



# Asymptotically Exact Data Augmentation: Models, Properties, and Algorithms

Maxime Vono<sup>a</sup> , Nicolas Dobigeon<sup>a</sup> , and Pierre Chainais<sup>b</sup> 

<sup>a</sup>University of Toulouse, IRIT/INP-ENSEEHT, Toulouse, France; <sup>b</sup>University of Lille, Centrale Lille, UMR CNRS 9189—CRISTAL, Lille, France

## ABSTRACT

Data augmentation, by the introduction of auxiliary variables, has become an ubiquitous technique to improve convergence properties, simplify the implementation or reduce the computational time of inference methods such as Markov chain Monte Carlo ones. Nonetheless, introducing appropriate auxiliary variables while preserving the initial target probability distribution and offering a computationally efficient inference cannot be conducted in a systematic way. To deal with such issues, this article studies a unified framework, coined asymptotically exact data augmentation (AXDA), which encompasses both well-established and more recent approximate augmented models. In a broader perspective, this article shows that AXDA models can benefit from interesting statistical properties and yield efficient inference algorithms. In non-asymptotic settings, the quality of the proposed approximation is assessed with several theoretical results. The latter are illustrated on standard statistical problems. Supplementary materials including computer code for this article are available online.

## ARTICLE HISTORY

Received July 2019  
Revised September 2020

## KEYWORDS

Approximation; Auxiliary variables; Bayesian inference; Divide-and-conquer; Robustness

## 1. Introduction

Starting at least from the 1960s with the seminal paper of Hartley (1958) on the expectation-maximization (EM) algorithm, introducing auxiliary variables has been a widely adopted strategy to derive iterative algorithms able to deal with possibly complicated inference problems. Indeed, either by coming from statistical physics (Swendsen and Wang 1987) or by the broad statistical community (Dempster, Laird, and Rubin 1977), auxiliary (also called latent) variables have been used to improve (Duane et al. 1987; Edwards and Sokal 1988; Marnissi et al. 2018) and/or simplify (Tanner and Wong 1987; Doucet, Godsill, and Robert 2002) inference methods, such as maximum likelihood (ML) estimation or simulation-based ones. Insightful reviews of these methods were conducted by Besag and Green (1993), van Dyk and Meng (2001), and Tanner and Wong (2010). Among many others, slice sampling and half-quadratic (HQ) methods are archetypal instances of such auxiliary variable-based methods. These methods, by introducing auxiliary variables, appear to be an interesting alternative when sampling cannot be performed directly from a target distribution  $\pi$ . Nonetheless, the superiority of simulation-based algorithms based on data augmentation (DA) over classical Markov chain Monte Carlo (MCMC) methods without DA is not obvious as pointed out by Polson (1996) and Damien, Wakefield, and Walker (1999). DA methods have been found to be slower than single-site update approaches in some cases (Hurn 1997) and some improvements have been derived to cope with these problems such as partial decoupling (Higdon 1998) or the introduction of a working parameter (Meng and van Dyk 1997). Moreover, DA techniques are often used on a case-by-case basis (Geman and Reynolds 1992; Albert and Chib 1993;

Geman and Yang 1995; Polson, Scott, and Windle 2013) and could not be applied in general scenarios due to the absence of exact DA schemes yielding an efficient inference and low computation costs.

Similarly to approximate Bayesian computation (ABC) methods to circumvent intractable likelihoods (Beaumont, Zhang, and Balding 2002; Sisson, Fan, and Beaumont 2018b), these limitations can be tackled by considering approximate DA schemes that become exact asymptotically. For instance, inspired from the variable splitting technique used in the alternating direction method of multipliers (ADMM) (Boyd et al. 2011), Vono, Dobigeon, and Chainais (2019) and Rendell et al. (2020) recently and independently proposed a novel and broad Bayesian inference framework that can circumvent limitations of exact DA approaches. By introducing a collection of instrumental (also called *splitting*) variables, the aforementioned authors considered the inference from an approximate probability distribution which can be simpler, more efficient and distributed over multiple computational workers (e.g., machines or kernels).

This article aims at deeply investigating a broad framework coined asymptotically exact data augmentation (AXDA) which encompasses previously proposed special instances such as approximate models used in Vono, Dobigeon, and Chainais (2019), Rendell et al. (2020), among others. More precisely, Section 2 details how such models can be built in a quasi-systematic and simple way which is highly appreciable compared to the case-by-case search of computationally efficient DA schemes. In Section 3, we revisit some already-proposed special instances of AXDA models to show the potential benefits of AXDA on specific examples and to exhibit interesting properties

which can be generally inherited by AXDA approaches. In [Section 4](#), we assess quantitatively the bias of AXDA models with nonasymptotic theoretical results by considering Wasserstein and total variation distances. Then, [Section 5](#) illustrates the previous theoretical results and the benefits of the proposed methodology on several statistical problems. To facilitate the use of AXDA, we eventually point out that the supplementary materials involve a dedicated section ([Section 5](#)) presenting how such models can be instantiated to perform efficient inference through classical simulation-based, variational Bayes (VB), optimization or expectation-maximization (EM) methods. The proofs are also given in the supplementary materials, see [Section 1](#).

## 2. Asymptotically Exact Data Augmentation

This section introduces AXDA schemes that aim to circumvent exact DA main issue: *the art* (van Dyk and Meng 2001) of finding the exact DA associated to a statistical model and its inference limitations. For sake of simplicity, with little abuse, we shall use the same notations for a probability distribution and its associated probability density function (pdf).

### 2.1. Motivations

In this article, we are interested in performing the inference of a variable of interest  $\theta \in \Theta \subseteq \mathbb{R}^d$ , where  $\Theta$  is a closed convex set and  $\dim(\Theta) = d$ , by relying on a probability distribution with density  $\pi$  writing

$$\pi(\theta) \propto \exp(-f(\theta)), \quad \text{or} \quad \pi(\mathbf{y}|\theta) \propto \exp(-f(\mathbf{y};\theta)), \quad (1)$$

where the potential  $f$  taking values in  $\mathbb{R}$  is such that  $\pi$  defines a proper, bounded and continuous probability distribution. For sake of generality, note that  $\pi$  in (1) shall describe various quantities. First, with a little abuse of notations,  $\pi(\theta)$  may simply refer to a pdf associated to the random variable  $\theta$ , for example, its prior distribution  $\pi(\theta)$  or its posterior distribution  $\pi(\theta) \triangleq \pi(\theta|\mathbf{y})$  when referring to a set of observations denoted by  $\mathbf{y}$ . Depending on the problem, we also allow  $\pi$  to stand for a likelihood function  $\pi(\mathbf{y}|\theta)$ . We will work under this convention and write explicitly the form of  $\pi$  when required. For sake of simplicity and clarity, only the case corresponding to  $\pi(\theta)$  will be detailed in this section. The application of the proposed methodology to  $\pi(\mathbf{y}|\theta)$  is very similar and can be retrieved by a straightforward derivation.

We consider situations where direct inference from (1) is difficult because intractable or computationally prohibitive. To overcome these issues, an option is to rely on exact DA which introduces some auxiliary variables stacked into a vector  $\mathbf{z} \in \mathcal{Z} \subseteq \mathbb{R}^k$  and defines a new density, simpler to handle, such that

$$\int_{\mathcal{Z}} \pi(\theta, \mathbf{z}) d\mathbf{z} = \pi(\theta). \quad (2)$$

Much research has been devoted to these models to simplify an inference task or to improve the convergence properties of direct inference approaches (e.g., slice sampling and HQ methods introduced in [Section 1](#)). Nonetheless, these approaches have several limitations. Indeed, finding a convenient form for

the augmented density to satisfy (2) while leading to efficient algorithms generally requires some knowledge and can even be impossible in some cases (Geman and Yang 1995). For instance, the mixture representation of a binomial likelihood function based on the Polya-Gamma distribution has been used to derive a promising Gibbs sampler for logistic regression problems (Polson, Scott, and Windle 2013). Nonetheless, even if this algorithm has been proved to be uniformly ergodic by Choi and Hobert (2013), the corresponding ergodicity constant depends exponentially on the number of observations  $n$  and on the dimension of the regression coefficients vector  $d$ .

To tackle these limitations, we propose to relax the constraint (2) and consider an *approximate* DA model. This will permit the choice of an augmented density with more flexibility, fix the issues associated to the initial model and make inference more efficient in some cases. To this purpose, [Section 2.2](#) presents the so-called AXDA framework which embeds approximate DA models controlled by a positive scalar parameter  $\rho$ . These models become asymptotically exact when  $\rho$  tends toward 0. Of course, some assumptions will be required on the approximate augmented density to guarantee a good approximation. The quality of this approximation will be assessed in [Section 4](#) with nonasymptotic theoretical results.

### 2.2. Model

Instead of searching for an exact data augmentation scheme (2), some auxiliary variables  $\mathbf{z}$  can be introduced to define an approximate but asymptotically exact probability distribution. One possibility is to introduce an augmented distribution depending on a parameter  $\rho > 0$  and such that the associated marginal density defined by

$$\pi_{\rho}(\theta) = \int_{\mathcal{Z}} \pi_{\rho}(\theta, \mathbf{z}) d\mathbf{z}, \quad (3)$$

satisfies the following property.

**Property 1.** For all  $\theta \in \Theta$ ,  $\lim_{\rho \rightarrow 0} \pi_{\rho}(\theta) = \pi(\theta)$ .

By applying Scheffé's lemma (Scheffé 1947), this property yields the convergence in total variation, that is  $\|\pi_{\rho} - \pi\|_{TV} \rightarrow 0$  as  $\rho \rightarrow 0$ . A natural question is: how to choose the augmented density in (3) such that [Property 1](#) is met? In this article, we assume that  $\mathcal{Z} = \Theta$  and investigate AXDA schemes associated to an initial density (1) and defined by the approximate augmented density

$$\pi_{\rho}(\theta, \mathbf{z}) = \pi(\mathbf{z})\kappa_{\rho}(\mathbf{z}, \theta), \quad (4)$$

where  $\kappa_{\rho}$  is such that (4) defines a proper density.

**Remark 1.** When  $\pi$  stands for a product of  $J$  densities, that is  $\pi \propto \prod_{j=1}^J \pi_j$ , the proposed approximate model can naturally be generalized to  $\pi_{\rho}(\theta, \mathbf{z}_{1:J}) \propto \prod_{j=1}^J \pi_j(\mathbf{z}_j)\kappa_{\rho}(\mathbf{z}_j, \theta)$ . Such a generalization will for instance be considered in [Sections 3.1](#) and [3.2](#).

The introduction of the proposed model (4) is aimed at avoiding a case-by-case search of an appropriate augmented approach. Although there might exist other marginal densities

$\pi_\rho$  satisfying [Property 1](#), we restrict our analysis to models where  $\kappa_\rho(\cdot, \theta)$  weakly converges toward the Dirac measure at  $\theta$  as  $\rho \rightarrow 0$  (Aguirregabiria, Hernández, and Rivas 2002). This is a sufficient condition to satisfy [Property 1](#). In the sequel, we will call AXDA any approach based on (4) and satisfying these properties.

A natural choice for  $\kappa_\rho$  is to consider a standard kernel  $K$  (Wand and Jones 1995). Based on the latter, we define for all  $\mathbf{z}, \theta \in \Theta$ ,  $\kappa_\rho(\mathbf{z}, \theta) \propto_\rho \rho^{-d} K(\rho^{-1}(\theta - \mathbf{z}))$  (Dang and Ehrhardt 2012). Beyond standard kernels but motivated by the same idea of measuring the discrepancy between  $\mathbf{z}$  and  $\theta$ , one can also build on divergence functions widely used in the optimization literature to define a potentially asymmetric density  $\kappa_\rho$  such that for all  $\mathbf{z}, \theta \in \Theta$ ,  $\kappa_\rho(\mathbf{z}, \theta) \propto_\rho \exp(-\rho^{-1}\phi(\mathbf{z}, \theta))$  where  $\phi$  is a strictly convex function w.r.t.  $\mathbf{z}$  admitting a unique minimizer  $\mathbf{z}^* = \theta$  (Ben-Tal, Margalit, and Nemirovski 2001; Krichene, Bayen, and Bartlett 2015; Fellows et al. 2019). Specific instances of such potentials are Bregman divergences such as the logistic loss and the Kullback-Leibler divergence, see [Definition 1](#).

