour of model $\mathcal{M}_1$. They also confirm the good performance of the Bayesian model selection technique.

Comparison with proximal MALA. We conclude this first experiment by benchmarking our estimations against Px-MALA, which targets (22) exactly. Precisely, we recalculated the models' posterior probabilities (31) with Px-MALA and obtained that $p(\mathcal{M}_1|y) = 0.962$, $p(\mathcal{M}_2|y) = 0.038$, and $p(\mathcal{M}_3|y) < 0.001$, indicating that the MYULA estimate has an approximation error of the order of 0.5% (to obtain accurate estimates for Px-MALA we used $n = 10^7$ iterations with an adaptive time-step targeting an average acceptance rate of order 45%). Moreover, comparing the chains generated with MYULA and Px-MALA revealed that MYULA is significantly more computationally efficient than Px-MALA. For illustration, Fig.

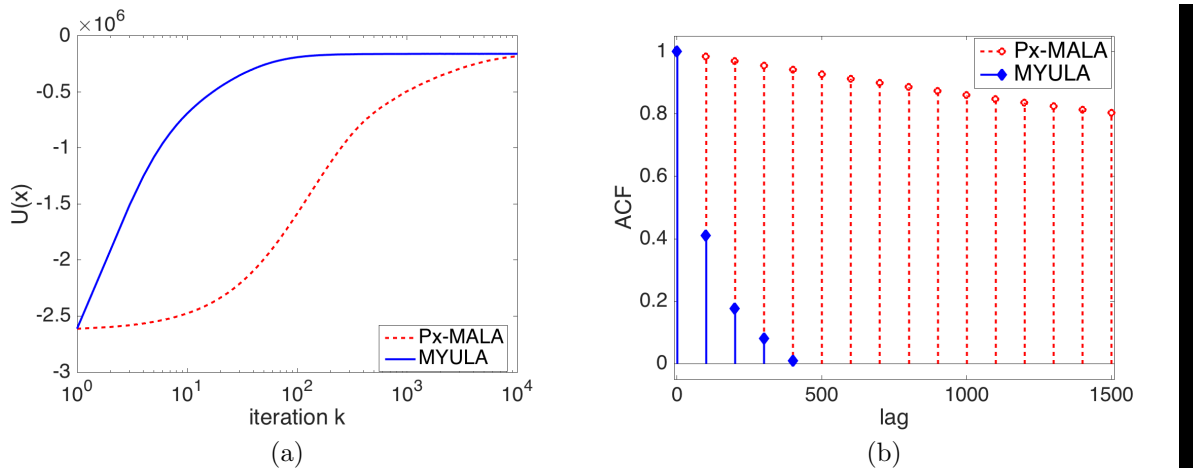

Figure 5. *MYULA and Px-MALA comparison: (a) Convergence of the chains to the typical set of* (22) *under model* $\mathcal{M}_1$ *(logarithmic scale), (b) chain autocorrelation function (ACF).*

FibBoat2

5(a) shows the transient regimes of the MYULA and Px-MALA chains related $\mathcal{M}_1$, where starting from a common initial condition the chains converge to the posterior typical set[2] of $p(x|y)$ (to improve visibility this is displayed in logarithmic scale). Observe that MYULA requires around 10^2 iterations to navigate the parameter space and reach the typical set, whereas Px-MALA requires 10^4 iterations. Furthermore, to compare the efficiency of the chains in stationarity, Fig. 5(b) shows the autocorrelation function of the chains generated by MYULA and Px-MALA. To highlight the efficiency of MYULA we have used the chains' slowest component[3] as summary statistic. Again, observe that MYULA is clearly significantly more efficient than Px-MALA. From a practitioner's viewpoint, this efficiency advantage is further accentuated by the fact that MYULA iterations are almost twice less computationally expensive than Px-MALA iterations, which include the MH step.

eriment-2:-image

4.1.3. Experiment 2: Image deconvolution with wavelet frame.

Experiment setup. The second model selection experiment we consider involves three alternative image deconvolution models with different priors. This experiment is more challenging than the previous one because priors operate indirectly on y through x. We consider three models of the form

nvolutionL1wave

$$\mathcal{M}_j : \quad p(x|y) \propto \exp\left[-(\|y - Hx\|^2/2\sigma^2) - \beta_j \|\Psi_j x\|_1\right] \tag{23}$$

where Ψ_j is a model dependent frame:

- $\mathcal{M}_1$: Ψ_1 is a redundant Haar frame with 6-level, and $\beta_1 = 0.02$ is selected automatically by using a hierarchical Bayesian method [38],

[2] In stationarity, $x|y$ is with very high probability in the neighbourhood of the $(d-1)$-dimensional shell $\{x : U(x) = \mathbb{E}[U(x)|y]\}$, see [37]

[3] The chain's slowest component was identified by doing an approximate singular value decomposition of the chain's covariance matrix and then projecting the samples on the dominant eigenvector.

- $\mathcal{M}_2$: Ψ_2 is a redundant Haar frame with 3-level, and $\beta_2 = 0.02$ is selected automatically by using a hierarchical Bayesian method [38],
- $\mathcal{M}_3$: Ψ_3 is a redundant Haar frame with 3-level, and $\beta_3 = 0.003$ is selected automatically by using the L-curve method [20].

To make the selection problem even more challenging, in this experiment we use a higher noise level $\sigma = 1.76$, related to a blurred signal-to-noise ratio of 30dB. We note that (23) is log-concave and MAP estimation can be performed efficiently by proximal convex optimisation.

Fig. 6 presents an experiment with the `Flinstones` test image of size $d = 256 \times 256$ pixels. Fig. 4(a) shows the blurred and noisy observation y used in this experiment, which we generated by using a 5×5 uniform blur and $\sigma = 1.76$, and Fig. 6(b)-(d) show the MAP estimates obtained with $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$ by convex optimisation (we used a forward-backward splitting algorithm [17] that is algorithmically similar to MYULA, and which required approximately 2 500 iterations and 2 minutes per model[4]). We observe in Figure 4 that models $\mathcal{M}_1$ and $\mathcal{M}_2$ have produced sharp images with fine detail, whereas $\mathcal{M}_3$ is misspecified. In terms of estimation performance with respect to the truth, the estimate of Figure 6(c) corresponding to model $\mathcal{M}_2$ achieves the highest peak signal-to-noise-ratio (PSNR) of 20.8dB, $\mathcal{M}_1$ scores 20.6dB, and $\mathcal{M}_3$ scores 11.6dB.

Model selection in the absence of ground truth. Similarly to the previous experiment, we used MYULA to perform Bayesian model selection intrinsically. Precisely, we used MYULA to generate three sets of $n = 10^5$ samples $X_1^M, \ldots, X_n^M$ approximately distributed according to (23) with the parameters corresponding to $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$. For this experiment we implemented MYULA with $f(x) = \|y - Hx\|^2/2\sigma^2$ and $g(x) = \beta_j \|\Psi_j x\|_1$, with fixed algorithm parameters $\lambda = L_f^{-1} = 4.5$ and $\gamma = L_f^{-1}/5 = 0.9$. Computing these samples required 50 minutes per model. Following on from this, we used the samples to calibrate the high-posterior-density regions $\mathcal{C}_j^\star$ of each model at level 20%, and then computed the Bayes factors between the models by using (21) (see C.1 for details).

By applying this procedure we obtained that $\mathcal{M}_2$ has the highest posterior probability $p(\mathcal{M}_2|y) = 0.42$, followed by $p(\mathcal{M}_1|y) = 0.32$ and $p(\mathcal{M}_3|y) = 0.26$ (the values of the Bayes factors for this experiment are $\hat{B}_{2,1}(y) = 1.31$ and $\hat{B}_{2,3}(y) = 1.62$). Note that these results, which have been computing without using any form of ground truth, are in agreement with the PSNR values calculated by using the true image and indicate that $\mathcal{M}_2$ is the most appropriate model for data y.

Comparison with proximal MALA. Again, we conclude our second experiment by benchmarking our estimations against Px-MALA, which targets (23) exactly. Precisely, we recalculated the models' posterior probabilities (31) with Px-MALA and obtained that $p(y|\mathcal{M}_1) = 0.41$, $p(y|\mathcal{M}_2) = 0.33$, and $p(y|\mathcal{M}_3) = 0.26$, indicating that the MYULA estimate has an approximation error of the order of 0.5% (to obtain accurate estimates for Px-MALA we used $n = 10^7$ iterations with an adaptive time-step targeting an average acceptance rate of order 45%). Moreover, efficiency analyses indicate that in this case MYULA is approximately an order of magnitude more efficient per iteration than Px-MALA, with an additional advantage in terms of time-normalised computational efficiency because of a lower computational cost

[4]The computation of the MAP estimates with the SALSA convex optimisation algorithm [1], which is faster than the forward-backward splitting algorithm, required 4 seconds per model.

