

Quantifying the mini-batching error in Bayesian inference for Adaptive Langevin dynamics

Sekkat Inass

CERMICS,

Ecole des Ponts,

Marne-la-Vallée, France

INASS.SEKKAT@ENPC.FR

Gabriel Stoltz

CERMICS, Ecole des Ponts,

Marne-la-Vallée, France

MATERIALS team-project,

Inria Paris, France

GABRIEL.STOLTZ@ENPC.FR

Editor: David Wipf

Abstract

Bayesian inference allows to obtain useful information on the parameters of models, either in computational statistics or more recently in the context of Bayesian Neural Networks. The computational cost of usual Monte Carlo methods for sampling posterior laws in Bayesian inference scales linearly with the number of data points. One option to reduce it to a fraction of this cost is to resort to mini-batching in conjunction with unadjusted discretizations of Langevin dynamics, in which case only a random fraction of the data is used to estimate the gradient. However, this leads to an additional noise in the dynamics and hence a bias on the invariant measure which is sampled by the Markov chain. We advocate using the so-called Adaptive Langevin dynamics, which is a modification of standard inertial Langevin dynamics with a dynamical friction which automatically corrects for the increased noise arising from mini-batching. We investigate the practical relevance of the assumptions underpinning Adaptive Langevin (constant covariance for the estimation of the gradient, Gaussian minibatching noise), which are not satisfied in typical models of Bayesian inference, and quantify the bias induced by minibatching in this case. We also suggest a possible extension of AdL to further reduce the bias on the posterior distribution, by considering a dynamical friction depending on the current value of the parameter to sample.

Keywords: Mini-batching, Langevin dynamics, Markov Chain Monte Carlo

1. Introduction

Bayesian modeling allows to determine the distribution of parameters in statistical models, and hence to estimate functions of these parameters and their uncertainties. The Bayesian approach to neural networks in particular has recently gained some attention from the Machine Learning community; see for instance Wang and Yeung (2020); Izmailov et al. (2021) and references therein. Since the distributions of parameters are given by (possibly very) high dimensional probability measures, Markov Chain Monte Carlo (MCMC) techniques (see Robert and Casella (2004); Brooks et al. (2011)) are the default methods to sample

these target measures. In this work, we restrict ourselves to probability distributions having a density with respect to the Lebesgue measure. Quantities of interest are then expectations with respect to the target distribution, which are approximated by Monte Carlo estimates based on ergodicity results for the Markov chains under consideration. Let us however mention, that, beyond sampling from distributions of parameters, MCMC methods can also be used for optimization when run at small target temperature, which may be beneficial to explore complex non-convex energy landscapes (see for instance Ma et al. (2019); Leimkuhler et al. (2019); Dieuleveut et al. (2020)), or to ensure some form of regularization in an attempt to avoid overfitting (as already noted in Welling and Teh (2011)).

Two major classes of MCMC techniques can be distinguished. The first one gathers methods based on the Metropolis–Hastings algorithm (see Metropolis et al. (1953); Hastings (1970)). The second class of MCMC techniques relies on discretizations of stochastic differential equations (SDEs) which are ergodic with respect to the target measure. The SDEs used to sample from a probability distribution were originally introduced in molecular dynamics, where atomic configurations in a system are typically distributed according to the Boltzmann–Gibbs distribution $\pi(d\theta) = Z^{-1} \exp(-V(\theta)) d\theta$, with Z the normalization constant and V the potential energy function of the system; see Frenkel and Smit (2002); Tuckerman (2010); Leimkuhler and Matthews (2015); Allen and Tildesley (2017). Two prominent examples of such dynamics are Langevin dynamics (which we consider in this work as being the underdamped, or kinetic version) and its overdamped limit. In practice, methods from the two classes can be blended, as for the Metropolis-adjusted Langevin algorithm (see Roberts and Tweedie (1996); Rossky et al. (1978); Roberts and Rosenthal (1998)), which is a Metropolis–Hastings algorithm whose proposal function is provided by a Euler–Maruyama discretization of overdamped Langevin dynamics. We focus in this work on the second class of techniques, namely the discretization of SDEs. Langevin-like dynamics are gaining increasing popularity in Machine Learning and related application fields (see for instance Chen et al. (2014); Mou et al. (2021); Gürbüzbalaban et al. (2021) to quote just a few works).