**Definition 1 (Bregman divergence).** Let  $\psi$  a continuously differentiable and strictly convex function defined on a closed convex set. The Bregman divergence associated to  $\psi$  is defined by

$$d_\psi(\mathbf{z}, \theta) = \psi(\mathbf{z}) - \psi(\theta) - \nabla \psi(\theta)^T(\mathbf{z} - \theta). \quad (5)$$

Additional details associated to standard kernels and Bregman divergences are given in Section 2 in the supplementary materials.

### 3. Benefits of AXDA by Revisiting Existing Models

Before providing theoretical guarantees for AXDA models, this section proposes to review some important state-of-the-art works from the AXDA perspective described in [Section 2](#). We do not pretend to give new insights about these approaches. We rather use them to illustrate potential benefits that can be gained by resorting to the proposed framework. For sake of clarity, these benefits are directly highlighted in the title of the following sections before being discussed in the latter.

#### 3.1. Tractable Posterior Inference

This first section illustrates how an AXDA approach can alleviate the intractability of an initial posterior distribution  $\pi$  and significantly aid in the computations.

To this purpose, we consider the case where the posterior distribution  $\pi$  is intractable. Such a model for instance appears when  $\pi$  involves a constraint on some set (Liechty, Liechty, and Müller 2009), admits a nonstandard potential function such as the total variation norm (Chambolle et al. 2010; Pereyra 2016; Vono, Dobigeon, and Chainais 2019) or yields complicated conditional posterior distributions (Holmes and Mallick 2003). To simplify the inference, the aforementioned authors have considered special instances of AXDA by relying on an additional level involving latent variables  $\mathbf{z}$ , leading a hierarchical Bayesian model. In these cases, AXDA has been invoked to move a difficulty to the conditional posterior of  $\mathbf{z}$  where it can be dealt with more easily by using standard inference algorithms, see

Section 5 in the supplementary materials for more details. The following example, derived from Holmes and Mallick (2003), illustrates this idea.

**Example 1.** Let  $\mathbf{y} \in \mathbb{R}^n$  be a set of observations and  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T \in \mathbb{R}^{n \times d}$  a design matrix filled with covariates. We consider a generalized nonlinear model which writes

$$y_i | \theta \sim p(y_i | g^{-1}(h(\mathbf{x}_i, \theta)), \sigma^2), \quad \forall i \in [n], \quad (6)$$

$$\theta \sim \mathcal{N}(\theta | \mathbf{0}_d, v^2 \mathbf{I}_d), \quad (7)$$

where  $p$  belongs to the exponential family and has mean  $g^{-1}(h(\mathbf{x}_i, \theta))$  and variance  $\sigma^2$  where  $g$  is a link function. As in classical regression problems, we are interested in inferring the regression coefficients  $\theta \in \mathbb{R}^d$ . In the sequel, we set the nonparametric model  $h$  to be

$$h(\mathbf{x}_i, \theta) = \sum_{j=1}^k \theta_j B(\mathbf{x}_i, \mathbf{k}_j), \quad (8)$$

where  $B(\mathbf{x}_i, \mathbf{k}_j)$  is a nonlinear function of  $\mathbf{x}_i$  (e.g., regression splines) and  $\mathbf{k}_j$  is the knot location of the  $j$ th basis. The difficulty here is the nonlinearity of  $h$  which, combined with the non-Gaussian likelihood, rules out the use of efficient simulation schemes to sample from the posterior  $\pi(\theta | \mathbf{y})$ . To mitigate this issue, Holmes and Mallick (2003) proposed to rely on an additional level which boils down to consider the approximate model (4). More specifically, the aforementioned authors treated the nonlinear predictor  $h$  as a Gaussian random latent variable which leads to the approximate model


$$y_i | z_i \sim p(y_i | g^{-1}(z_i), \sigma^2), \quad \forall i \in [n], \quad (9)$$

$$z_i | \theta \sim \mathcal{N}(z_i | h(\mathbf{x}_i, \theta), \rho^2), \quad \forall i \in [n], \quad (10)$$

$$\theta \sim \mathcal{N}(\theta | \mathbf{0}_d, v^2 \mathbf{I}_d). \quad (11)$$

Here, AXDA has been applied only to the likelihood function with  $\kappa_\rho$  chosen as the univariate normal distribution (10) leading to a smoothed likelihood function. The main advantage of relying on such a model is that the posterior conditional distribution  $\pi_\rho(\theta | \mathbf{z}, \mathbf{X})$ , with  $\mathbf{z} = [z_1, \dots, z_n]^T$ , is now a multivariate normal distribution. In addition, by moving the difficulty induced by  $h$  to the conditional posterior of  $z_i$ , we are now dealing with a generalized linear model (GLM) where standard techniques can be applied (Albert and Chib 1993; Polson, Scott, and Windle 2013).

Beyond the widely used Gaussian choice for  $\kappa_\rho$  (Holmes and Mallick 2003; Liechty, Liechty, and Müller 2009; Barbos et al. 2017; Vono, Dobigeon, and Chainais 2019), more general AXDA approaches can be built by taking inspiration from these works. To this purpose, we recommend to choose  $\kappa_\rho$  w.r.t. the prior and likelihood at stake. For instance, when a Poisson likelihood function and a complex prior distribution on its intensity  $\theta$  are considered, one option for  $\phi$  (see [Section 2.2](#)) would be an Itakura–Saito divergence since it preserves the positivity constraint on  $\theta$  and yields the well-known Gamma-Poisson model (Canny 2004).



**Figure 1.** Concept of localization. Comparison between the initial (left) and the localized hierarchical Bayesian (right) models with  $n$  the number of observations  $y_i$ .

### 3.2. Distributed Inference

When data are stored on multiple machines and/or one is interested in respecting their privacy, this section illustrates how AXDA can be resorted to perform distributed computations.

Let consider observed data  $\{y_i, \mathbf{x}_i\}_{i=1}^n$ , where  $\mathbf{x}_i$  stands for the covariates associated to observation  $y_i$ , which are distributed among  $B$  nodes within a cluster. By adopting a prior  $\nu(\theta)$  and by assuming that the likelihood can be factorized w.r.t. the  $B$  nodes, the posterior distribution of the variable of interest  $\theta$  writes

$$\pi(\theta | \mathbf{y}, \mathbf{X}) \propto \nu(\theta) \prod_{b=1}^B \prod_{i \in \text{node } b} \exp(-f_i(y_i; h(\mathbf{x}_i, \theta))). \quad (12)$$

Such models classically appear in statistical machine learning when GLMs (Dobson and Barnett 2008) are considered. In these cases,  $h(\mathbf{x}_i, \theta) = \mathbf{x}_i^T \theta$ . Due to the distributed environment, sampling efficiently from (12) is challenging and a lot of “divide-and-conquer” approaches have been proposed in the past few years to cope with this issue (Wang and Dunson 2013; Scott et al. 2016). These methods launch independent Markov chains on each node  $b$  and then combine the outputs of these local chains to obtain an approximation of the posterior of interest (12). Nonetheless, the averaging schemes used to combine the local chains might lead to poor approximations when  $\pi$  is high-dimensional and non-Gaussian. Instead, considering a special instance of AXDA circumvents the previously mentioned drawbacks by introducing local auxiliary variables on each node such that

$$\pi_\rho(\theta, \mathbf{z} | \mathbf{y}, \mathbf{X}) \propto \nu(\theta) \prod_{b=1}^B \prod_{i \in \text{node } b} \exp(-f_i(y_i; z_i)) \kappa_\rho(z_i, h(\mathbf{x}_i, \theta)). \quad (13)$$

The posterior distribution of the auxiliary variables conditionally to  $\theta$  only depends on the data available at a given node. Based on this nice property, the joint posterior can be sampled efficiently with a Gibbs sampler, see Rendell et al. (2020) for a comprehensive review. We emphasize that the benefits described in this section for Monte Carlo sampling also hold when one wants to use other types of algorithms (e.g., expectation-maximization or variational Bayes), see Section 5 in the supplementary materials.

### 3.3. Robust Inference

By noting that classical robust hierarchical models fall into the proposed framework, this section shows that AXDA is also a relevant strategy to perform robust inference by coping with model misspecification by modeling additional sources of uncertainty.

Considering a well-chosen *demarginalization* procedure is known to yield robustness properties in some cases (Robert and Casella 2004). Some approaches took advantage of this idea to build robust hierarchical Bayesian models w.r.t. possible outliers in the data. For instance, such models can be built by allowing each observation to be randomly drawn from a local statistical model, as described in the recent review of Wang and Blei (2018). This “localization” idea is illustrated in Figure 1. Many of these models can be viewed as particular instances of AXDA. Indeed, assume that  $n$  data points  $y_i$  are independently and identically distributed (iid) defining the likelihood function

$$\pi(\mathbf{y} | \theta) \propto \prod_{i=1}^n \pi(y_i | \theta), \quad (14)$$

where  $\theta \in \Theta$  is a common parameter. Applying AXDA as described in Section 2 by introducing  $n$   $d$ -dimensional auxiliary variables stacked into the vector  $\mathbf{z}_{1:n}$  leads to the augmented likelihood

$$\pi_\rho(\mathbf{y}, \mathbf{z}_{1:n} | \theta) \propto \prod_{i=1}^n \pi(y_i | z_i) \kappa_\rho(z_i, \theta). \quad (15)$$

The statistical model defined by (15) implies a hierarchical Bayesian model similar to the localized one depicted in Figure 1b and corresponds in general to an approximation of the initial one, see Example 2.

**Example 2 (Robust logistic regression).** Assume that for all  $i \in [n]$ ,  $\pi(y_i | \theta) = \mathcal{B}(\sigma(\mathbf{x}_i^T \theta))$ , where  $\mathcal{B}$  stands for the Bernoulli distribution,  $\sigma$  for the sigmoid function,  $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$  for the transpose of the design matrix and  $\theta$  for the regression coefficients vector to infer. Then as proposed by Wang and Blei (2018), one can robustify the inference by assuming that each observation  $y_i$  is drawn from a local and independent model  $\mathcal{B}(\sigma(\mathbf{x}_i^T z_i))$  associated to an auxiliary parameter  $z_i \sim \mathcal{N}(\theta, \rho^2 \mathbf{I}_d)$ . In this case,  $\kappa_\rho(\mathbf{z}, \theta) \propto \prod_{i=1}^n \mathcal{N}(z_i | \theta, \rho^2 \mathbf{I}_d)$ .

Beyond the convenient Gaussian prior  $\kappa_\rho$ , its choice can be motivated by robust loss functions. In the statistical machine learning literature, the absolute or Huber losses are of common use (She and Owen 2011). In Bayesian linear inverse problems considered in the signal processing community, it is classical to approximate a complicated forward physical model to yield tractable computations. If the latter can be written as  $\mathbf{y} = h(\theta) + \epsilon$ , with  $\epsilon \sim \pi(\epsilon)$ , then introducing a latent variable  $\mathbf{z} \sim \kappa_\rho(\mathbf{z}, h(\theta))$  such that  $\mathbf{y} = \mathbf{z} + \epsilon$  allows to take into consideration the model approximation. In those cases, one can set  $\kappa_\rho$  to be the distribution of the modeling error which could be adjusted thanks to some expertise.



### 3.4. Inheriting Sophisticated Inference Schemes From ABC

Finally, this section shows that AXDA models, by sharing strong connections with ABC, might inherit sophisticated algorithms to sample from (4).