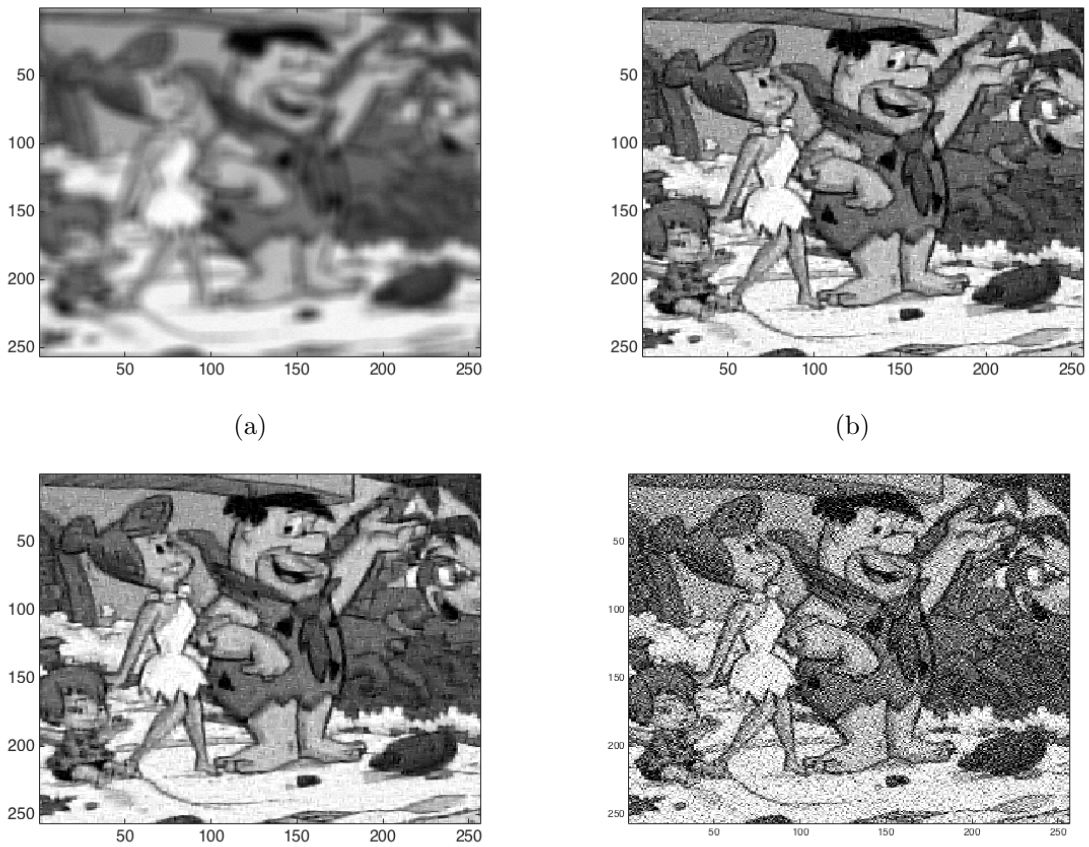

Figure 6. *Deconvolution experiment - `Flinstones` test image (256×256 pixels): (a) Blurred and noisy image y, (b)-(d) MAP estimators corresponding to models $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$.*

FibFlin

per iteration.

exp:BUQ

4.2. Bayesian uncertainty quantification via posterior credible sets.

4.2.1. Bayesian analysis and computation. As mentioned earlier, point estimators such as $\hat{x}_{MAP}$ deliver accurate results but do not provide information about the posterior uncertainty of x. Given the uncertainty that is inherent to ill-posed and ill-conditioned inverse problems, it would be highly desirable to complement point estimators with posterior credibility sets that indicate the region of the parameter space where most of the posterior probability mass of x lies. This is formalised in the Bayesian decision theory framework by computing *credible regions* [40]. A set C_α is a posterior credible region with confidence level $(1-\alpha)$ if

$$\mathbb{P}\left[x \in \mathcal{C}_\alpha | y\right] = 1 - \alpha.$$

It is easy to check that for any $\alpha \in (0,1)$ there are infinitely many regions of the parameter space that verify this property. Among all possible regions, the so-called *highest posterior*

density (HPD) region has minimum volume [40], and is given by

$$\mathcal{C}_\alpha^\star = \{\boldsymbol{x} : U(x) \leq \eta_\alpha\} \tag{24}$$

with $\eta_\alpha \in \mathbb{R}$ chosen such that $\int_{\mathcal{C}_\alpha^\star} p(x|y)\mathrm{d}\boldsymbol{x} = 1-\alpha$ holds. This joint credible set has the important advantage that it can be enumerated by simply specifying the scalar value η_α.

From a computation viewpoint, calculating credible sets for images is very challenging because it requires solving very high-dimensional integrals of the form $\int_{\mathcal{C}_\alpha^\star} p(x|y)\mathrm{d}x$. In this work, we use MYULA to approximate these integrals.

4.2.2. Experiment 3: Tomographic image reconstruction.

Experiment setup. The third experiment we consider is a tomographic image reconstruction problem with a total-variation prior. The goal is to recover the image $x \in \mathbb{R}^n$ from an incomplete and noisy set of Fourier measurements $y = AFx + w$, where F is the discrete Fourier transform operator, A is a tomographic sampling mask, and $w \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I}_n)$. This inverse problem is ill-posed, resulting in significant uncertainty about the true value of x. Similarly to Experiment 1, in this experiment we regularise the problem and reduce the uncertainty about x by using a total-variation prior promoting piecewise regular images. The resulting posterior $p(x|y)$ is

$$\pi(x) \propto \exp\left[-\|y - AFx\|^2/2\sigma^2 - \beta TV(x)\right]. \tag{25}$$

with fixed hyper-parameters $\sigma > 0$ and $\beta > 0$ set manually by an expert. We note that this density is log-concave and MAP estimation can be performed efficiently by proximal convex optimisation.

Figure 7 presents an experiment with the `Shepp-Logan phantom` magnetic resonance image (MRI) of size $d = 128 \times 128$ pixels presented in Figure 7(a). Figure 7(b) shows a noisy tomographic measurement y of this image, contaminated with Gaussian noise with $\sigma = 7 \times 10^{-2}$ (to improve visibility Figure 7(b) shows the amplitude of the Fourier coefficients in logarithmic scale, with black regions representing unobserved coefficients). Notice from Figure 7(b) that only 15% of the original Fourier coefficients are observed. Moreover, Figure 7(c) shows the Bayesian estimate $\hat{x}_{MAP}$ associated with (25) with hyper-parameter value $\beta = 5$.

Bayesian uncertainty analysis. We now conduct a simple Bayesian uncertainty analysis to illustrate how posterior credible sets can inform decision-making. For illustration, suppose that the structure highlighted in red in Figure 7(c) is relevant from a clinical viewpoint because it provides important information for diagnosis or treatment related decision-making. Also, suppose that we first observe this structure in the Bayesian estimate $\hat{x}_{MAP}$ and that, following on from this, we wish to explore the posterior uncertainty about x to learn more about the structure. In particular, here we conduct a simple analysis to show that there is lack of confidence regarding the presence of this structure in the true image (i.e., the structure could be an artefact). Precisely, this is achieved by computing the HDP credible region $\mathcal{C}_\alpha^\star$ and showing that it includes solutions that are essentially equivalent to $\hat{x}_{MAP}$ except for the fact that they do not have the structure of interest.

As alternative solution or "counter example" of $\hat{x}_{MAP}$, consider the image $x_\dagger$ displayed in Figure 7(d). This image is equivalent to $\hat{x}_{MAP}$ except for the fact that the structure of interest