Both Metropolis–Hastings methods and discretization of SDEs can be computationally expensive in the context of Bayesian inference since the log-likelihood and/or its gradient have to be evaluated at each step, either to perform a Metropolis test or to compute the forces in the dynamics. The cost of computing the log-likelihood and its gradient scales as $\mathcal{O}(N_{\text{data}})$, where N_{data} is the size of the data set, which may be large. Some variations of Metropolis-like algorithms, based on estimates of the log-likelihood obtained from a random subsample of the data, have been introduced to reduce the computational cost of one iteration; see for example Korattikara et al. (2014); Bardenet et al. (2014); Quiroz et al. (2019). These methods however require some prior knowledge on the target measure, and/or introduce some bias on the measure actually sampled by the algorithm. A similar approach was proposed for discretizations of SDEs by Welling and Teh (2011), who suggested to use a mini-batch of the data to construct an estimator of the gradient of the log-likelihood, leading to the so-called Stochastic Gradient Langevin Dynamics (SGLD).

SGLD however also induces biases on the sampled invariant measure, which have two origins: the finiteness of the time step and the fact that the mini-batch size n is smaller than N_{data} . It is possible to remove the bias by using decreasing time steps, as initially suggested in Welling and Teh (2011) and mathematically analyzed in Teh et al. (2016), but

this is not practical for the convergence of longtime averages, in particular when there is some metastability in the system (*i.e.* in situations when the posterior probability measure is multimodal, and the numerical methods remain temporarily stuck in one of the modes before finding their way to another one. Multimodality is generic in Bayesian inference as soon as some parameters are exchangeable. It can be argued that numerical methods which are able to switch between equivalent modes have good exploration properties, which allows them to find possibly unanticipated local modes; see Chopin et al. (2012)). An analysis of the asymptotic bias of SGLD for fixed step sizes is provided in Vollmer et al. (2016) (together with some analysis of the non-asymptotic bias, following the approach developed in Mattingly et al. (2010)). The results show that SGLD has a weak error of order one, similarly to the Euler-Maruyama discretization of overdamped Langevin dynamics, but with an extra term in the leading error arising from the mini-batching procedure. In fact, the magnitude of this term makes it the dominant one, unless n is very close to N_{data} .

Various extensions and refinements of SGLD were proposed, in an attempt to improve the performance of the method and/or to reduce its bias. A first trend is to apply the mini-batching philosophy to dynamics which are more efficient in terms of sampling than overdamped Langevin dynamics, in particular Langevin dynamics (see Matthews and Weare (2018)), Hybrid Monte Carlo algorithms (see Chen et al. (2014)) and piecewise deterministic Markov processes (see Pakman et al. (2017); Bierkens et al. (2019)). A second trend is to rely on control variate techniques to reduce the covariance of the stochastic estimator of the gradient of the log-likelihood, computed in practice by Gaussian approximations of the modes of the target probability measure; see for instance Nagapetyan et al. (2017); Brosse et al. (2018); Baker et al. (2019). Of course, both trends can (should) be combined.

Our emphasis in this work is on the adaptive Langevin dynamics (AdL), introduced in Jones and Leimkuhler (2011); Ding et al. (2014), which provides a way to reduce, and hopefully even almost remove the bias arising from mini-batching in SGLD. This dynamics is a modification of Langevin dynamics where the friction is considered as a dynamical variable that adjusts itself so that the distribution of the velocities is correct. In addition to correcting for the mini-batching bias, it can also be used to train neural networks; see Leimkuhler et al. (2019). AdL was mathematically studied in Leimkuhler et al. (2020), and further tested from a numerical viewpoint in Shang et al. (2015). The method completely removes the mini-batching bias when the covariance matrix $\Sigma_{\mathbf{x}}(\theta)$ of the estimator of the gradient is constant in the range of parameters explored by the method. This assumption is satisfied for Gaussian posterior distributions, as obtained when considering a Gaussian prior and Gaussian likelihoods, or when the number of data points is large enough and the posterior distribution concentrates around a Gaussian distribution according to the Bernstein–von Mises theorem (see for instance Section 10.2 in van der Vaart (1998)). There are however situations where this assumption does not hold (as we highlight on a numerical example in Section 3.3.2), in which case AdL fails to fully correct for the bias arising from mini-batching.