ABC stands for a family of methods that permit to cope with intractable likelihoods by sampling from the latter instead of evaluating them. In a nutshell, if one's goal is to infer a parameter  $\theta$  based on a posterior of interest, the simplest ABC rejection sampler is as follows. At iteration  $t$ , draw a candidate  $\theta^{(t)}$  from the prior, generate pseudo-observations  $\mathbf{z}$  from the likelihood given this candidate and accept  $\theta^{(t)}$  if  $\mathbf{z} = \mathbf{y}$  where  $\mathbf{y}$  is the observations vector. Many more sophisticated ABC samplers have been derived. We refer the interested reader to the recent review by Sisson, Fan, and Beaumont (2018a) for more information about ABC methods.

Among a huge literature on ABC (also called likelihood-free) methods, *noisy ABC* approaches proposed and motivated by Fearnhead and Prangle (2012) and Wilkinson (2013) are strongly related to AXDA. Indeed, only comparing the underlying models, AXDA with observation splitting is equivalent to noisy ABC. To see this, let  $\pi(\mathbf{y}|\theta)$  stand for an intractable likelihood. Noisy ABC replaces the exact inference based on  $\pi$  by considering the pseudo-likelihood with density

$$\pi_\rho(\mathbf{y}|\theta) \triangleq \int_{\Theta} \pi_\rho(\mathbf{y}, \mathbf{z}|\theta) d\mathbf{z} = \int_{\Theta} \pi(\mathbf{z}|\theta) \kappa_\rho(\mathbf{z}, \mathbf{y}) d\mathbf{z}. \quad (16)$$

This density has exactly the same formulation as the one defined in (4) except that noisy ABC splits the observations  $\mathbf{y}$  instead of the parameter of interest  $\theta$ . Capitalizing on this equivalence property, also pointed out by Rendell et al. (2020), one can derive efficient algorithms for AXDA from the ABC framework. For instance, Rendell et al. (2020) recently built on the works of Beaumont, Zhang, and Balding (2002) and Del Moral, Doucet, and Jasra (2012) in the ABC context to propose a bias correction approach and a sequential Monte Carlo (SMC) algorithm avoiding the tuning of the tolerance parameter  $\rho$ . Obviously, many other inspirations from ABC can be considered, such as the parallel tempering approach of Baragatti, Grimaud, and Pommeret (2013) among others, to make the inference from an AXDA model more flexible and efficient.

## 4. Theoretical Guarantees

By building on existing approaches, Section 3 showed that AXDA can be used in quite general and different settings depending on one's motivations. To further promote the use of such approximate augmented models, this section goes beyond the empirical bias analysis performed by previous works and provides quantitative bounds on the error between the initial and the approximate model. More precisely, for a fixed tolerance parameter  $\rho > 0$ , nonasymptotic results on the error associated to densities, potentials and credibility regions are derived. We will assume all along this section that  $\Theta = \mathbb{R}^d$ . The proofs of the results of this section can be found in Section 1 of the supplementary materials.

### 4.1. Results for Standard Kernels

In this section, we consider the case  $\kappa_\rho(\mathbf{z}, \theta) \propto \rho^{-d} K(\rho^{-1}(\theta - \mathbf{z}))$  where  $K$  is a kernel, see Section 2.2. Under this model,  $\pi_\rho$  stands for the convolution of  $\pi$  and  $\kappa_\rho$  and the following results hold.

**Proposition 1.** Let  $\pi \in L^1$ . The marginal with density  $\pi_\rho$  in (3) has the following properties.

- (i) Let  $\pi$  stand for a pdf associated to the random variable  $\theta$  and  $\mathbb{E}_{\kappa_\rho}(X) = 0$ . Then, the expectation and variance under  $\pi_\rho$  are given by

$$\mathbb{E}_{\pi_\rho}(\theta) = \mathbb{E}_\pi(\theta), \quad (17)$$

$$\text{var}_{\pi_\rho}(\theta) = \text{var}_\pi(\theta) + \text{var}_{\kappa_\rho}(\theta). \quad (18)$$

- (ii)  $\text{supp}(\pi_\rho) \subseteq S$  where  $S$  is the closure of  $\{\mathbf{x} + \mathbf{z}; \mathbf{x} \in \text{supp}(\pi), \mathbf{z} \in \text{supp}(\kappa_\rho)\}$ . The notation  $\text{supp}(h) = \{\mathbf{x} \in \mathcal{X} \mid h(\mathbf{x}) \neq 0\}$  refers to the support of a function  $h : \mathcal{X} \rightarrow \mathbb{R}$ .
- (iii) If both  $\pi$  and  $\kappa_\rho$  are log-concave, then  $\pi_\rho$  is log-concave.
- (iv) If  $\kappa_\rho \in C^\infty(\mathbb{R}^d)$  and  $|\partial^k \kappa_\rho|$  is bounded for all  $k \geq 0$ , then  $\pi_\rho$  is infinitely differentiable w.r.t.  $\theta$ .

**Proposition 1** permits to draw several conclusions about the inference based on  $\pi_\rho$ . First, the infinite differentiability of  $\pi_\rho$  (Property (iv)) implies that it stands for a smooth approximation of  $\pi$ , see Figure 5 in Section 5.2. Second, Property (i) of **Proposition 1** is reassuring regarding the inference task. Indeed, if  $\pi$  stands for a prior distribution, then considering the approximation  $\pi_\rho$  simply corresponds to a more diffuse prior knowledge around the same expected value, see Section 5.2. Thus, more weight will be given to the likelihood if a posterior distribution is derived with this prior. On the other hand, if  $\pi$  stands for a likelihood, then considering the approximation  $\pi_\rho$  yields the opposite behavior: the likelihood becomes less informative w.r.t. the prior. This idea is directly related to robust hierarchical Bayesian models discussed in Section 3.3.

We now provide quantitative bounds on the approximation implied by considering the marginal  $\pi_\rho$  instead of  $\pi$ . For  $p \geq 1$ , we define the  $p$ -Wasserstein distance between  $\pi$  and  $\pi_\rho$  by

$$W_p(\pi, \pi_\rho) = \left( \min_{\mu} \left\{ \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \|\theta - \mathbf{z}\|_2^p d\mu(\mathbf{z}, \theta); \mu \in \Gamma(\pi_\rho, \pi) \right\} \right)^{1/p}, \quad (19)$$

where  $\Gamma(\pi_\rho, \pi)$  is the set of probability distributions  $\mu(\theta, \mathbf{z})$  with marginals  $\pi_\rho$  and  $\pi$  w.r.t.  $\theta$  and  $\mathbf{z}$ , respectively. Under mild assumptions on the kernel  $K$ , **Proposition 2** gives a simple and practical upper bound on (19).

**Proposition 2.** Assume that  $\pi_\rho$  in (3) stands for a pdf associated to the variable  $\theta$ . Let  $p \geq 1$  such that  $m_p \triangleq \left( \int_{\mathbb{R}^d} \|\mathbf{u}\|_2^p K(\mathbf{u}) d\mathbf{u} \right)^{1/p} < \infty$ . Then, we have

$$W_p(\pi, \pi_\rho) \leq \rho m_p. \quad (20)$$

Note that (20) holds without assuming additional assumptions on the initial density  $\pi$  such as infinite differentiability.

**Table 1.** Closed-form expressions of  $m_2$  appearing in (20) for multivariate generalizations of standard kernels where  $d$  denotes the dimension.

	Gaussian	Cauchy	Laplace	Dirichlet	Uniform	Triangular	Epanechnikov
$m_2$	$\sqrt{d}$	–	$\sqrt{2d}$	–	$\sqrt{d/3}$	$\sqrt{d/6}$	$\sqrt{d/5}$

If the latter is assumed w.r.t. the parameter of interest  $\theta$ , then one can estimate the bias  $\pi - \pi_\rho$  with a Taylor expansion of  $\pi$  similarly to bias analysis in ABC, see Sisson, Fan, and Beaumont (2018b). Table 1 gives closed-form expressions of  $m_2$  for the multivariate generalizations of standard kernels. One can see that the constant  $m_2$  has the same dependence w.r.t. the dimension  $d$  for the considered standard kernels  $K$ . Hence, in high-dimensional scenarios, the approximation quality will be more affected by an inappropriate value for the tolerance parameter  $\rho$  rather than by the choice of  $K$ . In Section 5, we illustrate Proposition 2 with numerical experiments.

#### 4.2. Pointwise Bias for Bregman Divergences

In complement to Section 4.1 where  $\kappa_\rho$  was built using kernels, we now analyze the bias induced by considering  $\pi_\rho$  when  $\kappa_\rho$  is derived from a Bregman divergence  $d_\psi$  (see Definition 1), that is

$$\kappa_\rho(\mathbf{z}, \theta) \propto \exp\left(-\frac{d_\psi(\mathbf{z}, \theta)}{\rho}\right). \quad (21)$$

Under regularity assumptions on both  $\pi$  and  $\kappa_\rho$ , Proposition 3 shows the dependence of the pointwise bias  $\pi_\rho - \pi$  w.r.t. to the tolerance parameter  $\rho$  when the latter is sufficiently small.

**Proposition 3.** Assume that  $\pi$  is analytic and twice differentiable on  $\mathbb{R}^d$  and so does  $d_\psi$  w.r.t. its first argument. Let  $\theta \in \mathbb{R}^d$  such that both  $\mathbf{H}_\pi(\theta)$  and  $\mathbf{H}_{d_\psi}(\theta)^{-1}$  exist and are continuous, where  $\mathbf{H}_\pi(\theta)$  is the Hessian matrix of  $\pi$  and  $\mathbf{H}_{d_\psi}(\theta) \triangleq \frac{\partial^2 d_\psi(\mathbf{z}, \theta)}{\partial \mathbf{z}^2} \Big|_{\mathbf{z}=\theta}$  is the Hessian matrix associated to  $d_\psi(\cdot, \theta)$ . Then, if

- $\|\mathbf{H}_\pi\| \leq C < \infty$
- $\|\mathbf{H}_{d_\psi}\| \geq c > 0$ ,

it follows that

$$\pi_\rho(\theta) - \pi(\theta) = \mathcal{O}(\sqrt{\rho}). \quad (22)$$

In addition, if we have  $\int_{\mathbb{R}^d} \mathbf{u} \kappa_\rho(\theta - \sqrt{\rho} \mathbf{u}, \theta) d\mathbf{u} = \mathbf{0}_d$ , then

$$\pi_\rho(\theta) - \pi(\theta) = \frac{\rho}{2} \text{Trace}(\mathbf{H}_\pi(\theta) \mathbf{H}_{d_\psi}(\theta)^{-1}) + o(\rho). \quad (23)$$

Note that when  $\psi(\mathbf{z}) = \|\mathbf{z}\|_2^2/2$ ,  $\kappa_\rho$  stands for a Gaussian smoothing kernel, see Section 4.1. In that case, we have the sanity check that the dependence w.r.t.  $\rho$  of the bias between  $\pi$  and  $\pi_\rho$  in (23) is the same as the one derived by Sisson, Fan, and Beaumont (2018b) when interpreting  $\kappa_\rho$  as a kernel.

#### 4.3. A Detailed Nonasymptotic Analysis for Gaussian Smoothing

The previous sections gave quantitative approximation results for a large class of densities  $\kappa_\rho$  built either via a kernel or a

Bregman divergence. In this section, we provide complementary results by restricting our analysis on the case

$$\kappa_\rho(\mathbf{z}, \theta) = \mathcal{N}(\mathbf{z}|\theta, \rho^2 \mathbf{I}_d). \quad (24)$$

This particular yet convenient assumption will allow to complement and sharpen results of Section 4.1 by deriving quantitative bounds which take into account the regularity properties of  $f$ . Furthermore, these bounds can be extended to a sum of potential functions  $f = \sum_i f_i$  and used to assess the bias associated to both log-densities and credibility regions. This analysis is also motivated by the fact that the Gaussian smoothing case has been widely advocated in the literature since it generally leads to simple inference steps (Holmes and Mallick 2003; Giovannelli 2008; Liechty, Liechty, and Müller 2009; Dümbgen and Rufibach 2009), and can be related to both the ADMM in optimization (Boyd et al. 2011; Vono, Dobigeon, and Chainais 2019) and the approximation involved in proximal MCMC methods (Pereyra 2016; Durmus, Moulines, and Pereyra 2018; Salim, Koralev, and Richtarik 2019). Unfortunately, a straightforward generalization of the proof techniques used in the sequel does not give informative upper bounds for smoothing associated to other Bregman divergences.