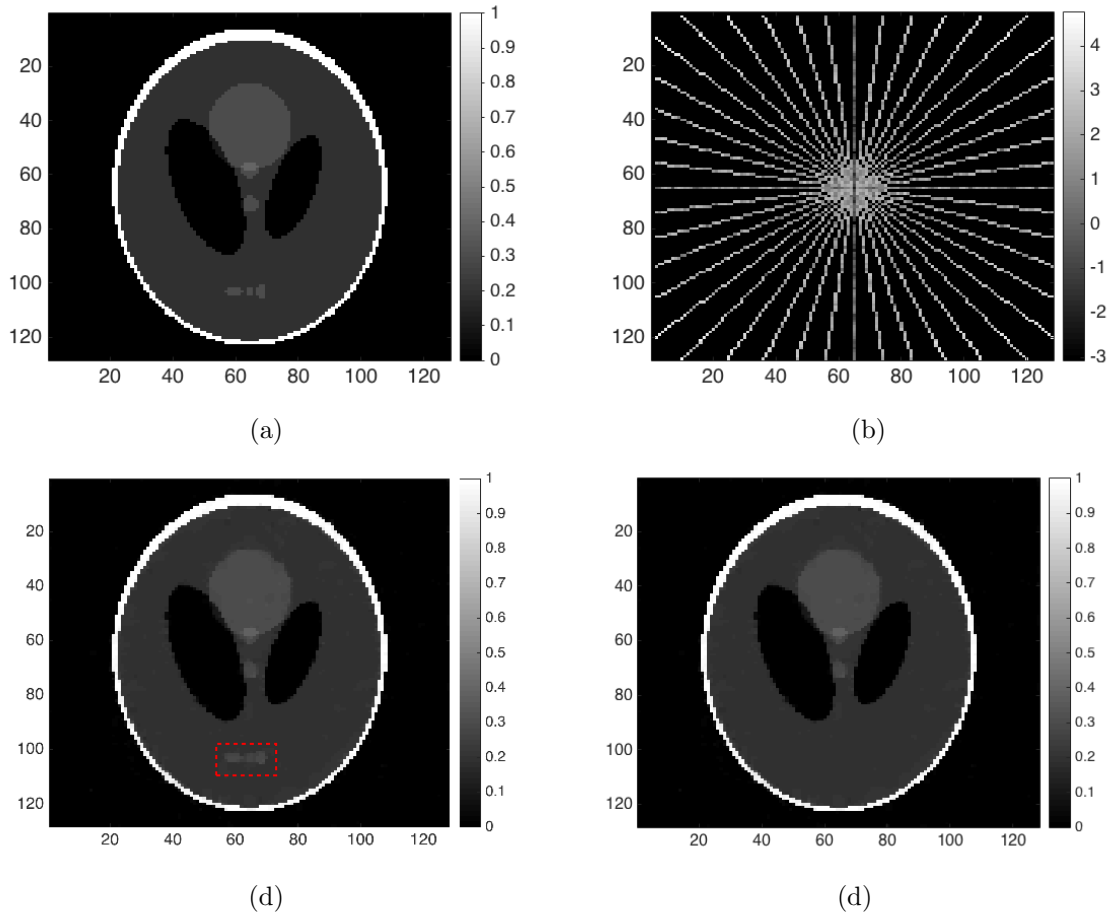

Figure 7. *Tomography experiment: (a)* `Shepp-Logan phantom` *image (*128×128 *pixels), (b) tomographic observation y (amplitude of Fourier coefficients in logarithmic scale), (c) MAP estimator* $\hat{x}_{MAP}$*, (d) counter example image* $x_\dagger$*.*

FigMRI1

has been removed (we generated this image by modifying $\hat{x}_{MAP}$ by applying a segmentation-inpainting process to replace the structure with the surrounding intensity level). Of course, clinicians observing $x_\dagger$ images would potentially arrive to significantly different conclusions about the diagnosis or the treatment required. This test image scores $U(x_\dagger) = 1.27 \times 10^4$.

To determine if $x_\dagger$ belongs to $\mathcal{C}^\star_\alpha$ we used MYULA to generate $n = 10^5$ samples from (25), and calculated the HPD threshold η_α by estimating the $(1-\alpha)$-quantile of $U(x)$ (we implemented the algorithm with $f(x) = \|y - AFx\|^2/2\sigma^2$ and $g(x) = \beta TV(x)$, with fixed parameters $\lambda = L_f^{-1} = 1 \times 10^{-4}$ and $\gamma_k = L_f^{-1}/10 = 10^{-5}$, and by using Chambolle's algorithm [7] to evaluate the proximal operator of the TV-norm). Fig. 8(a) shows the threshold values η_α for a range of values of $\alpha \in [0.01, 0.99]$. Observe that $U(x_\dagger) = 1.27 \times 10^4$ is significantly lower than the values displayed in Fig. 8(a), indicating that the counter example image $x_\dagger$ belongs to set of likely solutions to the inverse problem (e.g., at level 90% $\eta_{0.10} = 2.34 \times 10^4$ hence $x_\dagger \in \mathcal{C}^\star_{0.10}$, for information $U(\hat{x}_{MAP}) = 1.21 \times 10^4$). Based on this we conclude that, with the

current number of observations and noise level, it is not possible to assert confidently that the structure considered is present in the true image. Consequently, we would recommend that this data is not used as primary evidence to support decision-making about this structure. Generating the Monte Carlo samples and computing the HPD threshold values required 15 minutes.

Comparison with proximal MALA. We conclude this experiment by benchmarking our estimations against Px-MALA, which targets (25) exactly (to obtain accurate estimates for Px-MALA we use $n = 10^7$ iterations with an adaptive time-step targeting an average acceptance rate of order 45%). The HPD threshold values η_α obtained with Px-MALA are reported in Fig. 8(a), notice the approximation error of order of 3% due to MYULA's estimation bias (this does not affect the conclusions of the experiment). With regards to computational performance, an efficiency analysis of the two algorithms indicates that for this model MYULA is approximately two orders of magnitude more efficient than Px-MALA in terms of integrated autocorrelation time (for illustration Fig. 8(b) compares the autocorrelation functions for slowest component[5] of the MYULA and Px-MALA chains.

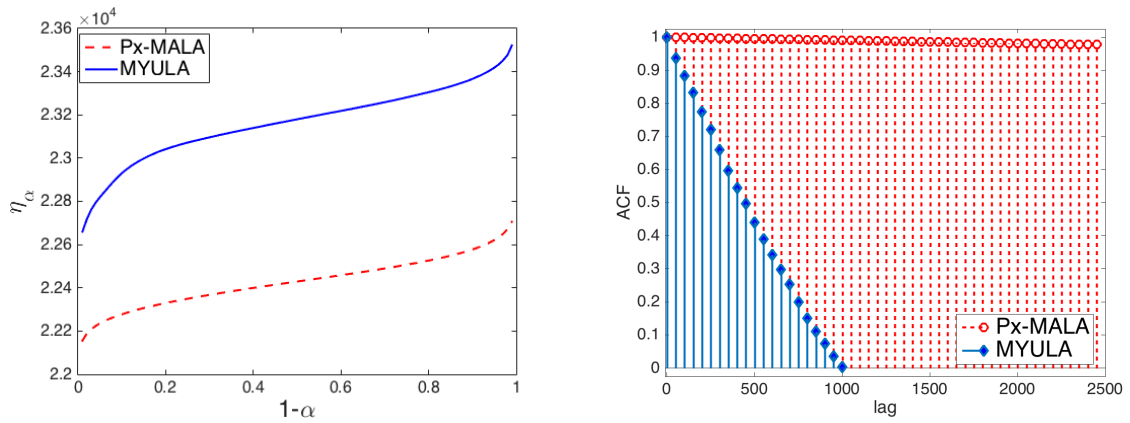

(a) (b)

Figure 8. *Tomography experiment: (a) HDP region thresholds η_α for MYULA and Px-MALA, (b) chain autocorrelation functions for MYULA and Px-MALA.*

FigMRI3

ssec:exp4

4.2.3. Experiment 4: Sparse image deconvolution with an ℓ_1 prior.

Experiment setup. The fourth experiment we consider is a sparse image deconvolution problem with a Laplace or ℓ_1 prior. Again, we aim to recover $x \in \mathbb{R}^n$ from $y = Hx + \boldsymbol{w}$, where H is a circulant blurring matrix and $\boldsymbol{w} \sim \mathcal{N}(0, \sigma^2 \boldsymbol{I}_n)$. We expect sparse solutions and use a Laplace prior related to the ℓ_1 norm of x. The resulting posterior $p(x|y)$ is

deconvolution

$$\pi(x) \propto \exp\left[-\|y - Hx\|^2/2\sigma^2 - \beta\|x\|_1\right]. \tag{26}$$

with fixed hyper-parameters $\sigma > 0$ and $\beta > 0$ set manually by an expert. Similarly to the previous experiments, we notice that this density is log-concave and MAP estimation can be

[5] Again, the chain's slowest component was identified by doing an approximate singular value decomposition of the chain's covariance matrix and then projecting the samples on the dominant eigenvector.

performed efficiently by proximal convex optimisation.

Figure 9 presents an experiment with a microscopy dataset of [49] related to high-resolution live cell imaging. Figure 9(a) shows an observation y of field of size $4\mu m \times 4\mu m$ containing 100 molecules. This low-resolution observation has been acquired with an instrument specific point-spread-function of size 16×16 pixels and a blurred signal-to-noise ratio of 20dB (see [49] for more details). Figure 9(b) shows the Bayesian estimate $\hat{x}_{MAP}$ associated with (26) with hyper-parameter value $\alpha = 0.01$ (notice that $\hat{x}_{MAP}$ is displayed in logarithmic scale to improve visibility). Computing this estimate with a forward-backward splitting optimisation algorithm, which is algorithmically similar to MYULA, required approximately 5 minutes[6].