In this paper, we consider the case of extreme mini-batching procedures, corresponding to n as small as 1, as in initial works on stochastic approximation in Robbins and Monro (1951). In this situation, no central limit theorem can be invoked to precisely characterize the statistical properties of the stochastic gradient estimator. This is in contrast with various works which assume the mini-batching noise to be Gaussian. This is however not needed

to make precise the weak error of numerical methods, and hence the bias on the invariant probability measure, as we show in our estimates on the bias for SGLD and Langevin dynamics with stochastic gradient estimators. One of our contributions is to precisely quantify how the bias on the invariant measure depends on the type (with or without replacement) and amount of mini-batching, through some key parameter $\varepsilon(n)$ defined in Section 2.2. We also carefully study the covariance matrix of the stochastic gradient estimator in illustrative numerical examples, highlighting that this matrix may exhibit substantial variations in the range of parameters explored by the dynamics under consideration, so that the key assumption underlying AdL is not satisfied in general. Nonetheless, our numerical analysis explains why AdL still succeeds to substantially reduce the bias compared to SGLD and Langevin dynamics. We finally introduce an extended version of AdL allowing to even further reduce the bias incurred by non constant covariance matrices.

One of our main contributions is to prove that, for all the Langevin-type SDEs considered in this paper (standard Langevin, adaptive Langevin, extended adaptive Langevin), the bias introduced by the minibatching procedure is controlled by

$$\varepsilon(n)\Delta t \min_{S \in \mathcal{S}} \|\Sigma_{\mathbf{x}} - S\|_{L^2(\pi)} = \varepsilon(n)\Delta t \|\Sigma_{\mathbf{x}} - S^*\|_{L^2(\pi)}, \quad (1)$$

where Δt is the time step used to discretize the SDE at hand, and π is the target posterior distribution. The vector space of matrix valued functions \mathcal{S} depends on the chosen dynamics. For discretizations of standard Langevin dynamics, $S^* = 0$ (see Sections 2.3.2 and 2.4.2, in particular (37)), whereas S^* corresponds to the average of $\Sigma_{\mathbf{x}}$ with respect to π for AdL (see Section 3.3.1, in particular (67) and Remark 15) and to the $L^2(\pi)$ -projection of $\Sigma_{\mathbf{x}}$ onto the vector space of symmetric matrices generated by some basis of functions for the extended version of AdL we introduce (see Section 4.2, in particular (76)). We numerically verify that the reduction in the bias is proportional to the quality of the approximation of the covariance matrix on a basis of functions (*e.g.* constant matrices, piecewise constant scalar functions, ...).

The paper is organized as follows. We start in Section 2 by reviewing various results related to the numerical analysis of SDEs and the quantification of the bias on the invariant measure sampled by SGLD and Langevin dynamics with stochastic gradient estimators. We next turn to AdL in Section 3: we illustrate that some residual bias remains present due to the fact that the covariance of the gradient estimator is not constant in general, and we quantify it. Section 4 is dedicated to the introduction of an extended version of AdL that allows to accommodate non constant covariance matrices for the gradient estimator and further reduce the bias. Some conclusions and perspectives are gathered in Section 5.

2. Stochastic gradient Markov Chain Monte Carlo

We consider a Bayesian inference problem where we denote by $\mathbf{x} = (x_i)_{i=1, \dots, N_{\text{data}}} \in \mathcal{X}^{N_{\text{data}}}$ a set of N_{data} data points, with $\mathcal{X} \subset \mathbb{R}^{d_{\text{data}}}$. The data points are assumed to be independent and identically distributed (i.i.d.) with respect to an elementary likelihood probability measure $P_{\text{elem}}(\cdot|\theta)$, parameterized by $\theta \in \Theta = \mathbb{R}^d$. The likelihood of \mathbf{x} is then given by

$$P_{\text{likelihood}}(\mathbf{x}|\theta) = \prod_{i=1}^{N_{\text{data}}} P_{\text{elem}}(x_i|\theta).$$