##### 4.3.1. Assumptions

To derive nonasymptotic bounds between quantities related to  $\pi_\rho$  defined in (3) and  $\pi$  in (1), some complementary assumptions on  $f = -\log \pi$  will be required. They are detailed hereafter. For simplicity and with a little abuse of notations, we also denote here by  $f(\theta)$  the potential associated to (1) when  $\pi(\mathbf{y}|\theta)$  stands for a likelihood.

- (A<sub>1</sub>)  $f$  is  $L_f$ -Lipschitz w.r.t.  $\|\cdot\|_2$ , that is  $\exists L_f \geq 0$  such that for all  $\theta, \eta \in \mathbb{R}^d$ ,  $|f(\theta) - f(\eta)| \leq L_f \|\theta - \eta\|_2$ . When  $\pi$  is a likelihood, it is further assumed that  $L_f$  is independent of  $\mathbf{y}$ .
- (A<sub>2</sub>)  $f$  is continuously differentiable and has an  $M_f$ -Lipschitz continuous gradient w.r.t.  $\|\cdot\|_2$ , that is  $\exists M_f \geq 0$  such that for all  $\theta, \eta \in \mathbb{R}^d$ ,  $\|\nabla f(\theta) - \nabla f(\eta)\|_2 \leq M_f \|\theta - \eta\|_2$ .
- (A<sub>3</sub>)  $f$  is convex, that is for every  $\alpha \in [0, 1]$ ,  $\theta, \eta \in \mathbb{R}^d$ ,  $f(\alpha\theta + (1-\alpha)\eta) \leq \alpha f(\theta) + (1-\alpha)f(\eta)$ .
- (A<sub>4</sub>)  $M_f = \int_{\mathbb{R}^d} \|\nabla f(\theta)\|_2^2 \pi(\theta) d\theta < \infty$ .

Assumptions (A<sub>1</sub>), (A<sub>2</sub>), and (A<sub>3</sub>) on the potential  $f$  stand for standard regularity assumptions in the optimization literature and cover a large class of functions  $f$  (Beck and Teboulle 2009). In the broad statistical community, (A<sub>1</sub>) has been used by Durmus, Moulines, and Pereyra (2018) to derive non-asymptotic bounds on the total variation distance between probability distributions while (A<sub>2</sub>) stands for a sufficient condition to have a strong solution to the overdamped Langevin stochastic differential equation (Durmus and Moulines 2017).

Under the previous assumptions (not used all at once), nonasymptotic upper bounds on the total variation distance between  $\pi_\rho$  and  $\pi$  are derived in Section 4.3.2. Then, Sections 4.3.3 and 4.3.4 take advantage of this bound to state theoretical properties on the potential functions and credibility regions.

#### 4.3.2. Nonasymptotic Bounds on the Total Variation Distance

In this section, we make additional regularity assumptions on the potential  $f$  to show quantitative results depending explicitly on regularity constants associated to  $f$ . Two different cases will be considered, namely Lipschitz potentials, and differentiable, gradient-Lipschitz and convex ones.

**4.3.2.1. Lipschitz Potential.** When the potential function  $f$  is assumed to be Lipschitz continuous but not necessarily continuously differentiable, the following result holds.

**Theorem 1.** Let a potential function  $f$  satisfy  $(A_1)$ . Then,

$$\|\pi_\rho - \pi\|_{\text{TV}} \leq 1 - \Delta_d(\rho), \quad (25)$$

where

$$\Delta_d(\rho) = \frac{D_{-d}(L_f \rho)}{D_{-d}(-L_f \rho)}. \quad (26)$$

The function  $D_{-d}$  is a parabolic cylinder function defined for all  $d > 0$  and  $z \in \mathbb{R}$  by

$$D_{-d}(z) = \frac{\exp(-z^2/4)}{\Gamma(d)} \int_0^{+\infty} e^{-xz - x^2/2} x^{d-1} dx. \quad (27)$$

As expected from Property 1, note that this bound tends toward zero when  $\rho \rightarrow 0$ . Additionally, this bound depends on few quantities that can be computed, bounded or approximated in real applications: the dimension of the problem  $d$ , the Lipschitz constant  $L_f$  associated to the regularized potential  $f$  and the tolerance parameter  $\rho$ . In the limiting case  $\rho \rightarrow 0$ , the following equivalent function for the upper bound derived in (25) holds.

**Corollary 1.** In the limiting case  $\rho \rightarrow 0$ , we have

$$\|\pi_\rho - \pi\|_{\text{TV}} \leq \rho L_f \frac{2\sqrt{2}\Gamma\left(\frac{d+1}{2}\right)}{\Gamma\left(\frac{d}{2}\right)} + o(\rho), \quad (28)$$

where for all  $z > 0$  as  $\Gamma(z) = \int_0^{+\infty} x^{z-1} e^{-x} dx$ .

Under some regularity conditions (here Lipschitz continuity) on the potential function  $f$ , Proposition 1 states that  $\|\pi_\rho - \pi\|_{\text{TV}}$  grows at most linearly w.r.t. the parameter  $\rho$  and w.r.t.  $L_f$  when  $\rho$  is sufficiently small. Moreover, using Stirling-like approximations when  $d$  is large in the equivalence relation (28) may give a mild dependence on the dimensionality of the problem in  $\mathcal{O}(L_f d^{1/2})$ . Potential functions verifying the hypothesis of Theorem 1 are common in machine learning and signal/image processing problems, see Section 3 in the online

supplementary materials. As an archetypal example, the sparsity promoting potential function defined for all  $\theta \in \mathbb{R}^d$  by  $f(\theta) = \tau \|\theta\|_1$  with  $\tau > 0$  is Lipschitz continuous with Lipschitz constant  $L_f = \tau\sqrt{d}$  and satisfies Theorem 1 and Proposition 1. In this case, the dependence of (28) is linear w.r.t.  $d$  when  $d$  is large and  $\rho$  is small. Note also that continuously differentiable functions on a compact set are Lipschitz continuous.

**4.3.2.2. Convex and Gradient-Lipschitz Potential.** We now show a complementary result by assuming  $f$  to be convex and continuously differentiable with a Lipschitz-continuous gradient.

**Theorem 2.** Let a potential function  $f$  satisfy  $(A_2)$ ,  $(A_3)$  and  $(A_4)$ . Then, when  $\pi$  stands for a pdf associated to  $\theta$ , we have

$$\|\pi_\rho - \pi\|_{\text{TV}} \leq 1 - \frac{1}{(1 + 2\rho^2 M_f)^{d/2}} \left(1 - \frac{\rho^4 M_f M_f}{1 + 2\rho^2 M_f}\right). \quad (29)$$

In the limiting case  $\rho \rightarrow 0$ , the upper bound in (29) has a simpler expression as shown hereafter.

**Corollary 2.** In the limiting case  $\rho \rightarrow 0$ , we have

$$\|\pi_\rho - \pi\|_{\text{TV}} \leq \rho^2 d M_f + o(\rho^2). \quad (30)$$

Note that the dependences w.r.t. both  $\rho$  and  $d$  in Corollaries 1 and 2 are similar to the ones found by Nesterov and Spokoiny (2017) for optimization purposes.

Figure 2 gives the behavior of the upper bounds in (25) and (29) w.r.t. the dimensionality  $d$  of the problem ranging from 1 to  $10^6$  and as a function of  $\rho$  in log-log scale. The linear (resp. quadratic) relation between this upper bound and  $\rho$  shown in (28) (resp. (30)) is clearly observed for small values of  $\rho$ . Nonetheless, these upper bounds are not a silver bullet. Indeed, as expected, for a fixed value of the parameter  $\rho$ , the approximation error increases as the dimension  $d$  grows. Thus, these bounds suffer from the curse of dimensionality and become non-informative in high-dimension if  $\rho$  is not sufficiently small.


Theorem 1 is easily extended to the case where the initial density  $\pi$  is expressed as a product of several terms. If  $\pi$  stands for the pdf associated to the variable  $\theta$ , this boils down to considering

$$\pi(\theta) \propto \prod_{j=1}^J \pi_j(\theta) \propto \exp\left(-\sum_{j=1}^J f_j(\theta)\right), \quad (31)$$

where for all  $j \in [J]$ ,  $f_j : \mathbb{R}^d \rightarrow \mathbb{R}$ , and a natural generalization of AXDA when applied to each  $\pi_j$ , which writes

$$\begin{aligned} \pi_\rho(\theta, \mathbf{z}_{1:J}) &\propto \prod_{j=1}^J \pi_j(\mathbf{z}_j) \kappa_{\rho_j}(\mathbf{z}_j; \theta) \\ &\propto \exp\left(-\sum_{j=1}^J f_j(\mathbf{z}_j) + \frac{1}{2\rho_j^2} \|\mathbf{z}_j - \theta\|_2^2\right). \end{aligned} \quad (32)$$

Under this product form, we have the following corollary.



**Figure 2.** Behavior of the quantitative bounds shown in [Theorems 1](#) and [2](#) w.r.t.  $\rho$  in log-log scale for a set of dimensions  $d$ . The other quantities appearing in the bounds have been set to 1.

**Corollary 3.** For all  $j \in [J]$ , let  $f_j$  satisfy  $(A_1)$ . Then,

$$\|\pi_\rho - \pi\|_{\text{TV}} \leq 1 - \prod_{j=1}^J \Delta_d^{(j)}(\rho_j), \quad (33)$$

where  $\Delta_d^{(j)}(\rho_j) = D_{-d}(L_{f_j}\rho_j)/D_{-d}(-L_{f_j}\rho_j)$ .

Unfortunately, [Theorem 2](#) cannot be extended to the multiple splitting scenario. We are nevertheless confident that quantitative bounds can be found with different proof techniques but this task goes beyond the scope of this article.

#### 4.3.3. Uniform Bounds on Potentials

From an optimization point of view, it is quite common to consider potential functions associated to densities. For such applications, we give hereafter a quantitative uniform bound on the difference between the potential functions associated to  $\pi$  and  $\pi_\rho$ . Similarly to the definition of the potential function  $f$  in [\(1\)](#), we define the potential function  $f_\rho$  associated to the approximate marginal  $\pi_\rho$  in [\(3\)](#), for all  $\theta \in \mathbb{R}^d$ , by

$$f_\rho(\theta) = -\log \int_{\mathbb{R}^d} \exp(-f(\mathbf{z})) \kappa_\rho(\mathbf{z}, \theta) d\mathbf{z}. \quad (34)$$

By considering a Gaussian smoothing kernel  $\kappa_\rho$ , the potential  $f_\rho$  becomes

$$f_\rho(\theta) = -\log \int_{\mathbb{R}^d} \exp\left(-f(\mathbf{z}) - \frac{1}{2\rho^2} \|\mathbf{z} - \theta\|_2^2\right) d\mathbf{z} + \frac{d}{2} \log(2\pi\rho^2). \quad (35)$$

Note that  $f_\rho(\theta)$  appears as a regularized version of  $f(\theta)$ .

**Proposition 4.** Let  $f$  satisfy  $(A_1)$ . Then, for all  $\theta \in \mathbb{R}^d$ ,

$$L_\rho \leq f_\rho(\theta) - f(\theta) \leq U_\rho, \quad (36)$$

with

$$L_\rho = \log N_\rho - \log D_{-d}(-L_f\rho), \quad (37)$$

$$U_\rho = \log N_\rho - \log D_{-d}(L_f\rho), \quad (38)$$

and

$$N_\rho = \frac{2^{d/2-1} \Gamma(d/2)}{\Gamma(d) \exp(L_f^2 \rho^2 / 4)}. \quad (39)$$

It is easily observed that these bounds are informative in the limiting case  $\rho \rightarrow 0$  since they both tend toward zero.