Bayesian uncertainty analysis. As second example of Bayesian uncertainty quantification, we use $\mathcal{C}^\star_\alpha$ to examine the uncertainty about the position of the group of molecules highlighted in red in Fig. 9, which we assume to be relevant for an application considered. Precisely, we used $n = 10^5$ samples generated with MYULA to compute $\mathcal{C}^\star_\alpha$ with $\alpha = 0.01$ related to the 99% confidence level, and obtained the threshold value $\eta_{0.01} = 9.69 \times 10^4$. Following on from this, to explore $\mathcal{C}^\star_{0.01}$ to quantify the uncertainty about the exact position of the molecules, we generated several surrogate test images by modifying $\hat{x}_{MAP}$ by displacing the molecules in different directions until these surrogates exit $\mathcal{C}^\star_{0.01}$ (similarly to the previous experiment, the resulting empty space was filled by inpainting). Figure 9(c) shows the posterior uncertainty of the molecule positions (note that for visibility the figure focuses on the region of interest). This analysis reveals that the uncertainty at level 99% is of the order of ± 5 pixels vertically and ± 8 pixels horizontally, corresponding to $\pm 78nm$ and $\pm 125nm$. It is worth mentioning that these results are in close in agreement with the experimental precision results reported in [49], which identified an average precision of the order of $80nm$ for the one hundred molecules.

Comparison with proximal MALA. Again, we conclude the experiment by benchmarking our estimations against Px-MALA, which targets (26) exactly (to obtain accurate estimates for Px-MALA we use $n = 2 \times 10^7$ iterations with an adaptive step-size targeting an acceptance rate of the order of 45%). Figure 9(d) compares the estimations of the threshold values η_α obtained with MYULA and Px-MALA for different values of α, indicating that the approximation errors of MYULA are of the order of 0.1%. Moreover, performance analyses based on the chains generated with each algorithm indicate that in this case MYULA is approximately one order of magnitude more computationally efficient than Px-MALA.

sec:conclusion

5. Discussion and conclusion.

This paper presented a new and general proximal MCMC methodology to perform Bayesian computation in log-concave models, with a focus on enabling advanced Bayesian analyses for imaging inverse problems that are convex and not smooth, and currently solved mainly by convex optimisation. The methodology is based on a Moreau-Yoshida-type regularised approximation of the target density that is by construction is log-concave and Lipchitz continuously differentiable, and which can be addressed efficiently by using an unadjusted Langevin MCMC algorithm. We provided a detailed theoretical analysis of this scheme, including asymptotic as well as non-asymptotic convergence results, and bounds on the convergence rate of the chains with explicit dependence on model dimension. In addition to being highly computational efficient and having a strong theoretical underpinning,

[6]The computation of the MAP estimate with the SALSA [1] convex optimisation algorithm, which is faster than the forward-backward splitting algorithm, required 2.3 seconds.

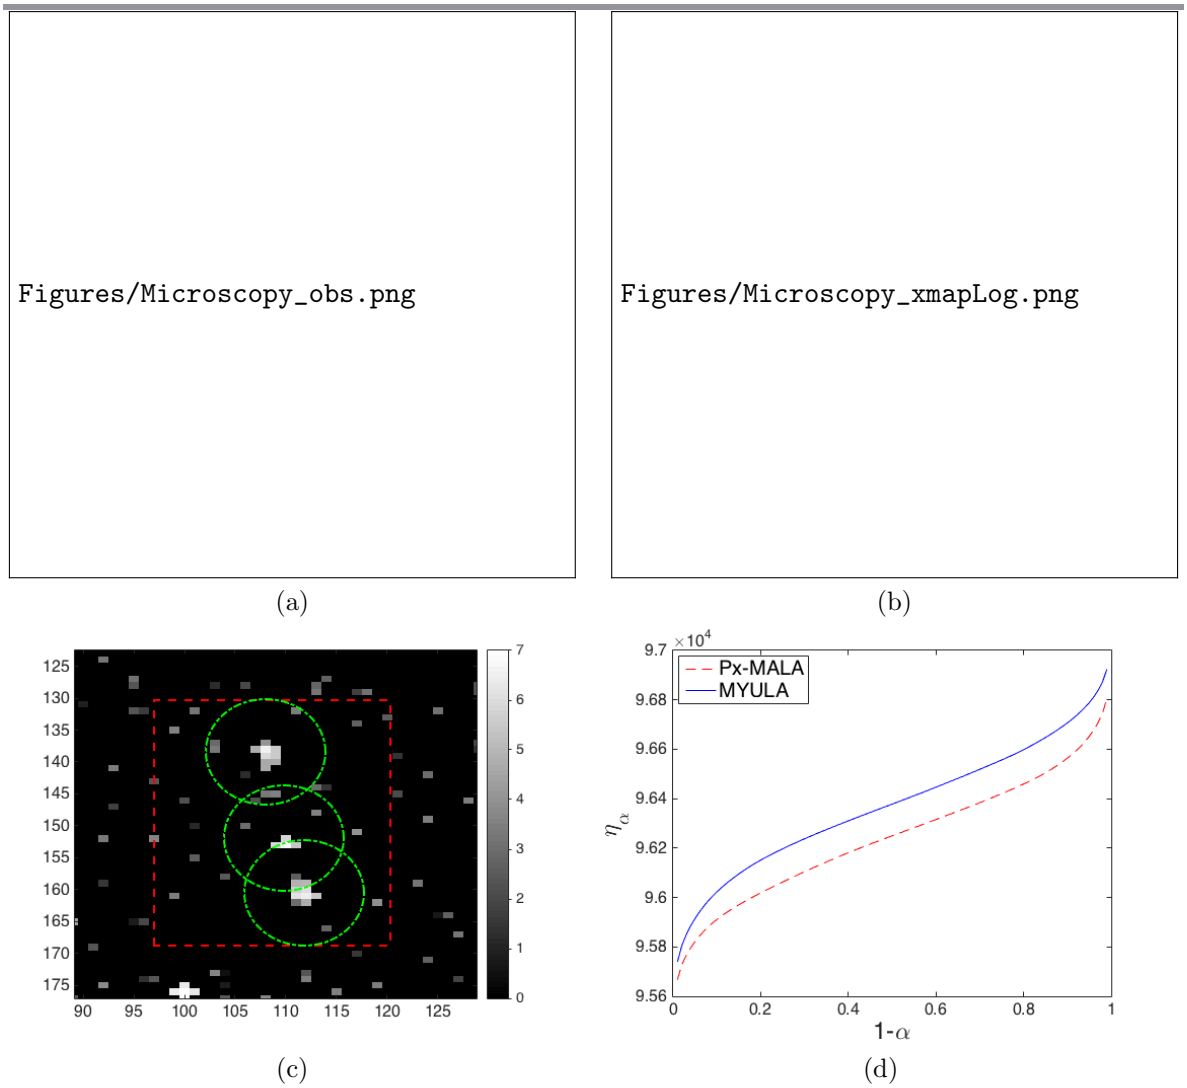

Figure 9. *Microscopy experiment: (a) Blurred image y (256×256 pixels, $4\mu m \times 4\mu m$)), (b) MAP estimate $\hat{x}_{MAP}$ (logarithmic scale), (c) molecule position uncertainty quantification (vertical: $\pm 78nm$, horizontal $\pm 125nm$), (d) HDP region thresholds η_α for MYULA and Px-MALA.*

FigMicro

this new methodology is general and can be applied straightforwardly to most problems solved by proximal optimisation, particularly all problems solved by using forward-backward splitting techniques. The proposed methodology was finally demonstrated with four experiments related to image deconvolution and tomographic reconstruction with total-variation and ℓ_1 priors, where we conducted a range of challenging analyses related to model comparison and uncertainty quantification, and where we reported estimation accuracy and computational efficiency comparisons with the proximal Metropolis-adjusted Langevin algorithm.

Furthermore, observe that the regularisation strategy used in this paper, based on the Moreau-Yoshida envelope, can also be applied straightforwardly to the Hamiltonian Monte

Carlo algorithm [17]. The theoretical and empirical properties of this algorithm are currently under investigation and will be reported separately.

Finally, it is worth mentioning that MYULA can also be applied to some models that are not log-concave, for example multi-modal models where the smooth term f is not convex. In this case, the non-smooth term g must be associated with a proper prior to guarantee that the approximation π^λ is proper. It is possible to derive non-asymptotic convergence results for these models; however, unlike the log-concave case, here the dependence w.r.t. to the dimension of the model is difficult to analyse (see [45] for details). The analysis of the performance of MYULA in non-convex settings is an important perspective for future work. Also, another important perspective for future work is to investigate ways in which the proposed regularisation approach, combined with an appropriate convex relaxation, could enable Langevin and Hamiltonian sampling in high-dimensional spaces that are discrete.