In the Bayesian framework, a prior distribution P_{prior} is considered on the parameters. The aim is to sample from the posterior distribution

$$\pi(\theta|\mathbf{x}) \propto P_{\text{prior}}(\theta)P_{\text{likelihood}}(\mathbf{x}|\theta), \quad (2)$$

in order, for instance, to compute expectations with respect to this distribution. Sampling is usually done with MCMC methods. This however requires, at each iteration, to compute either the log-likelihood $\pi(\cdot|\mathbf{x})$ or its gradient

$$\nabla_{\theta}(\log \pi(\theta|\mathbf{x})) = \nabla_{\theta}(\log P_{\text{prior}}(\theta)) + \sum_{i=1}^{N_{\text{data}}} \nabla_{\theta}(\log P_{\text{elem}}(x_i|\theta)). \quad (3)$$

The cost of computing these quantities scales as $\mathcal{O}(N_{\text{data}})$, and is usually the computational bottleneck in implementations of MCMC algorithms. As discussed in the introduction, we focus here on stochastic gradient dynamics, which are discretizations of stochastic dynamics admitting $\pi(\cdot|\mathbf{x})$ as invariant probability measure, and where the cost of evaluating (3) is reduced by approximating the gradient through some estimator based on a mini-batch of the complete data set.

We first review in Section 2.1 elements on the error analysis of discretization of SDEs, with some emphasis on the bias induced on the invariant measure, as these results will be repeatedly used throughout this paper, and are key to understanding the performance and limitations of all the methods we discuss. We then recall in Section 2.2 the mini-batching method introduced in Welling and Teh (2011) and focus on the case where the size of the mini-batch is small, possibly limited to a single point (as done in stochastic approximation algorithms in Robbins and Monro (1951)). We then recall two classical SDEs upon which our method is based: overdamped Langevin dynamics in Section 2.3 (known as SGLD when using mini-batching), and then Langevin dynamics in Section 2.4. In each of these sections, we analyze how mini-batching affects the posterior distribution and quantify the bias incurred on it by studying the associated effective dynamics. These results are illustrated by numerical examples in Section 2.5.

2.1 Some elements on error analysis for discretizations of SDEs

Consider a SDE $(\theta_t)_{t \geq 0} \subset \Theta$ with generator \mathcal{L} admitting $\pi(\cdot|\mathbf{x})$ as invariant probability measure: For any smooth function ϕ with compact support,

$$\int_{\Theta} \mathcal{L}\phi \, d\pi(\cdot|\mathbf{x}) = 0.$$

Denote by $(\theta^m)_{m \geq 0}$ a time discretization of the SDE with a fixed time step Δt (so that θ^m is an approximation of $\theta_{m\Delta t}$). We assume that the Markov chain corresponding to the time discretization of the SDE admits a unique invariant probability measure, denoted by $\pi_{\Delta t}$. This is for instance the case for Langevin-type dynamics when the drift of the dynamics is globally Lipschitz or when Lyapunov conditions are satisfied; see Mattingly et al. (2002). For a given observable ϕ , the target expectation

$$\mathbb{E}_{\pi}(\phi) = \int_{\Theta} \phi(\theta)\pi(\theta|\mathbf{x}) \, d\theta$$

is approximated by $\mathbb{E}_{\pi_{\Delta t}}(\phi)$, which is itself typically estimated by the trajectory average

$$\widehat{\phi}_{\Delta t, N_{\text{iter}}} = \frac{1}{N_{\text{iter}}} \sum_{m=1}^{N_{\text{iter}}} \phi(\theta^m).$$

The total error on averages with respect to $\pi(\cdot|\mathbf{x})$ can then be written as:

$$\widehat{\phi}_{\Delta t, N_{\text{iter}}} - \mathbb{E}_{\pi}(\phi) = (\mathbb{E}_{\pi_{\Delta t}}(\phi) - \mathbb{E}_{\pi}(\phi)) + \left(\widehat{\phi}_{\Delta t, N_{\text{iter}}} - \mathbb{E}_{\pi_{\Delta t}}(\phi) \right).$$

The first term on the right hand side corresponds to the bias on the invariant probability measure resulting from taking finite step sizes. The second term in the error has two origins: (i) a bias coming from the initial distribution of θ^0 when this random variable is not distributed according to $\pi_{\Delta t}$; (ii) a statistical error, which is dictated by the central limit theorem for N_{iter} large.