#### 4.3.4. Uniform Bounds on Credibility Regions

When  $\pi$  stands for the density associated to a posterior distribution, one advantage of Bayesian analysis is its ability to derive the underlying probability distribution of the variable of interest  $\theta$  and thereby to provide credibility information under this distribution. This uncertainty information is particularly relevant and essential for real-world applications. Since the marginal  $\pi_\rho$  stands for an approximation of the original target distribution  $\pi$ , it is important to control the credibility regions under  $\pi_\rho$  w.r.t. those drawn under  $\pi$ . The control in total variation distance given by [Theorem 1](#) is already a good indication. However, it is possible to quantify more precisely the difference between the credible regions ([Robert 2001](#)) with confidence level  $(1 - \alpha)$  under  $\pi_\rho$  and  $\pi$ , as stated below.

**Proposition 5.** Let  $\pi$  be a posterior distribution associated to  $\theta$  and  $f$  such that  $(A_1)$  is verified. Let  $C_\alpha^\rho$  an arbitrary  $(1 - \alpha)$ -credibility region under  $\pi_\rho$ , that is  $\mathbb{P}_{\pi_\rho}(\theta \in C_\alpha^\rho) = 1 - \alpha$  with  $\alpha \in (0, 1)$ . Then,

$$(1 - \alpha) \frac{N_\rho}{D_{-d}(-L_f\rho)} \leq \int_{C_\alpha^\rho} \pi(\theta) d\theta \leq \min\left(1, (1 - \alpha) \frac{N_\rho}{D_{-d}(L_f\rho)}\right), \quad (40)$$

where  $N_\rho$  is defined in [\(39\)](#).

[Proposition 5](#) states that the coverage of  $\pi$  under  $C_\alpha^\rho$  can be determined for a fixed value of  $\rho$ . Thus, it is even possible to obtain a theoretical comprehensive description of  $C_\alpha^\rho$  w.r.t. the initial target density  $\pi$  before conducting an AXDA-based inference. The bounds in [\(40\)](#) permit to choose a parameter  $\rho$  to ensure a prescribed coverage property. The behavior of these bounds w.r.t.  $\rho$  is the same as in [Section 4.3.2](#), that is, linear behavior w.r.t.  $\rho$  when this parameter is sufficiently small.

## 5. Numerical Illustrations

This section illustrates the quantitative results shown in [Sections 4.1](#) and [4.3](#) on three different examples which classically appear in statistical signal processing and machine learning. As shown in [Table 1](#), the bias induced by considering  $\pi_\rho$  is mostly driven by the value of the tolerance parameter  $\rho$  rather than by the choice of  $\kappa_\rho$ . Hence, for simplicity, most of the numerical illustrations hereafter consider the case where  $\kappa_\rho$  is a Gaussian



smoothing kernel. Additional illustrations can be found in the online supplementary materials.

### 5.1. Multivariate Gaussian Example

We start by performing a sanity check with the simple case where  $\pi$  stands for a multivariate Gaussian density that is

$$\pi(\theta) = \mathcal{N}(\theta | \mu, \Sigma), \quad (41)$$

where  $\Sigma$  is assumed to be positive definite. If  $\kappa_\rho(\cdot, \theta)$  is taken to be Gaussian density with mean  $\theta$  and covariance matrix  $\rho^2 \mathbf{I}_d$ , then one can show that

$$\pi_\rho(\theta) = \mathcal{N}(\theta | \mu, \Sigma + \rho^2 \mathbf{I}_d). \quad (42)$$

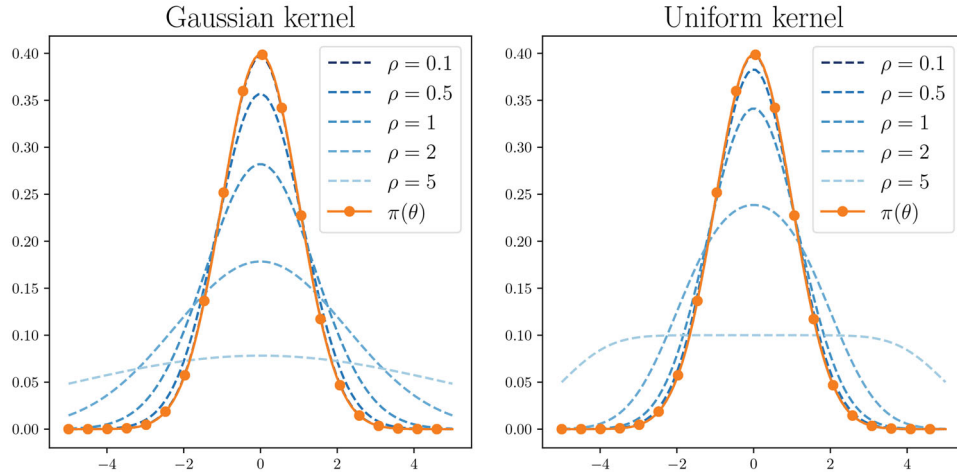
In particular, let consider the univariate setting, that is  $\Theta = \mathbb{R}$ ,  $\Sigma = \sigma^2$ . In this case, the variance under  $\pi_\rho$  is  $\sigma^2 + \rho^2$  and simply corresponds to the variance under  $\pi$  inflated by a factor  $\rho^2$ . Therefore, the approximation will be reasonable if  $\rho^2/\sigma^2$  is sufficiently small, see Figure 3. In this figure, we also show the approximation induced by considering a uniform kernel instead of a Gaussian one. The smoothing via the uniform kernel performs slightly better than Gaussian smoothing due to its lower variance. In both cases, the approximation is reasonable for small  $\rho$  although  $\pi_\rho$ , built with a uniform kernel, no longer belongs to the Gaussian family.

To illustrate the proposed upper bounds on both 2-Wasserstein and total variation distances, we consider a covariance matrix  $\Sigma$  which stands for a squared exponential matrix commonly used in applications involving Gaussian processes (Higdon 2007) and which writes

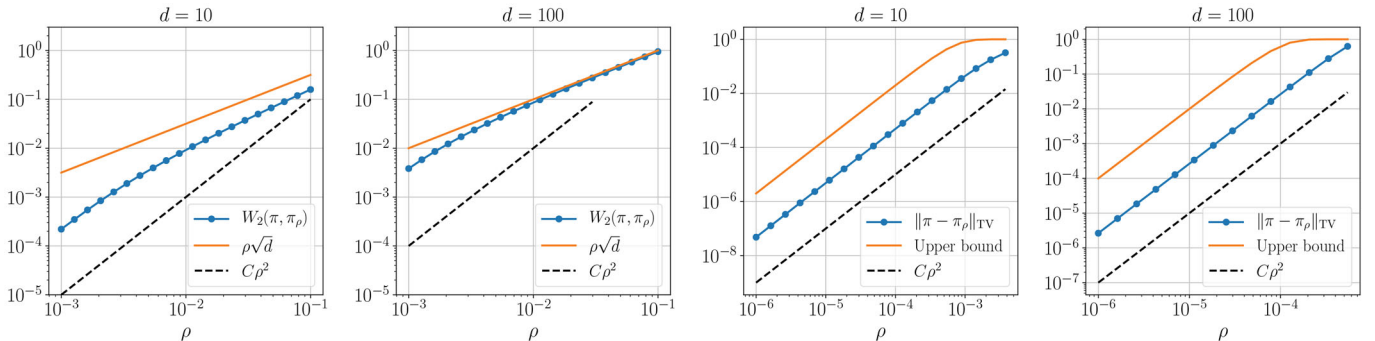
$$\Sigma_{ij} = 2 \exp\left(-\frac{(s_i - s_j)^2}{2a^2}\right) + 10^{-6} \delta_{ij}, \forall i, j \in [d], \quad (43)$$

where  $a = 1.5$ ,  $s_{i,i \in [d]}$  are regularly spaced scalars on  $[-3, 3]$  and  $\delta_{ij} = 1$  if  $i = j$  and zero otherwise.

Figure 4 shows the behavior of the quantitative bounds derived in Proposition 2 and Theorem 2 for  $d \in \{10, 100\}$ . The Gaussian case allows to compute exactly  $W_2(\pi, \pi_\rho)$  by noting that  $W_2^2(\pi, \pi_\rho) = \text{Trace}(\Sigma + \rho^2 \mathbf{I}_d - 2\rho \Sigma^{1/2})$ . On the other hand,  $\|\pi - \pi_\rho\|_{\text{TV}}$  has been estimated by using a Monte Carlo approximation. One can note that the general upper bound on the 2-Wasserstein distance is quite conservative for small  $\rho$  since it does not catch the behavior in  $\mathcal{O}(\rho^2)$  when  $\rho$  is small. This is essentially due to the fact that this bound only assumes a finite moment property and does not require any regularity assumptions on  $\pi$  such as differentiability or strong convexity of its potential. On the contrary, the bound on the total variation distance, derived under stronger assumptions, manages to achieve a rate of the order  $\mathcal{O}(\rho^2)$  for small  $\rho$ .



**Figure 3.** Bias between  $\pi_\rho$  and  $\pi$  in the case  $\Theta = \mathbb{R}$ ,  $\pi = \mathcal{N}(\mu, \sigma^2)$  with  $\mu = 0$  and  $\sigma = 1$ . (left)  $\pi_\rho$  is built with a Gaussian kernel  $\mathcal{N}(0, \rho^2)$  and (right) with a uniform kernel on  $[-\rho, \rho]$ . Note that the curves associated to  $\pi$  and  $\pi_\rho$  for  $\rho = 0.1$  are overlapping.



**Figure 4.** For  $d \in \{10, 100\}$ , illustration of the quantitative bounds (20) and (29) associated to 2-Wasserstein and total variation distances, respectively. The decay in  $\mathcal{O}(\rho^2)$  is shown via the dashed line  $C\rho^2$  where  $C$  is a constant.

## 5.2. Sparse Linear Regression

We study here a generalized version of the least absolute shrinkage and selection operator (lasso) regression problem analyzed by Park and Casella (2008). We assume a standard linear regression problem where centered observations  $\mathbf{y} \in \mathbb{R}^n$  are related to the unknown parameters  $\boldsymbol{\theta} \in \mathbb{R}^d$  via the model  $\mathbf{y} = \mathbf{X}\boldsymbol{\theta} + \boldsymbol{\varepsilon}$ , where  $\mathbf{X} \in \mathbb{R}^{n \times d}$  stands for a known standardized design matrix and  $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}_n, \sigma^2 \mathbf{I}_n)$ . By considering a generalized Laplacian prior distribution for  $\boldsymbol{\theta}$ , the target posterior distribution has density for all  $\boldsymbol{\theta} \in \mathbb{R}^d$ ,

$$\pi(\boldsymbol{\theta}) \triangleq \pi(\boldsymbol{\theta}|\mathbf{y}) \propto \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\theta}\|_2^2 - g(\mathbf{B}\boldsymbol{\theta})\right), \quad (44)$$

where  $g(\mathbf{B}\boldsymbol{\theta}) = \tau \|\mathbf{B}\boldsymbol{\theta}\|_1$  with  $\tau > 0$  and  $\mathbf{B} \in \mathbb{R}^{k \times d}$  an arbitrary matrix acting on  $\boldsymbol{\theta}$ . The choice of such a prior may promote a form of sparsity (lasso). For instance, this matrix  $\mathbf{B}$  might stand for a  $p$ th order difference operator (Bredies, Kunisch, and Pock 2010) which is highly used in signal and image processing problems. As an archetypal example, the case  $p = 1$  leads to the well-known total variation regularization function (Chambolle et al. 2010) used to recover piecewise constant signals.