6. Acknowledgements. Part of this work was conducted when Marcelo Pereyra held a Marie Curie Intra-European Research Fellowship for Career Development at the School of Mathematics of the University of Bristol, and a Visiting Scholar at the School of Mathematical and Computer Sciences of Heriot-Watt University.

References.

[1] M. Afonso, J. M. Bioucas-Dias, and M. A. T. Figueiredo, *An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems*, IEEE. Trans. on Image Process., 20 (2011), pp. 681–695.

[2] Y. F. Atchadé, *A Moreau-Yosida approximation scheme for a class of high-dimensional posterior distributions*, ArXiv e-prints, (2015), arXiv:1505.07072.

[3] D. Bakry, F. Barthe, P. Cattiaux, and A. Guillin, *A simple proof of the Poincaré inequality for a large class of probability measures.*, Electronic Communications in Probability [electronic only], 13 (2008), pp. 60–66, http://eudml.org/doc/225690.

[4] J. M. Bardsley, A. Solonen, H. Haario, and M. Laine, *Randomize-then-optimize: A method for sampling from posterior distributions in nonlinear inverse problems*, SIAM Journal on Scientific Computing, 36 (2014), pp. A1895–A1910.

[5] S. Bonettini, A. Cornelio, and M. Prato, *A new semiblind deconvolution approach for Fourier-based image restoration: An application in astronomy*, SIAM J. Imaging Sci., 6 (2013), pp. 1736–1757, doi:10.1137/120873169, http://dx.doi.org/10.1137/120873169, arXiv:http://dx.doi.org/10.1137/120873169.

[6] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, *Phase retrieval via matrix completion*, SIAM J. Imaging Sci., 6 (2013), pp. 199–225.

[7] A. Chambolle, *An algorithm for total variation minimization and applications*, Journal of Mathematical Imaging and Vision, 20 (2004), pp. 89–97, doi:10.1023/B:JMIV.0000011325.36760.1e, http://dx.doi.org/10.1023/B:JMIV.0000011325.36760.1e.

[8] R. H. Chan, J. Yang, and X. Yuan, *Alternating direction method for image inpainting in wavelet domains*, SIAM J. Imaging Sci., 4 (2011), pp. 807–826, doi:10.1137/100807247, http://dx.doi.org/10.1137/100807247, arXiv:http://dx.doi.org/10.1137/100807247.

[9] P. L. Combettes and J.-C. Pesquet, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, Springer New York, New York, NY, 2011, ch. Proximal

Splitting Methods in Signal Processing, pp. 185–212.

pesquet:2011 [10] P. L. Combettes and J.-C. Pesquet, *Proximal splitting methods in signal processing*, in Fixed-Point Algorithms for Inverse Problems in Science and Engineering, H. H. Bauschke, R. S. Burachik, P. L. Combettes, V. Elser, D. R. Luke, and H. Wolkowicz, eds., Springer New-York, 2011, pp. 185–212.

dalalyan:2014 [11] A. S. Dalalyan, *Theoretical guarantees for approximate sampling from smooth and log-concave densities*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), (2016), pp. n/a–n/a, doi:10.1111/rssb.12183, http://dx.doi.org/10.1111/rssb.12183.

Donoho2006 [12] D. L. Donoho, *Compressed sensing*, IEEE Trans. Inf. Theory, 52 (2006), pp. 1289–1306.

us:moulines:2015 [13] A. Durmus and E. Moulines, *Non-asymptotic convergence analysis for the unadjusted langevin algorithm*, Accepted for publication in Ann. Appl. Probab. 1507.05021, arXiv, July 2015, http://arxiv.org/pdf/1507.05021v1.pdf.

eberle:2015 [14] A. Eberle, *Reflection couplings and contraction rates for diffusions*, Probab. Theory Related Fields, (2015), pp. 1–36, doi:10.1007/s00440-015-0673-1, http://dx.doi.org/10.1007/s00440-015-0673-1.

ermak:1975 [15] D. L. Ermak, *A computer simulation of charged particles in solution. i. technique and equilibrium properties*, The Journal of Chemical Physics, 62 (1975), pp. 4189–4196.

nzano:levan:2001 [16] M. Florenzano and C. Le Van, *Finite dimensional convexity and optimization*, vol. 13 of Studies in Economic Theory, Springer-Verlag, Berlin, 2001, doi:10.1007/978-3-642-56522-9, http://dx.doi.org/10.1007/978-3-642-56522-9. In cooperation with Pascal Gourdel.

Green2015 [17] P. J. Green, K. Łatuszyński, M. Pereyra, and C. P. Robert, *Bayesian computation: a summary of the current state, and samples backwards and forwards*, Statistics and Computing, 25 (2015), pp. 835–862, doi:10.1007/s11222-015-9574-5, http://dx.doi.org/10.1007/s11222-015-9574-5.

grenander:1983 [18] U. Grenander, *Tutorial in pattern theory*. Division of Applied Mathematics, Brown University, Providence, 1983.

nder:miller:1994 [19] U. Grenander and M. I. Miller, *Representations of knowledge in complex systems*, J. Roy. Statist. Soc. Ser. B, 56 (1994), pp. 549–603, http://links.jstor.org/sici?sici=0035-9246(1994)56:4⟨549:ROKICS⟩2.0.CO;2-2&origin=MSN. With discussion and a reply by the authors.

Hanke1993 [20] M. Hanke and C. Hansen, *Regularization methods for large-scale problems*, Surv. Math. Ind., 3 (1993), pp. 253–315.

Haro:2012 [21] G. Haro, A. Buades, and J.-M. Morel, *Photographing paintings by image fusion*, SIAM J. Imaging Sci., 5 (2012), pp. 1055–1087, doi:10.1137/120873923, http://dx.doi.org/10.1137/120873923, arXiv:http://dx.doi.org/10.1137/120873923.

somersalo:2005 [22] J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*, Springer, New-York, 2005.

khasminskii:1960 [23] R. Z. Khas'minskii, *Ergodic properties of recurrent diffusion processes and stabilization of the solution to the cauchy problem for parabolic equations*, Theory of Probability & Its Applications, 5 (1960), pp. 179–196, doi:10.1137/1105016, http://dx.doi.org/10.1137/1105016, arXiv:http://dx.doi.org/10.1137/1105016.

Lebrun:2013 [24] M. Lebrun, A. Buades, and J. M. Morel, *A nonlocal Bayesian image denoising algorithm*, SIAM J. Imaging Sci., 6 (2013), pp. 1665–1688.

[25] F. Lucka, *Fast markov chain monte carlo sampling for sparse bayesian inference in high-dimensional inverse problems using l1-type priors*, Inverse Problems, 28 (2012), p. 125012.

[26] F. Lucka, *Fast gibbs sampling for high-dimensional bayesian inversion*, Inverse Problems, 32 (2016), p. 115019, http://stacks.iop.org/0266-5611/32/i=11/a=115019.

[27] M. Lustig, D. Donoho, and J. M. Pauly, *Sparse mri: The application of compressed sensing for rapid mr imaging*, Magnetic Resonance in Medicine, 58 (2007), pp. 1182–1195.

[28] J.-M. Marin and C. Robert, *Bayesian core: a practical approach to computational Bayesian statistics*, Springer Science & Business Media, 2007.

[29] S. Meyn and R. Tweedie, *Markov Chains and Stochastic Stability*, Cambridge University Press, New York, NY, USA, 2nd ed., 2009.

[30] V. I. Morgenshtern and E. J. Candès, *Super-resolution of positive sources: The discrete setup*, SIAM J. Imaging Sci., 9 (2016), pp. 412–444.

[31] R. M. Neal, *Bayesian learning via stochastic dynamics*, in Advances in Neural Information Processing Systems 5, [NIPS Conference], San Francisco, CA, USA, 1993, Morgan Kaufmann Publishers Inc., pp. 475–482, http://dl.acm.org/citation.cfm?id=645753.667903.

[32] N. Parikh and S. Boyd, *Proximal algorithms*, Foundations and Trends in Optimization, 1 (2013), pp. 123–231.

[33] G. Parisi, *Correlation functions and computer simulations*, Nuclear Physics B, 180 (1981), pp. 378–384.

[34] T. Park and G. Casella, *The Bayesian lasso*, J. Amer. Statist. Assoc., 103 (2008), pp. 681–686, doi:10.1198/016214508000000337, http://dx.doi.org/10.1198/016214508000000337.

[35] M. Pereyra, *Maximum-a-posteriori estimation with Bayesian confidence regions*, SIAM J. Imaging Sci. to appear.

[36] M. Pereyra, *Proximal Markov chain Monte Carlo algorithms*, Statistics and Computing, (2015). open access paper, http://dx.doi.org/10.1007/s11222-015-9567-4.

[37] M. Pereyra, *Maximum-a-posteriori estimation with Bayesian confidence regions*, ArXiv e-prints, (2016), arXiv:1602.08590.