We focus in this work on the bias on the invariant probability measure, which can be bounded using the weak order of the scheme, provided some ergodicity conditions are satisfied. Recall that a numerical scheme is of weak order s if for any smooth and compactly supported function ϕ and final time $T > 0$, there exists $C \in \mathbb{R}_+$ such that

$$\forall m \in \{1, \dots, \lceil T/\Delta t \rceil\}, \quad |\mathbb{E}[\phi(\theta_{m\Delta t})] - \mathbb{E}[\phi(\theta^m)]| \leq C\Delta t^s. \quad (4)$$

When this condition holds, and under appropriate ergodicity conditions (see Talay and Tubaro (1990) for a pioneering work, as well as Talay (2002); Mattingly et al. (2002); Abdulle et al. (2014); Leimkuhler et al. (2016); Bou-Rabee and Owhadi (2010) for subsequent works on Langevin-like dynamics), the following bound is obtained on the bias on the invariant probability measure of the numerical scheme: For any smooth and compactly supported function ϕ , there exist $\Delta t_{\star} > 0$ and L such that

$$\forall \Delta t \in (0, \Delta t_{\star}], \quad |\mathbb{E}_{\pi_{\Delta t}}(\phi) - \mathbb{E}_{\pi}(\phi)| \leq L\Delta t^s. \quad (5)$$

In order to write a sufficient local consistency condition to obtain an estimate such as (4), we introduce the evolution operator $P_{\Delta t}$ associated with the numerical scheme at hand, defined as follows: For any smooth and compactly supported function ϕ ,

$$(P_{\Delta t}\phi)(\theta) = \mathbb{E}[\phi(\theta^{m+1}) | \theta^m = \theta].$$

Under appropriate technical conditions, including moment conditions on the iterates of the numerical scheme (see Theorem 2.1 in Milstein and Tretyakov (2013) for a precise statement), a sufficient condition for (4) to hold is

$$P_{\Delta t} = e^{\Delta t\mathcal{L}} + \mathcal{O}(\Delta t^{s+1}), \quad (6)$$

where $(e^{t\mathcal{L}}\phi)(\theta) = \mathbb{E}[\phi(\theta_t) | \theta_0 = \theta]$ is the evolution operator associated with the underlying SDE. Here and in the sequel, the above equality has to be understood as follows: For any smooth and compactly supported function ϕ , there exist $\Delta t_{\star} > 0$ and $K \in \mathbb{R}_+$ such that, for any $\Delta t \in (0, \Delta t_{\star}]$, there is a function $R_{\phi, \Delta t}$ for which

$$P_{\Delta t}\phi = e^{\Delta t\mathcal{L}}\phi + \Delta t^{s+1}R_{\phi, \Delta t}, \quad \sup_{\Delta t \in (0, \Delta t_{\star}]} \sup_{\theta \in \Theta} |R_{\phi, \Delta t}(\theta)| \leq K; \quad (7)$$

see for instance Section 3.3 in Lelièvre and Stoltz (2016) for a more precise discussion of this point.

Let us conclude this section by introducing the concept of effective dynamics (also called modified SDEs in Shardlow (2006); Zygalkakis (2011)), a tool inspired by backward numerical analysis. The idea is to construct a continuous dynamics, parameterized by the time step, which better coincides with the numerical scheme than the reference dynamics. More precisely, given a numerical scheme of weak order s , characterized by the local consistency condition (6), we look for a modified SDE with generator $\mathcal{L}_{\text{mod},\Delta t}$ such that

$$P_{\Delta t} - e^{\Delta t \mathcal{L}_{\text{mod},\Delta t}} = \mathcal{O}(\Delta t^{s+2}). \quad (8)$$

An alternative reformulation is the following: denoting by $(\tilde{\theta})_{t \geq 0}$ the solution to the SDE with generator $\mathcal{L}_{\text{mod},\Delta t}$, then

$$\mathbb{E}[\phi(\theta^1) | \theta_0 = \theta] = \mathbb{E}[\phi(\tilde{\theta}_{\Delta t}) | \theta_0 = \theta] + \mathcal{O}(\Delta t^{s+2}),$$

while

$$\mathbb{E}[\phi(\theta^1) | \theta_0 = \theta] = \mathbb{E}[\phi(\theta_{\Delta t}) | \theta_0 = \theta] + \mathcal{O}(\Delta t^{s+1}).$$