Note that because of the presence of the matrix  $\mathbf{B}$ , finding an exact data augmentation leading to an efficient sampling scheme is not possible for the general case  $\mathbf{B} \neq \mathbf{I}_d$ . Instead, an AXDA model makes the posterior sampling task possible. Indeed, by regularizing the prior with a Gaussian term, the joint density  $\pi_\rho$  writes

$$\pi_\rho(\boldsymbol{\theta}, \mathbf{z}) \propto \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\theta}\|_2^2 - g(\mathbf{z}) - \frac{1}{2\rho^2} \|\mathbf{B}\mathbf{x} - \mathbf{z}\|_2^2\right). \quad (45)$$

By resorting to a Gibbs algorithm to sample from (45), one can now use a simple data augmentation scheme (Park and Casella 2008) to sample from the  $\mathbf{z}$ -conditional. On the other hand, sampling from the  $\boldsymbol{\theta}$ -conditional, which is a multivariate Gaussian distribution, can be undertaken efficiently with state-of-the-art approaches (Papandreou and Yuille 2010; Barbos et al. 2017; Marnissi et al. 2018).

In this specific case, the potential  $g_\rho$  associated to the smoothed prior distribution (see (35)) has a closed-form expression given for all  $\boldsymbol{\theta} \in \mathbb{R}^d$ , by

$$\begin{aligned} g_\rho(\boldsymbol{\theta}) &= \frac{k}{2} \log(2\pi\rho^2) \\ &\quad - \log \prod_{i=1}^k \int_{\mathbb{R}} \exp\left(-\tau|z_i| - \frac{1}{2\rho^2} (\mathbf{b}_i^T \boldsymbol{\theta} - z_i)^2\right) dz_i \\ &= \frac{k}{2} \log(2\pi\rho^2) \\ &\quad - \log \prod_{i=1}^k (a(\boldsymbol{\theta}) \{ \exp(b(\boldsymbol{\theta})^2) \{1 - \operatorname{erf}(b(\boldsymbol{\theta}))\} \\ &\quad + \exp(c(\boldsymbol{\theta})^2) \{1 - \operatorname{erf}(c(\boldsymbol{\theta}))\} \}) \} \end{aligned} \quad (46)$$

with  $a(\boldsymbol{\theta}) = \sqrt{\pi\rho^2/2} \exp(-(\mathbf{b}_i^T \boldsymbol{\theta})^2/(2\rho^2))$ ,  $b(\boldsymbol{\theta}) = \sqrt{\rho^2/2}(\tau - \mathbf{b}_i^T \boldsymbol{\theta}/\rho^2)$ ,  $c(\boldsymbol{\theta}) = \sqrt{\rho^2/2}(\tau + \mathbf{b}_i^T \boldsymbol{\theta}/\rho^2)$  and  $\mathbf{b}_i \in \mathbb{R}^d$

**Table 2.** Illustration of the bound derived in (40) for the marginal posterior  $\pi_\rho$  depicted in Section 5.2.

$\rho$	$\mathcal{C}_\alpha$	$\mathcal{C}_\alpha^\rho$	$\int_{\mathcal{C}_\alpha^\rho} \pi(\theta_1) d\theta_1$	$\mathcal{I}_\alpha^\rho$
$10^{-3}$	$[-0.47, 1.24]$	$[-0.47, 1.24]$	0.95	$[0.949, 0.951]$
$10^{-2}$	idem	$[-0.47, 1.24]$	0.95	$[0.948, 0.952]$
$10^{-1}$	idem	$[-0.47, 1.24]$	0.95	$[0.88, 1]$
$10^0$	idem	$[-0.47, 1.37]$	0.96	$[0.34, 1]$

NOTE: The  $(1-\alpha)$ -credibility intervals  $\mathcal{C}_\alpha$  and  $\mathcal{C}_\alpha^\rho$  are the highest posterior density regions associated to each density with  $\alpha = 0.05$ .

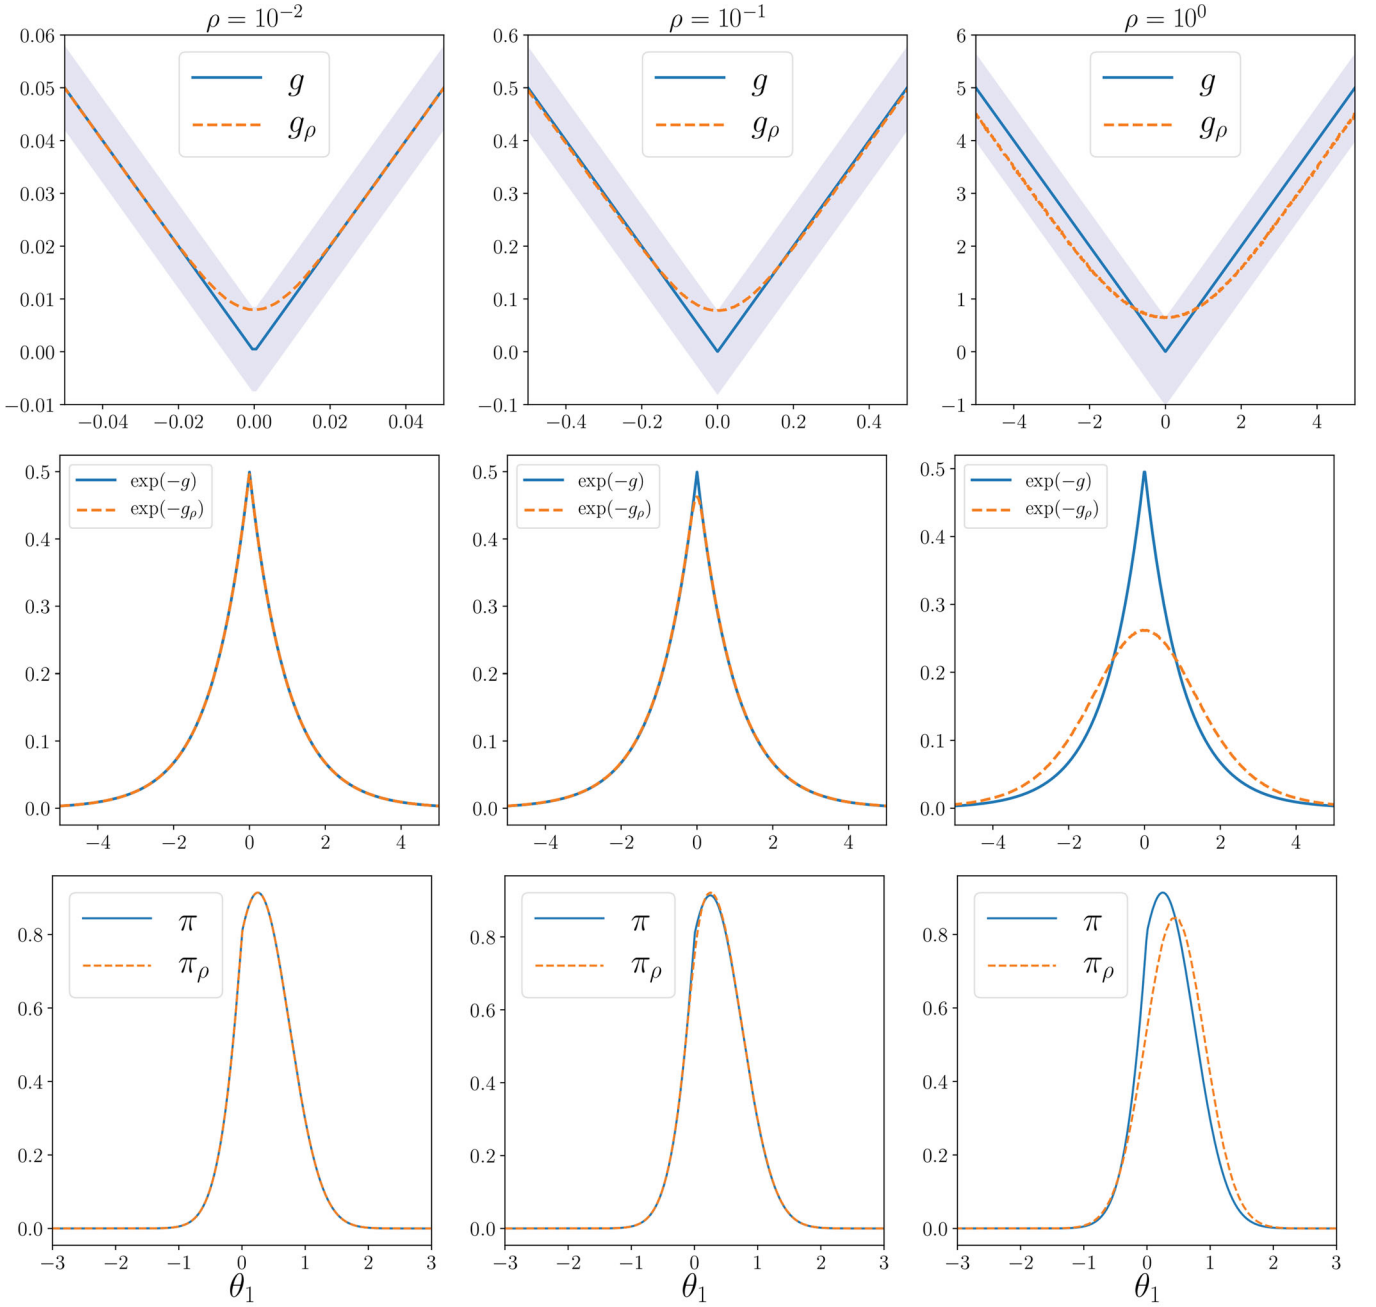
standing for the  $i$ th row of  $\mathbf{B}$ . Note that in more general cases where  $g_\rho$  has no closed form, one can estimate it by a Monte Carlo approximation.

Figure 5 shows the behavior of the regularized potential  $g_\rho$  defined in (46) for several values of the parameter  $\rho$  along with the associated smoothed prior and posterior distributions. For simplicity and pedagogical reasons, the univariate case corresponding to  $\boldsymbol{\theta} = \theta_1 \in \mathbb{R}$  and  $\mathbf{B} = 1$  has been considered. The regularization parameter  $\tau$  has been set to  $\tau = 1$ . The contours of the shaded area correspond to  $g + L_\rho$  and  $g + U_\rho$ . The potential  $g_\rho$  is a smooth approximation of the potential  $g$  associated to the initial prior as expected, see Property (iv) in Proposition 1. Note that the inequalities derived in (36) are verified. Although this approximation seems similar to the Moreau–Yosida regularization of a nonsmooth potential function (Combettes and Pesquet 2011), the rationale behind this approximation is different. Indeed, the Moreau–Yosida envelope stands for a particular instance of the infimal convolution between two convex functions (an initial potential and a Gaussian one). On the other hand,  $g_\rho$  is the potential associated to a smoothed density obtained by convolution with a Gaussian kernel. In addition, the third row of Figure 5 shows the form of the posterior of  $\theta_1$  defined in (45) for  $y = 1$ ,  $x = 2$ , and  $\sigma = 1$  and derived from the smoothed prior distributions shown in Figure 5. For sufficiently small values of  $\rho$ , the marginal  $\pi_\rho$  stands for a quite accurate approximation of the original target  $\pi$ .

Table 2 illustrates the bounds derived in (40) for  $\rho \in \{1, 10^{-1}, 10^{-2}, 10^{-3}\}$ . For each case, the values of the bounds are summarized in the interval

$$\mathcal{I}_\alpha^\rho = [(1-\alpha)N_\rho/D_{-d}(-L_f\rho), \min(1, (1-\alpha)N_\rho/D_{-d}(L_f\rho))], \quad (47)$$

and the real coverage  $\int_{\mathcal{C}_\alpha^\rho} \pi(\theta_1) d\theta_1$  is also reported. The  $(1-\alpha)$ -credibility intervals  $\mathcal{C}_\alpha$  and  $\mathcal{C}_\alpha^\rho$  have been chosen to be the highest posterior density regions associated to each density with  $\alpha = 0.05$ . Note that the theoretical coverage interval  $\mathcal{I}_\alpha^\rho$  becomes informative only if  $\rho$  is sufficiently small which is not surprising since the assumptions on the potential of  $\pi_\rho$  are weak. Indeed, the form of the density (e.g. symmetry or unimodality) is not taken into account in the derived bounds. Regarding the empirical value of the coverage  $\int_{\mathcal{C}_\alpha^\rho} \pi(\theta_1) d\theta_1$ , we emphasize that the marginal  $\pi_\rho$  stands for a conservative approximation of  $\pi$  in this example. Indeed, in each case, the  $(1-\alpha)$ -credibility interval under  $\pi_\rho$  denoted  $\mathcal{C}_\alpha^\rho$  covers at least  $100(1-\alpha)\%$  of the probability mass under  $\pi$ .