[38] M. Pereyra, J. M. Bioucas-Dias, and M. A. T. Figueiredo, *Maximum-a-posteriori estimation with unknown regularisation parameters*, in Proc. European Signal Proc. Conf. (EUSIPCO), Nice, France, Sep. 2015., Aug 2015, pp. 230–234.

[39] M. Pereyra, P. Schniter, E. Chouzenoux, J.-C. Pesquet, J.-Y. Tourneret, A. Hero, and S. Mclaughlin, *A survey of stochastic simulation and optimization methods in signal processing*, IEEE. J. Selected Topics in Signal Process., 10 (2016), pp. 224–241.

[40] C. P. Robert, *The Bayesian Choice (second edition)*, Springer Verlag, New-York, 2001.

[41] C. P. Robert and D. Wraith, *Computational methods for Bayesian model choice*, AIP Conf. Proc., 1193 (2009), pp. 251–262.

[42] G. O. Roberts and R. L. Tweedie, *Exponential convergence of Langevin distributions and their discrete approximations*, Bernoulli, 2 (1996), pp. 341–363, doi:10.2307/3318418, http://dx.doi.org/10.2307/3318418.

[43] R. T. Rockafellar and R. J.-B. Wets, *Variational analysis*, vol. 317 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences],

Springer-Verlag, Berlin, 1998, doi:10.1007/978-3-642-02431-3, http://dx.doi.org/10.1007/978-3-642-02431-3.

[44] P. J. Rossky, J. D. Doll, and H. L. Friedman, *Brownian dynamics as smart Monte Carlo simulation*, The Journal of Chemical Physics, 69 (1978), pp. 4628–4633, doi:http://dx.doi.org/10.1063/1.436415, http://scitation.aip.org/content/aip/journal/jcp/69/10/10.1063/1.436415.

[45] D. Talay and L. Tubaro, *Expansion of the global error for numerical schemes solving stochastic differential equations*, Stochastic Anal. Appl., 8 (1990), pp. 483–509 (1991), doi:10.1080/07362999008809220, http://dx.doi.org/10.1080/07362999008809220.

[46] S. J. Vollmer and K. C. Zygalakis, *(Non-) asymptotic properties of Stochastic Gradient Langevin Dynamics*, ArXiv e-prints, (2015), arXiv:1501.00438.

[47] Z. Wang, J. M. Bardsley, A. Solonen, T. Cui, and Y. M. Marzouk, *Bayesian inverse problems with l_1 priors: a randomize-then-optimize approach*, arXiv preprint arXiv:1607.01904, (2016).

[48] Z. Xing, M. Zhou, A. Castrodad, G. Sapiro, and L. Carin, *Dictionary learning for noisy and incomplete hyperspectral images*, SIAM J. Imaging Sci., 5 (2012), pp. 33–56, doi:10.1137/110837486, http://dx.doi.org/10.1137/110837486, arXiv:http://dx.doi.org/10.1137/110837486.

[49] L. Zhu, W. Zhang, D. Elnatan, and B. Huang, *Faster STORM using compressed sensing*, Nat. Meth., 9 (2012), pp. 721–723.

Appendix A. Proof of Proposition 1.

We preface the proof by a Lemma.

Lemma 4. *Let* $\mathrm{g} : \mathbb{R}^d \to (-\infty, +\infty]$ *be a lower bounded, l.s.c convex function satisfying* $0 < \int_{\mathbb{R}^d} \mathrm{e}^{-\mathrm{g}(y)} \mathrm{d}y < +\infty$. *Then there exists* $x_{\mathrm{g}} \in \mathbb{R}^d$, $R_{\mathrm{g}}, \rho_{\mathrm{g}} > 0$ *such that for all* $x \in \mathbb{R}^d$, $x \notin \mathrm{B}(x_{\mathrm{g}}, R_{\mathrm{g}})$, $\mathrm{g}(x) - \mathrm{g}(x_{\mathrm{g}}) \geq \rho_{\mathrm{g}} \|x - x_{\mathrm{g}}\|$.

Proof. The proof is a simple extension of the one of [3, Theorem 2.2.2], where g is assumed to be continuously differentiable.

We first show that g is finite on a non-empty open set of $\mathbb{R}^d$. Note since $\int_{\mathbb{R}^d} \mathrm{e}^{-\mathrm{g}(y)} \mathrm{d}y > 0$, the set $\{\mathrm{g} < \infty\}$ can not be contained in a k-dimensional hyperplane, for $k \in \{0, \cdots, d-1\}$. Then, there exists $d+1$ points $\{\mathrm{v}_i\}_{0 \leq i \leq d} \subset \{\mathrm{g} < \infty\}$ such that the vectors $\{\mathrm{v}_i - \mathrm{v}_0\}_{1 \leq i \leq d}$ are linearly independent. Denote by $\mathrm{co}(\mathrm{v}_0, \cdots, \mathrm{v}_d)$ the convex hull of $\{\mathrm{v}_i\}_{0 \leq i \leq d}$ defined by

$$\mathrm{co}(\mathrm{v}_0, \cdots, \mathrm{v}_d) = \left\{ \sum_{i=0}^{d} \alpha_i \mathrm{v}_i \mid \sum_{i=0}^{d} \alpha_i = 1 \ , \forall i \in \{0, \cdots, d\} \ , \ \alpha_i \geq 0 \right\} .$$

Since g is convex and $\mathrm{co}(\mathrm{v}_0, \cdots, \mathrm{v}_d) \subset \{\mathrm{g} < \infty\}$, we have

$$\sup_{y \in \mathrm{co}(\mathrm{v}_0, \cdots, \mathrm{v}_d)} |\mathrm{g}(y)| \leq M_{\mathrm{co}} = \max_{i \in \{0, \cdots, d\}} \{|\mathrm{g}(\mathrm{v}_i)|\} . \tag{27}$$

It follows from $\{\mathrm{v}_i\}_{0 \leq i \leq d} \subset \{\mathrm{g} < \infty\}$ and g is lower bounded that M_{co} is finite. Finally by [16, Lemma 1.2.1], $\mathrm{co}(\mathrm{v}_0, \cdots, \mathrm{v}_d)$ has non empty interior.

Consider now the set $\{\mathrm{g} \leq M_{\mathrm{co}} + 1\}$. We prove by contradiction that it is a bounded subset of $\mathbb{R}^d$. Assume that for all $R \geq 0$, there exists $x_R \in \{\mathrm{g} \leq M_{\mathrm{co}} + 1\}$ and $x_R \notin \mathrm{B}(\mathrm{v}_0, R)$. Then since $\{\mathrm{g} \leq M_{\mathrm{co}} + 1\}$ is convex, it contains the convex hull of $\{\mathrm{v}_0, \cdots, \mathrm{v}_d, x_R\}$. Since

$\mathrm{co}(\mathrm{v}_0, \cdots, \mathrm{v}_d)$ has non empty interior, the volume of $\mathrm{co}(\mathrm{v}_0, \cdots, \mathrm{v}_d, x_R)$ grows at least linearly in R and the volume corresponding to $\{\mathrm{g} \leq M_{\mathrm{co}} + 1\}$ is infinite taking the limit as R goes to ∞. On the other hand, by assumption and since $\{\mathrm{v}_0, \cdots, \mathrm{v}_d, x_R\} \subset \{\mathrm{g} \leq M_{\mathrm{co}} + 1\}$, we have using the Markov inequality

$$\mathrm{Vol}\left(\{\mathrm{g} \leq M_{\mathrm{co}} + 1\}\right) \leq \mathrm{e}^{M_{\mathrm{co}}+1} \int_{\{\mathrm{g} \leq M_{\mathrm{co}}+1\}} \mathrm{e}^{-\mathrm{g}(y)} \mathrm{d}y < +\infty \, ,$$

which leads to a contradiction. Then there exists $R_{\mathrm{g}} \geq 0$, such that $\{\mathrm{g} \leq M_{\mathrm{co}}+1\} \subset \mathrm{B}(\mathrm{v}_0, R_{\mathrm{g}})$. For all $x \notin \mathrm{B}(\mathrm{v}_0, R_{\mathrm{g}})$, consider $y = R_{\mathrm{g}}(x - \mathrm{v}_0) \left\|x - \mathrm{v}_0\right\|^{-1} + \mathrm{v}_0$. Note that $y \notin \{\mathrm{g} \leq M_{\mathrm{co}} + 1\}$, so $\mathrm{g}(y) \geq M_{\mathrm{co}} + 1$. Now using the convexity of g, we have for all $x \notin \mathrm{B}(\mathrm{v}_0, R_{\mathrm{g}})$,

$$M_{\mathrm{co}} + 1 \leq \mathrm{g}(y) \leq R_{\mathrm{g}} \left\|x - \mathrm{v}_0\right\|^{-1} (\mathrm{g}(x) - \mathrm{g}(\mathrm{v}_0)) + \mathrm{g}(\mathrm{v}_0) \, .$$