The prime interest of effective dynamics in our work is to provide some interpretation of the behavior of numerical schemes. Following the standard philosophy of backward analysis, it is indeed possible to obtain information on the behavior of numerical schemes by studying the properties of the effective dynamics, in particular their invariant probability measures. This is also useful for instance to understand the impact of mini-batching on SGLD and Langevin dynamics (see respectively Vollmer et al. (2016) and Section 2.4), and is also the foundation of Adaptive Langevin dynamics. In particular, when (8) holds with $\mathcal{L}_{\text{mod},\Delta t} = \mathcal{L} + \Delta t^s \mathcal{A}_{\text{mod}}$, then (5) can be rewritten more precisely in the form of a Talay–Tubaro expansion (see Talay and Tubaro (1990)):

$$\forall \Delta t \in (0, \Delta t_*], \quad |\mathbb{E}_{\pi_{\Delta t}}(\phi) - \mathbb{E}_{\pi}(\phi) - \Delta t^s \mathbb{E}_{\pi}(f\phi)| \leq \tilde{L} \Delta t^{s+1}, \quad (9)$$

with $f = (-\mathcal{L}^*)^{-1} \mathcal{A}_{\text{mod}}^* \mathbf{1}$, adjoints being taken on $L^2(\pi)$ (provided the inverse of \mathcal{L} can be defined on an appropriate subspace of $L^2(\pi)$, see again Section 3.3 in Lelièvre and Stoltz (2016)). The formula (9) allows to make precise the dominant term in the error $\mathbb{E}_{\pi_{\Delta t}}(\phi) - \mathbb{E}_{\pi}(\phi)$.

2.2 Mini-Batching procedure

The idea of using mini-batching in the context of SDEs has first been introduced through SGLD in Welling and Teh (2011) to reduce the calculation time of one step of the discretization of the overdamped Langevin dynamics (see (17) below) from $\mathcal{O}(N_{\text{data}})$ to a fraction of this cost. It is inspired from the classical stochastic gradient descent method in Robbins and Monro (1951) in the sense that a consistent approximation of $\nabla_{\theta}(\log \pi(\cdot | \mathbf{x}))$ is used in the discretization of overdamped Langevin dynamics; see Section 2.3 below.

Mini-batching relies on computing at each iteration an approximation of the gradient of the log-likelihood using a random subset of size n of the data points, which reduces the

cost from $\mathcal{O}(N_{\text{data}})$ to $\mathcal{O}(n)$. More precisely, the gradient of the posterior distribution is approximated by

$$\widehat{F}_n(\theta) = \nabla_{\theta}(\log P_{\text{prior}}(\theta)) + \frac{N_{\text{data}}}{n} \sum_{i \in I_n} \nabla_{\theta}(\log P_{\text{elem}}(x_i|\theta)), \quad (10)$$

where I_n is a random subset of size n generated by sampling uniformly indices from $\{1, \dots, N_{\text{data}}\}$, with or without replacement. Sampling with or without replacement produces similar results when $n \ll N_{\text{data}}$ with N_{data} large enough. Sampling with replacement however leads to estimators (10) with larger variances for a given value of n (as quantified by (13) below), and should therefore be avoided when larger batches are considered. In any case, it is easily shown that $\widehat{F}_n(\theta)$ in (10) is a consistent approximation of $\nabla_{\theta}(\log \pi(\cdot|\mathbf{x}))$ given by (3).

Properties of $\widehat{F}_n(\theta)$. In view of (10), the covariance matrix of the stochastic gradient, which is nonnegative and symmetric, is given by:

$$\text{cov}(\widehat{F}_n(\theta)) = \varepsilon(n)\Sigma_{\mathbf{x}}(\theta), \quad (11)$$

with $\Sigma_{\mathbf{x}}(\theta)$ the empirical covariance of the gradient estimator for $n = 1$ (*i.e.* with expectations computed with respect to the random variable \mathcal{I} uniformly distributed in $\{1, \dots, N_{\text{data}}\}$):