**Figure 5.** From left to right,  $\rho = 0.01, \rho = 0.1$ , and  $\rho = 1$ . (1st row) Behaviors of  $g$  (blue) and  $g_\rho$  (orange) where the contours of the shaded area correspond to  $g + L_\rho$  and  $g + U_\rho$ ; (2nd row) the corresponding normalized smoothed prior densities proportional to  $\exp(-g)$  and  $\exp(-g_\rho)$ ; (3rd row) posterior densities  $\pi_\rho$  w.r.t.  $\rho$ .

### 5.3. Illustration on an Image Inpainting Problem

We illustrate here the correctness of the proposed approach on a multidimensional and non-Gaussian example which classically appears in image processing. To this purpose, we consider the observation of a damaged and noisy image  $\mathbf{y} \in \mathbb{R}^n$  (represented as a vector by lexicographic ordering) related to the unknown original image  $\boldsymbol{\theta} \in \mathbb{R}^d$  by the linear model

$$\mathbf{y} = \mathbf{H}\boldsymbol{\theta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}_n, \sigma^2 \mathbf{I}_n), \quad (48)$$

where  $n < d$ ,  $\mathbf{H} \in \mathbb{R}^{n \times d}$  stands for a decimation binary matrix. The dimension  $d$  being typically large (e.g.,  $10^3 \leq d \leq 10^9$ ), these problems require scalable inference algorithms. Since the matrix  $\mathbf{H}$  is not invertible, the linear inverse problem (48) is ill-posed. To cope with this issue, we assign the total variation

prior distribution to the unknown parameter  $\boldsymbol{\theta}$ , leading to the posterior distribution

$$\pi(\boldsymbol{\theta}|\mathbf{y}) \propto \exp \left( -\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{H}\boldsymbol{\theta}\|_2^2 - \tau \sum_{1 \leq i \leq d} \|(\mathbf{D}\boldsymbol{\theta})_i\|_2 \right), \quad (49)$$

where  $\tau > 0$  is a regularization parameter,  $\mathbf{D}\boldsymbol{\theta} = (\mathbf{D}_1\boldsymbol{\theta}, \mathbf{D}_2\boldsymbol{\theta}) \in \mathbb{R}^{2 \times d}$  is the two-dimensional discrete gradient associated to the image  $\boldsymbol{\theta}$  and the notation  $\mathbf{M}_i$  stands for the  $i$ th column of the matrix  $\mathbf{M}$ , see Chambolle et al. (2010) for more details about the total variation regularization. The presence of the operator  $\mathbf{D}$  and the nondifferentiability of the total variation norm rule out the use of common data augmentation schemes and simulation-based algorithms (e.g., Hamiltonian




Figure 6. From left to right: original image, minimum mean square estimate (MMSE) under  $\pi_\rho$  and absolute bias between the posterior means under  $\pi_\rho$  and  $\pi$ .

and Langevin Monte Carlo methods). Possible surrogates are proximal MCMC methods (Pereyra 2016; Durmus, Moulines, and Pereyra 2018) which replace the nondifferentiable posterior distribution by a smooth approximation based on the proximity operator (Combettes and Pesquet 2011) of the total variation norm. However, the latter does not admit a closed-form expression and iterative routines are commonly used to approximate the latter (Chambolle 2004) leading to higher computational costs.

To mitigate these issues, we propose to rely on a particular instance of AXDA by smoothing the total variation prior with a Gaussian term, leading to the approximate joint posterior density

$$\pi_\rho(\theta, \mathbf{z}|\mathbf{y}) \propto \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{H}\theta\|_2^2 - \tau \sum_{1 \leq i \leq d} \|\mathbf{z}_i\|_2 - \frac{1}{2\rho^2} \|\mathbf{Z} - \mathbf{D}\theta\|_2^2\right), \quad (50)$$

where  $\mathbf{Z} = (\mathbf{z}_1, \mathbf{z}_2) \in \mathbb{R}^{2 \times d}$ . By relying on (50), the inference is now simplified and can be conducted with a Gibbs sampler, see Section 4 in the supplementary materials. Since  $\ker(\mathbf{H}) \cap \ker(\mathbf{D}) = \{\mathbf{0}_d\}$ , the conditional posterior distribution of  $\theta$  is a nondegenerate multivariate Gaussian distribution. Samples from the latter can be obtained efficiently with the two-dimensional discrete Fourier transform by exploiting the periodic boundary conditions for  $\theta$  (Wang et al. 2008; Marnissi et al. 2018). On the other hand, samples from  $\pi_\rho(\mathbf{Z}|\theta)$  can be drawn efficiently using exact data augmentation, see Kyung et al. (2010). All the inference details are given in Section 4 of the supplementary materials.

We illustrate the proposed approximate model  $\pi_\rho$  by considering the Shepp–Logan phantom magnetic resonance image of size  $100 \times 100$  ( $d = 10^4$ ), see Figure 6. We artificially damaged and added noise to this image to build a noisy observation  $\mathbf{y}$  consisting of 90% randomly selected pixels of the initial image. The standard deviation of the Gaussian noise and the regularization parameter have been set to  $\sigma = 7 \times 10^{-2}$  (corresponding to a SNR of 58 dB) and  $\tau = 5$ , respectively. The tolerance parameter has been set to  $\rho = 0.1$ .

To assess the bias of the proposed approach, we implemented the Moreau–Yosida Metropolis-adjusted Langevin algorithm (MYMALA) of Durmus, Moulines, and Pereyra (2018), specifically designed to sample exactly from high-dimensional and non-smooth posterior distributions. For all the MCMC algorithms, the initialization has been set to  $\theta^{[0]} = \mathbf{0}_d$ . We generated  $10^5$  samples and kept the last  $5 \times 10^4$  ones.

Figure 6 shows the minimum mean square estimate (MMSE) under  $\pi_\rho$  along with the original image. One can denote that the MMSE under  $\pi_\rho$  is visually similar to the original image and hence coherent with the reconstruction task. The relative residual error between the former and the MMSE under  $\pi$  is of order 2%. The main differences are located on the boundaries of the image, as depicted in the figure on the left which shows the absolute difference between the pixels of the two posterior means  $\mathbb{E}_\pi(\theta)$  and  $\mathbb{E}_{\pi_\rho}(\theta)$ .

To emphasize the correctness of the proposed approach beyond the comparison between pointwise estimates, we also paid attention to the comparison between posterior credibility sets induced by both  $\pi$  and  $\pi_\rho$ . To this purpose, we considered the highest posterior density region given by

$$\mathcal{C}_\alpha^* = \{\theta \in \mathbb{R}^d \mid f(\theta) \leq \gamma_\alpha\}, \quad (51)$$

where  $\gamma_\alpha \in \mathbb{R}$  is such that  $\int_{\mathcal{C}_\alpha^*} \pi(\theta|\mathbf{y})d\theta = 1 - \alpha$  and  $f$  is the potential function associated to  $\pi(\theta|\mathbf{y})$ .


Figure 7 shows the different values of the scalar summary  $\gamma_\alpha$  estimated using  $\pi$  and the scalar  $\gamma_\alpha^\rho$  estimated using  $\pi_\rho$  for  $\alpha \in [0.01, 0.99]$ . Note that the approximation error associated to  $\gamma_\alpha$  is of order 2.6% whatever the value of  $\alpha$ , which supports the use of  $\pi_\rho$  to conduct Bayesian uncertainty analysis in this problem. After the burn-in period, the efficiency of the Gibbs algorithm used to sample from  $\pi_\rho$  has been measured by comparing the effective sample size (ESS) associated to the slowest component of  $\theta$  to the one obtained with MYMALA. We found that the two ESS were roughly similar but the cost per iteration of the Gibbs sampler (0.079 sec/iteration) is almost two times lower than that of MYMALA (0.144 sec/iteration).<sup>1</sup> In addition, the number of iterations required to reach high-probability regions is much less important for the Gibbs sampler than for MYMALA, showing the interest of AXDA, see Figure 7.

## 6. Conclusion

This article presented a unifying framework for AXDA schemes. AXDA introduces approximate densities to simplify the inference. By building on existing works which considered special instances of AXDA, we illustrate potential benefits that can be inherited by the proposed framework such as distributed computations, robustness or sophisticated inference schemes from the ABC literature. On top of these qualitative properties, we derived a set of theoretical guarantees on the bias involved in

<sup>1</sup>Both algorithms have been implemented in Matlab with the same level of efficiency.





**Figure 7.** (Left) Relative error between the threshold value estimated with  $\pi$  denoted  $\gamma_\alpha$  and the one estimated with  $\pi_\rho$  denoted  $\gamma_\alpha^\rho$  and (right) potential  $f = -\log \pi$  w.r.t. the number of iterations  $t$  for both MYMALA and the Gibbs sampler targeting  $\pi_\rho$ .

the proposed methodology. The latter encompass a large class of AXDA models and a detailed nonasymptotic analysis has been done for Gaussian smoothing. These results have been illustrated on several cases that can arise in statistical learning or signal processing showing the broad scope of application of the proposed approach. In practice, we emphasize that AXDA models can remarkably improve the inference task in big data and high-dimensional settings. In summary, at the price of an approximation which comes with theoretical guarantees, AXDA approaches appear to be a general, systematic and efficient way to conduct simple inference in a wide variety of large-scale problems. They provide accurate estimates with relevant confidence intervals that are crucial in many applications, in particular when no ground truth is available.

## Supplementary Materials

**Appendices:** The supplementary materials include the proofs of [Theorems 1](#) and [2](#); [Propositions 1–5](#); and [Corollary 1](#) and [3](#). It also includes additional details about standard kernels and Bregman divergences, inference details associated to the image inpainting example in [Section 5.3](#) and derivations of classical inference algorithms to target AXDA models. (supplementary\_material.pdf, pdf file)

**Package for AXDA:** The computer code associated to the illustrations and experiments described in this article is also available online. More precisely, the package “AXDA” contains a PYTHON jupyter notebook to reproduce all the tables and figures of the article and a MATLAB-code associated to the image inpainting example along with a README file. (axda.zip, zipped file)

## Funding

Part of this work has been supported by the ANR-3IA Artificial and Natural Intelligence Toulouse Institute (ANITI) under grant agreement ANITI ANR-19-PI3A-0004.