Since $\mathrm{g}(\mathrm{v}_0) \leq M_{\mathrm{co}}$, we get

$$(\mathrm{g}(x) - \mathrm{g}(\mathrm{v}_0)) \geq R_{\mathrm{g}}^{-1} \left\|x - \mathrm{v}_0\right\|$$

and the proof is concluded setting $x_{\mathrm{g}} = \mathrm{v}_0$. ∎

Proof of Proposition 1. a) We first assume that **H** 2-(i) holds. By (6), $U \geq U^\lambda$ and therefore $0 < \int_{\mathbb{R}^d} \mathrm{e}^{-U(y)} \mathrm{d}y < \int_{\mathbb{R}^d} \mathrm{e}^{-U^\lambda(y)} \mathrm{d}y$. We now prove e^{-g^λ} is integrable with respect to the Lebesgue measure, which implies $y \mapsto \mathrm{e}^{-U^\lambda(y)}$ is integrable as well since f is assumed to be lower bounded. By **H**1 and Lemma 4, there exist $\rho_g > 0$, $x_g \in \mathbb{R}^d$ and $M_1 \in \mathbb{R}$ such that for all $x \in \mathbb{R}^d$, $g(x) - g(x_g) \geq M_1 + \rho_g \left\|x - x_g\right\|$. Thus, for all $x \in \mathbb{R}^d$, we have by (6) and (9)

$$\begin{aligned} g^\lambda(x) - g(x_g) &= g(\mathrm{prox}_g^\lambda(x)) - g(x_g) + (2\lambda)^{-1} \left\|x - \mathrm{prox}_g^\lambda(x)\right\|^2 \\ &\geq M_1 + \rho_g \left\|\mathrm{prox}_g^\lambda(x) - x_g\right\| + (2\lambda)^{-1} \left\|x - \mathrm{prox}_g^\lambda(x)\right\|^2 \\ &\geq M_1 + \inf_{y \in \mathbb{R}^d} \{\rho_g \left\|y - x_g\right\| + (2\lambda)^{-1} \left\|x - y\right\|^2\} \geq M_1 + \mathrm{h}^\lambda(x) \, , \end{aligned} \tag{28}$$

where $\mathrm{h}^\lambda(x)$ is the λ-Moreau Yosida envelope of $\mathrm{h}(x) = \rho_g \left\|x - x_g\right\|$. By [32, Section 6.5.1], the proximal operator associated with the norm is the block soft thresholding given for all $\lambda > 0$ and $x \in \mathbb{R}^d \setminus \{0\}$ by $\mathrm{prox}_{\mathrm{h}}^\lambda(x) = \max(0, 1 - \lambda / \left\|x\right\|) x$ and $\mathrm{prox}_{\mathrm{h}}^\lambda(0) = 0$. Therefore using again (6), it follows that there exists $M_2 \in \mathbb{R}$ such that for all $x \in \mathbb{R}^d$,

$$\mathrm{h}^\lambda(x) \geq \rho_g \left\|x - x_g\right\| + M_2 \, .$$

Combining this inequality with (28) concludes the proof.
We now assume that **H**2-(ii) holds. First, we show that for all $\lambda > 0$

$$\sup_{x \in \mathbb{R}^d} \{g(x) - g^\lambda(x)\} \leq \lambda \left\|g\right\|_{\mathrm{Lip}}^2 / 2 \, , \tag{29}$$

Indeed if this inequality holds, then for all $x \in \mathbb{R}^d$, we have

$$f(x) + g(x) - \lambda \left\|g\right\|_{\mathrm{Lip}}^2 / 2 \leq f(x) + g^\lambda(x) \, .$$

Therefore by assumption

$$\int_{\mathbb{R}^d} \mathrm{e}^{-U^\lambda(x)} \mathrm{d}x \leq \mathrm{e}^{\lambda \|g\|^2_{\mathrm{Lip}}/2} \int_{\mathbb{R}^d} \mathrm{e}^{-U(x)} \mathrm{d}x < +\infty \, .$$

We now prove (29). Using that g is Lipschitz, we have by (6), for all $x \in \mathbb{R}^d$

$$g(x) - g^\lambda(x) = g(x) - \inf_{y \in \mathbb{R}^d} \left\{ g(y) + (2\lambda)^{-1} \|x - y\|^2 \right\} = \sup_{y \in \mathbb{R}^d} \left\{ g(x) - g(y) - (2\lambda)^{-1} \|x - y\|^2 \right\}$$

$$\leq \sup_{y \in \mathbb{R}^d} \left\{ \|g\|_{\mathrm{Lip}} \|x - y\| - (2\lambda)^{-1} \|x - y\|^2 \right\} \leq \lambda \|g\|^2_{\mathrm{Lip}} /2 \, ,$$

where we have used that the maximum of $u \mapsto au - bu^2$, for $a, b \geq 0$, is given by $a^2/(4b)$.
b) This point is a straightforward consequence of (8) and (7).
c) Since π has also a density with respect to the Lebesgue measure and $U^\lambda(x) \leq U(x)$ for all $x \in \mathbb{R}^d$, we have for all $\lambda > 0$

ound˙2˙TV˙MY (30)
$$\|\pi^\lambda - \pi\|_{\mathrm{TV}} = \int_{\mathbb{R}^d} \left| \pi^\lambda(x) - \pi(x) \right| \mathrm{d}x \leq 2 A_\lambda \, ,$$

where $A_\lambda = \int_{\mathbb{R}^d} \{1 - \mathrm{e}^{g^\lambda(x) - g(x)}\} \pi^\lambda(x) \mathrm{d}x = 1 - \left\{ \int_{\mathbb{R}^d} \mathrm{e}^{-U^\lambda(x)} \mathrm{d}x \right\}^{-1} \int_{\mathbb{R}^d} \mathrm{e}^{-U(x)} \mathrm{d}x$. By (10), for all $x \in \mathbb{R}^d$, we get $\lim_{\lambda \downarrow 0} \uparrow U^\lambda(x) = U(x)$. We conclude by applying the monotone convergence theorem.
d) Using that for all $x \in \mathbb{R}^d$, $g^\lambda(x) \leq g(x)$ and $1 - \mathrm{e}^{-u} \leq u$ for all $u \geq 0$, (30) shows that

$$\|\pi^\lambda - \pi\|_{\mathrm{TV}} \leq 2 \int_{\mathbb{R}^d} \{g(x) - g^\lambda(x)\} \pi^\lambda(x) \mathrm{d}x \, .$$

Then the proof follows from (29).

proper-imp˙prior **Appendix B. Model selection using improper priors.** Model selection using improper priors can lead to tedious considerations [40]. Indeed, in that case the joint density of each model is not defined. However, this difficulty can be avoided when the considered models share the same improper prior distribution see [28]. Let $\mathcal{M}_1, \ldots, \mathcal{M}_K$ be K alternative Bayesian models having the same improper distribution with density $\tilde{p}(x)$ on $\mathbb{R}^d$ and associated to the family of likelihood functions $p_i(y|x)$ such that for all $i \in \{1, \ldots, K\}$, $\int_{\mathbb{R}^d} p_i(y|x) \tilde{p}(x) \mathrm{d}x < +\infty$. The marginal posterior probabilities of $\mathcal{M}_1, \ldots, \mathcal{M}_K$ are then defined by

−margPost (31)
$$\tilde{p}(\mathcal{M}_j|y) = \frac{\tilde{p}(y|\mathcal{M}_j) K^{-1}}{\sum_{k=1}^K \tilde{p}(y|\mathcal{M}_k) K^{-1}}, \quad j \in \{1, \ldots, K\} \, ,$$

where for all $j \in \{1, \ldots, K\}$,

$$\tilde{p}(y|\mathcal{M}_j) = \int_{\mathbb{R}^d} p_i(y|x) \tilde{p}(x) \mathrm{d}x \, .$$

HME **Appendix C. Truncated harmonic mean estimator.**

case-proper-prior

C.1. Case of proper prior distributions. Consider a positive probability density p on $\mathbb{R}^d \times \mathbb{R}^m$ for $d, m \in \mathbb{N}^*$ of the form: $p(x,y) = f(x,y)/\int_{\mathbb{R}^d\times\mathbb{R}^m} f(z,w)\mathrm{d}z\mathrm{d}w$. Assume that f is known but not the normalization constant of p. Here p plays the role of a joint distribution of the data and the parameters. It can be defined if we take a proper prior distribution for the parameters. Define for any bounded Borel set $\mathrm{A} \in \mathcal{B}(\mathbb{R}^d)$