## ORCID

Maxime Vono <http://orcid.org/0000-0003-4859-965X>  
 Nicolas Dobigeon <http://orcid.org/0000-0001-8127-350X>  
 Pierre Chainais <http://orcid.org/0000-0003-4377-7584>

## References

Aguirregabiria, J. M., Hernández, A., and Rivas, M. (2002), “ $\delta$ -Function Converging Sequences,” *American Journal of Physics*, 70, 180–185. [3]

- Albert, J. H., and Chib, S. (1993), “Bayesian Analysis of Binary and Polychotomous Response Data,” *Journal of the American Statistical Association*, 88, 669–679. [1,3]
- Baragatti, M., Grimaud, A., and Pommeret, D. (2013), “Likelihood-Free Parallel Tempering,” *Statistics and Computing*, 23, 535–549. [5]
- Barbos, A.-C., Caron, F., Giovannelli, J.-F., and Doucet, A. (2017), “Clone MCMC: Parallel High-Dimensional Gaussian Gibbs Sampling,” in *Advances in Neural Information Processing Systems*, pp. 5020–5028. [3,10]
- Beaumont, M. A., Zhang, W., and Balding, D. J. (2002), “Approximate Bayesian Computation in Population Genetics,” *Genetics*, 162, 2025–2035. [1,5]
- Beck, A., and Teboulle, M. (2009), “A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems,” *SIAM Journal on Imaging Sciences*, 2, 183–202. [6]
- Ben-Tal, A., Margalit, T., and Nemirovski, A. (2001), “The Ordered Subsets Mirror Descent Optimization Method With Applications to Tomography,” *SIAM Journal on Optimization*, 12, 79–108. [3]
- Besag, J., and Green, P. J. (1993), “Spatial Statistics and Bayesian Computation,” *Journal of the Royal Statistical Society, Series B*, 55, 25–37. [1]
- Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011), “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers,” *Foundations and Trends in Machine Learning*, 3, 1–122. [1,6]
- Bredies, K., Kunisch, K., and Pock, T. (2010), “Total Generalized Variation,” *SIAM Journal on Imaging Sciences*, 3, 492–526. [10]
- Canny, J. (2004), “GaP: A Factor Model for Discrete Data,” in *Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 122–129. [3]
- Chambolle, A. (2004), “An Algorithm for Total Variation Minimization and Applications,” *Journal of Mathematical Imaging and Vision*, 20, 89–97. [12]
- Chambolle, A., Novaga, M., Cremers, D., and Pock, T. (2010), “An Introduction to Total Variation for Image Analysis,” in *Theoretical Foundations and Numerical Methods for Sparse Recovery*, ed. Massimo Fornasier, Berlin: De Gruyter, pp. 263–340. [3,10,11]
- Choi, H. M., and Hobert, J. P. (2013), “The Pólya-Gamma Gibbs Sampler for Bayesian Logistic Regression Is Uniformly Ergodic,” *Electronic Journal of Statistics*, 7, 2054–2064. [2]
- Combettes, P. L., and Pesquet, J.-C. (2011), “Proximal Splitting Methods in Signal Processing,” in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, eds. H. H. Bauschke, R. Burachik, P. Combettes, V. Elser, D. Luke, and H. Wolkowicz, New York: Springer, pp. 185–212. [10,12]
- Damien, P., Wakefield, J., and Walker, S. (1999), “Gibbs Sampling for Bayesian Non-Conjugate and Hierarchical Models by Using Auxiliary Variables,” *Journal of the Royal Statistical Society, Series B*, 61, 331–344. [1]

- Dang, Q. A., and Ehrhardt, M. (2012), “On Dirac Delta Sequences and Their Generating Functions,” *Applied Mathematics Letters*, 25, 2385–2390. [3]
- Del Moral, P., Doucet, A., and Jasra, A. (2012), “An Adaptive Sequential Monte Carlo Method for Approximate Bayesian Computation,” *Statistics and Computing*, 22, 1009–1020. [5]
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977), “Maximum Likelihood From Incomplete Data via the EM Algorithm,” *Journal of the Royal Statistical Society, Series B*, 39, 1–38. [1]
- Dobson, A. J., and Barnett, A. G. (2008), *An Introduction to Generalized Linear Models*, Texts in Statistical Science (3rd ed.), Boca Raton, FL: Chapman & Hall/CRC Press. [4]
- Doucet, A., Godsill, S. J., and Robert, C. P. (2002), “Marginal Maximum a Posteriori Estimation Using Markov Chain Monte Carlo,” *Statistics and Computing*, 12, 77–84. [1]
- Duane, S., Kennedy, A., Pendleton, B. J., and Roweth, D. (1987), “Hybrid Monte Carlo,” *Physics Letters B*, 195, 216–222. [1]
- Dümbgen, L., and Rufibach, K. (2009), “Maximum Likelihood Estimation of a Log-Concave Density and Its Distribution Function: Basic Properties and Uniform Consistency,” *Bernoulli*, 15, 40–68. [6]
- Durmus, A., and Moulines, E. (2017), “Nonasymptotic Convergence Analysis for the Unadjusted Langevin Algorithm,” *The Annals of Applied Probability*, 27, 1551–1587. [6]
- Durmus, A., Moulines, E., and Pereyra, M. (2018), “Efficient Bayesian Computation by Proximal Markov Chain Monte Carlo: When Langevin Meets Moreau,” *SIAM Journal on Imaging Sciences*, 11, 473–506. [6,12]
- Edwards, R. G., and Sokal, A. D. (1988), “Generalization of the Fortuin-Kasteleyn-Swendsen-Wang Representation and Monte Carlo Algorithm,” *Physical Review D*, 38, 2009–2012. [1]
- Fearnhead, P., and Prangle, D. (2012), “Constructing Summary Statistics for Approximate Bayesian Computation: Semi-Automatic Approximate Bayesian Computation,” *Journal of the Royal Statistical Society, Series B*, 74, 419–474. [5]
- Fellows, M., Mahajan, A., Rudner, T. G. J., and Whiteson, S. (2019), “VIREL: A Variational Inference Framework for Reinforcement Learning,” in *Advances in Neural Information Processing Systems*, pp. 7120–7134. [3]
- Geman, D., and Reynolds, G. (1992), “Constrained Restoration and the Recovery of Discontinuities,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14, 367–383. [1]
- Geman, D., and Yang, C. (1995), “Nonlinear Image Recovery With Half-Quadratic Regularization,” *IEEE Transactions on Image Processing*, 4, 932–946. [1,2]
- Giovannelli, J. F. (2008), “Unsupervised Bayesian Convex Deconvolution Based on a Field With an Explicit Partition Function,” *IEEE Transactions on Image Processing*, 17, 16–26. [6]
- Hartley, H. O. (1958), “Maximum Likelihood Estimation From Incomplete Data,” *Biometrics*, 14, 174–194. [1]
- Higdon, D. (2007), *A Primer on Space-Time Modeling From a Bayesian Perspective*, Boca Raton, FL: Chapman & Hall/CRC, pp. 217–279. [9]
- Higdon, D. M. (1998), “Auxiliary Variable Methods for Markov Chain Monte Carlo With Applications,” *Journal of the American Statistical Association*, 93, 585–595. [1]
- Holmes, C. C., and Mallick, B. K. (2003), “Generalized Nonlinear Modeling With Multivariate Free-Knot Regression Splines,” *Journal of the American Statistical Association*, 98, 352–368. [3,6]
- Hurn, M. (1997), “Difficulties in the Use of Auxiliary Variables in Markov Chain Monte Carlo Methods,” *Statistics and Computing*, 7, 35–44. [1]
- Krichene, W., Bayen, A., and Bartlett, P. L. (2015), “Accelerated Mirror Descent in Continuous and Discrete Time,” in *Advances in Neural Information Processing Systems*, pp. 2845–2853. [3]
- Kyung, M., Gill, J., Ghosh, M., and Casella, G. (2010), “Penalized Regression, Standard Errors, and Bayesian Lassos,” *Bayesian Analysis*, 5, 369–411. [12]
- Liechty, M. W., Liechty, J. C., and Müller, P. (2009), “The Shadow Prior,” *Journal of Computational and Graphical Statistics*, 18, 368–383. [3,6]
- Marnissi, Y., Chouzenoux, E., Benazza-Benyahia, A., and Pesquet, J.-C. (2018), “An Auxiliary Variable Method for Markov Chain Monte Carlo Algorithms in High Dimension,” *Entropy*, 20, 110. [1,10,12]
- Meng, X.-L., and van Dyk, D. (1997), “The EM Algorithm—An Old Folk-Song Sung to a Fast New Tune,” *Journal of the Royal Statistical Society, Series B*, 59, 511–567. [1]
- Nesterov, Y., and Spokoiny, V. (2017), “Random Gradient-Free Minimization of Convex Functions,” *Foundations of Computational Mathematics*, 17, 527–566. [7]
- Papandreou, G., and Yuille, A. L. (2010), “Gaussian Sampling by Local Perturbations,” in *Advances in Neural Information Processing Systems*, pp. 1858–1866. [10]
- Park, T., and Casella, G. (2008), “The Bayesian Lasso,” *Journal of the American Statistical Association*, 103, 681–686. [10]
- Pereyra, M. (2016), “Proximal Markov Chain Monte Carlo Algorithms,” *Statistics and Computing*, 26, 745–760. [3,6,12]
- Polson, N. G. (1996), “Convergence of Markov Chain Monte Carlo Algorithms,” in *Bayesian Statistics* (Vol. 5), eds. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith. Oxford: Oxford University Press, pp. 297–321. [1]
- Polson, N. G., Scott, J. G., and Windle, J. (2013), “Bayesian Inference for Logistic Models Using Polya-Gamma Latent Variables,” *Journal of the American Statistical Association*, 108, 1339–1349. [1,2,3]
- Rendell, L. J., Johansen, A. M., Lee, A., and Whiteley, N. (2020), “Global Consensus Monte Carlo,” *Journal of Computational and Graphical Statistics*. [1,4,5]
- Robert, C. P. (2001), *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation* (2nd ed.), New York: Springer. [8]
- Robert, C. P., and Casella, G. (2004), *Monte Carlo Statistical Methods* (2nd ed.), Berlin: Springer. [4]
- Salim, A., Koralev, D., and Richtarik, P. (2019), “Stochastic Proximal Langevin Algorithm: Potential Splitting and Nonasymptotic Rates,” in *Advances in Neural Information Processing Systems*, pp. 6649–6661. [6]
- Scheffé, H. (1947), “A Useful Convergence Theorem for Probability Distributions,” *The Annals of Mathematical Statistics*, 18, 434–438. [2]
- Scott, S. L., Blocker, A. W., Bonassi, F. V., Chipman, H. A., George, E. I., and McCulloch, R. E. (2016), “Bayes and Big Data: The Consensus Monte Carlo Algorithm,” *International Journal of Management Science and Engineering Management*, 11, 78–88. [4]
- She, Y., and Owen, A. B. (2011), “Outlier Detection Using Nonconvex Penalized Regression,” *Journal of the American Statistical Association*, 106, 626–639. [4]
- Sisson, S. A., Fan, Y., and Beaumont, M. A. (eds.) (2018a), *Handbook of Approximate Bayesian Computation* (1st ed.), Boca Raton, FL: Chapman and Hall/CRC Press. [5]
- (2018b), “Overview of Approximate Bayesian Computation,” in *Handbook of Approximate Bayesian Computation* (1st ed.), eds. S. Sisson, Y. Fan, and M. Beaumont, Boca Raton, FL: Chapman and Hall/CRC Press, pp. 3–54. [1,6]
- Swendsen, R. H., and Wang, J.-S. (1987), “Nonuniversal Critical Dynamics in Monte Carlo Simulations,” *Physical Review Letters*, 58, 86–88. [1]
- Tanner, M. A., and Wong, W. H. (1987), “The Calculation of Posterior Distributions by Data Augmentation,” *Journal of the American Statistical Association*, 82, 528–540. [1]
- (2010), “From EM to Data Augmentation: The Emergence of MCMC Bayesian Computation in the 1980s,” *Statistical Science*, 25, 506–516. [1]
- van Dyk, D. A., and Meng, X.-L. (2001), “The Art of Data Augmentation,” *Journal of Computational and Graphical Statistics*, 10, 1–50. [1,2]
- Vono, M., Dobigeon, N., and Chainais, P. (2019), “Split-and-Augmented Gibbs Sampler—Application to Large-Scale Inference Problems,” *IEEE Transactions on Signal Processing*, 67, 1648–1661. [1,3,6]
- Wand, M. P., and Jones, M. C. (1995), *Kernel Smoothing*, London: Chapman & Hall/CRC. [3]
- Wang, C., and Blei, D. M. (2018), “A General Method for Robust Bayesian Modeling,” *Bayesian Analysis*, 13, 1163–1191. [4]
- Wang, X., and Dunson, D. B. (2013), “Parallelizing MCMC via Weierstrass Sampler,” Technical Report, arXiv no. 1312.4605. [4]
- Wang, Y., Yang, J., Yin, W., and Zhang, Y. (2008), “A New Alternating Minimization Algorithm for Total Variation Image Reconstruction,” *SIAM Journal on Imaging Sciences*, 1, 248–272. [12]
- Wilkinson, R. (2013), “Approximate Bayesian Computation (ABC) Gives Exact Results Under the Assumption of Model Error,” *Statistical Applications in Genetics and Molecular Biology*, 12, 1–13. [5]