-harmonicmean

$$I_{\mathrm{A}}(f,y) = \int_{\mathbb{R}^d} \mathbb{1}_{\mathrm{A}}(x)\frac{p(x|y)}{f(x,y)}\mathrm{d}x = \int_{\mathbb{R}^d} \mathbb{1}_{\mathrm{A}}(x)\frac{p(x|y)}{p(x,y)}\mathrm{d}x \Big/ \int_{\mathbb{R}^d\times\mathbb{R}^m} f(z,w)\mathrm{d}z\mathrm{d}w \ . \tag{32}$$

Since $p(x|y) = p(x,y)/p(y)$, the following identity holds

harmonic mean

$$p(y) = \mathrm{Vol}(\mathrm{A})\left\{ I_{\mathrm{A}}(f,y)\int_{\mathbb{R}^d\times\mathbb{R}^m} f(z,w)\mathrm{d}z\mathrm{d}w \right\}^{-1} \ . \tag{33}$$

For all $y \in \mathbb{R}^m$ and $\mathrm{A} \in \mathcal{B}(\mathbb{R}^d)$, we define the truncated harmonic mean estimator of $I_{\mathrm{A}}(f,y)$ by

onicmean_est1

$$\hat{I}_{\mathrm{A}}(f,y) = \sum_{k=1}^{n} \frac{\mathbb{1}_{\mathrm{A}}(X_k)}{f(X_k,y)} \ , \tag{34}$$

where $(X_k)_{k\geq 1}$ is an ergodic Markov chain targeting $p(x|y)$ to ensure that the defined estimator almost surely converges to $I_{\mathrm{A}}(f,y)$ given by (32).

Let p_1, p_2 be two positive distributions on $\mathbb{R}^d \times \mathbb{R}^m$, associated with their two unormalized versions $f_1, f_2 : \mathbb{R}^d \times \mathbb{R}^m \to \mathbb{R}_+$. We aim to estimate $p_1(y)/p_2(y)$. By (33), we have

$$\frac{p_1(y)}{p_2(y)} = \frac{\int_{\mathbb{R}^d\times\mathbb{R}^m} f_2(z,w)\mathrm{d}z\mathrm{d}w}{\int_{\mathbb{R}^d\times\mathbb{R}^m} f_1(z,w)\mathrm{d}z\mathrm{d}w}\frac{I_{\mathrm{A}}(f_2,y)}{I_{\mathrm{A}}(f_1,y)}$$

Using (34), we estimate this ratio by

$$\frac{p_1(y)}{p_2(y)} \approx \hat{B}_{1,2}(y) = \frac{\int_{\mathbb{R}^d\times\mathbb{R}^m} f_2(z,w)\mathrm{d}z\mathrm{d}w}{\int_{\mathbb{R}^d\times\mathbb{R}^m} f_1(z,w)\mathrm{d}z\mathrm{d}w}\frac{\hat{I}_{\mathrm{A}}(f_2,y)}{\hat{I}_{\mathrm{A}}(f_1,y)} \ .$$

However, we need to compute the ratio $\int_{\mathbb{R}^d\times\mathbb{R}^m} f_2(z,w)\mathrm{d}z\mathrm{d}w / \int_{\mathbb{R}^d\times\mathbb{R}^m} f_1(z,w)\mathrm{d}z\mathrm{d}w$.

Assume that for $i = 1, 2$, $f_i(x,y) = h_i(x,y)g_i(x)$, for some measurable functions $h_i : \mathbb{R}^d \times \mathbb{R}^m \to \mathbb{R}_+^*$, $g_i : \mathbb{R}^d \to \mathbb{R}_+^*$ such that $\int_{\mathbb{R}^m} h_i(x,y)\mathrm{d}y$ does not depend on x. Note that this assumption holds in Section 4.1.3. We distinguish two cases:

1. If for $i = 1, 2$, g_i is integrable, we get

$$\hat{B}_{1,2}(y) = \frac{\int_{\mathbb{R}^d} g_2(z)\mathrm{d}z}{\int_{\mathbb{R}^d} g_1(z)\mathrm{d}z}\frac{\hat{I}_{\mathrm{A}}(f_2)}{\hat{I}_{\mathrm{A}}(f_1)} \ .$$

In the case where the ratio $\int_{\mathbb{R}^d} g_2(z)\mathrm{d}z / \int_{\mathbb{R}^d} g_1(z)\mathrm{d}z$ is unknown, such as with the priors considered in the experiment reported in Section 4.1.3, we use a Monte Carlo algorithm

such as MYULA or Px-MALA to compute it. Observe that this computation can be performed offline when the ratio does not depend on the value of y.

2. If there exists a function $g : \mathbb{R}^d \to \mathbb{R}_+^*$ and two real numbers $\lambda_1, \lambda_2 > 0$ such that for $i = 1, 2$, $g_i(x) = g(\lambda_i x)$ for all $x \in \mathbb{R}^d$, we get for all $R > 0$

$$\int_{\mathbb{R}^d \times \mathbb{R}^m} \mathbb{1}_{\mathrm{B}(0,R)} f_2(z,w)\mathrm{d}z\mathrm{d}w \Big/ \int_{\mathbb{R}^d \times \mathbb{R}^m} \mathbb{1}_{\mathrm{B}(0,\lambda_1\lambda_2^{-1}R)} f_1(z,w)\mathrm{d}z\mathrm{d}w$$
$$= \int_{\mathbb{R}^d} \mathbb{1}_{\mathrm{B}(0,R)} g_2(z)\mathrm{d}z \Big/ \int_{\mathbb{R}^d} \mathbb{1}_{\mathrm{B}(0,\lambda_1\lambda_2^{-1}R)} g_1(z)\mathrm{d}z = (\lambda_1/\lambda_2)^d \, .$$

Since for all $a > 0$ and $i = 1, 2$,

$$\int_{\mathbb{R}^d \times \mathbb{R}^m} f_i(z,w)\mathrm{d}z\mathrm{d}w = \lim_{R \to +\infty} \int_{\mathbb{R}^d \times \mathbb{R}^m} \mathbb{1}_{\mathrm{B}(0,aR)} f_i(z,w)\mathrm{d}z\mathrm{d}w \, ,$$

we get

$$\int_{\mathbb{R}^d \times \mathbb{R}^m} f_2(z,w)\mathrm{d}z\mathrm{d}w \Big/ \int_{\mathbb{R}^d \times \mathbb{R}^m} f_1(z,w)\mathrm{d}z\mathrm{d}w = (\lambda_1/\lambda_2)^d \, .$$

proper-imp˙prior

C.2. Case of improper prior distributions. Let $f : \mathbb{R}^d \times \mathbb{R}^m \to \mathbb{R}_+$ such that for all $y \in \mathbb{R}^m$,

rginal˙improper

$$\tilde{p}(y) = \int_{\mathbb{R}^d} f(x,y)\mathrm{d}x < +\infty \, . \tag{35}$$

Here, f plays the role of an improper joint density of the data and the parameters as the prior distribution is improper. This setting corresponds to Section 4.1.2. Define for all $y \in \mathbb{R}^m$ the conditional distribution on $\mathbb{R}^d \times \mathbb{R}^m$ by $p(x|y) = f(x,y)/\tilde{p}(y)$, where $\tilde{p}$ is defined by (35). Then, define for any bounded Borel set $\mathrm{A} \in \mathcal{B}(\mathbb{R}^d)$

cmean˙improper

$$I_{\mathrm{A}}(f,y) = \int_{\mathbb{R}^d} \mathbb{1}_{\mathrm{A}}(x) \frac{p(x|y)}{f(x,y)} \mathrm{d}x \, . \tag{36}$$

Then by (35), we get

mean˙improper

$$\tilde{p}(y) = \mathrm{Vol}(\mathrm{A})/I_{\mathrm{A}}(f,y) \quad . \tag{37}$$

For all $y \in \mathbb{R}^m$ and $\mathrm{A} \in \mathcal{B}(\mathbb{R}^d)$, we define the truncated harmonic mean estimator of $I_{\mathrm{A}}(f,y)$ as in Appendix C.1 by (34).

Let now $f_1, f_2 : \mathbb{R}^d \times \mathbb{R}^m \to \mathbb{R}_+$, satisfying for all $i = 1, 2$ and $y \in \mathbb{R}^m$, $\tilde{p}_i(y) = \int_{\mathbb{R}^d} f_i(x,y)\mathrm{d}x < +\infty$. We aim to estimate $\tilde{p}_1(y)/\tilde{p}_2(y)$. But by (37), we have

$$\frac{\tilde{p}_1(y)}{\tilde{p}_2(y)} = \frac{I_{\mathrm{A}}(f_2,y)}{I_{\mathrm{A}}(f_1,y)} \, .$$

Using (36) and (34), we estimate this ratio by

$$\frac{\tilde{p}_1(y)}{\tilde{p}_2(y)} \approx \hat{B}_{1,2}(y) = \frac{\hat{I}_{\mathrm{A}}(f_2,y)}{\hat{I}_{\mathrm{A}}(f_1,y)} \, .$$