$$\begin{aligned} \Sigma_{\mathbf{x}}(\theta) &= \text{cov}_{\mathcal{I}}[\nabla_{\theta}(\log P_{\text{elem}}(x_{\mathcal{I}}|\theta))] \\ &= \frac{1}{N_{\text{data}} - 1} \sum_{i=1}^{N_{\text{data}}} [\nabla_{\theta}(\log P_{\text{elem}}(x_i|\theta)) - \mathcal{F}_{\mathbf{x}}(\theta)] [\nabla_{\theta}(\log P_{\text{elem}}(x_i|\theta)) - \mathcal{F}_{\mathbf{x}}(\theta)]^T, \end{aligned} \quad (12)$$

where uv^T is the matrix $(u_i v_j)_{1 \leq i, j \leq d} \in \mathbb{R}^{d \times d}$ for $u, v \in \mathbb{R}^d$, the average force reads

$$\mathcal{F}_{\mathbf{x}}(\theta) = \frac{1}{N_{\text{data}}} \sum_{i=1}^{N_{\text{data}}} \nabla_{\theta}(\log P_{\text{elem}}(x_i|\theta)),$$

and

$$\varepsilon(n) = \begin{cases} \frac{N_{\text{data}}(N_{\text{data}} - 1)}{n}, & \text{for sampling with replacement,} \\ \frac{N_{\text{data}}(N_{\text{data}} - n)}{n}, & \text{for sampling without replacement.} \end{cases} \quad (13)$$

We refer for instance to Chaudhari and Soatto (2018) for more details. We will see in the next sections that, for all the methods we consider, the bias due to mini-batching is controlled by $\varepsilon(n)$ whether sampling is performed with or without replacement. It is really this parameter (in fact, $\varepsilon(n)\Delta t$) which needs to be small in order for the asymptotic analysis on the bias to be correct. Note that $\varepsilon(n) \geq N_{\text{data}} - 1$ for sampling with replacement, which means that a bias is necessarily observed in this case. Values of $\varepsilon(n)$ of order N_{data} are obtained only when the mini-batch size n is a fraction of the total number of data points N_{data} . To obtain values of order 1, sampling has to be done without replacement




Figure 1: Histogram of $Z_{\mathbf{x}, N_{\text{data}}, n}$ when sampling is performed without replacement for the Gaussian model of Section 2.5.1, with $N_{\text{data}} = 100$, $\sigma_\theta = 1$, $\sigma_x = 1$, $\mu = 0$, $\theta = 0.5$. The reference standard Gaussian distribution is superimposed as a continuous line.

with n close to N_{data} . In contrast, $\varepsilon(n)$ is of order N_{data}^2 when the mini-batch size is small, whether sampling is performed with or without replacement.

The above definitions allow to rewrite (10) as

$$\widehat{F}_n(\theta) = \nabla_\theta(\log \pi(\theta|\mathbf{x})) + \sqrt{\varepsilon(n)} \Sigma_{\mathbf{x}}^{\frac{1}{2}}(\theta) Z_{\mathbf{x}, N_{\text{data}}, n}, \quad (14)$$

where $Z_{\mathbf{x}, N_{\text{data}}, n}$ is by construction a centered random variable with identity covariance. If n and N_{data} are large enough, and $n \ll N_{\text{data}}$, the central limit theorem holds so that $Z_{\mathbf{x}, N_{\text{data}}, n}$ asymptotically follows a centered reduced normal distribution (as assumed for instance in Chen et al. (2014); Ahn et al. (2012); Zhu et al. (2019)). Most of the works in the literature consider the regime where the central limit theorem holds. This is however unnecessary when it comes to quantifying the weak error of the numerical scheme and hence the error on the invariant measure. In our work, we consider explicitly the case when n is of order 1 and N_{data} is (moderately) large, so that $Z_{\mathbf{x}, N_{\text{data}}, n}$ is not necessarily close to a Gaussian distribution. This is illustrated in Figures 1 and 2, where we plot the distribution of $Z_{\mathbf{x}, N_{\text{data}}, n}$ for various values of n for the models introduced in Section 2.5, where the elementary likelihood is either a Gaussian or a mixture of Gaussians. It is clear that the distribution of $Z_{\mathbf{x}, N_{\text{data}}, n}$ is far from being Gaussian for small values of n , but is nonetheless centered and with covariance I_d . Distributions close to Gaussian are obtained for $n = 20-30$ for the situation considered in Figure 2.

2.3 Stochastic Gradient Langevin Dynamics

We introduce in this section the SGLD algorithm (see Section 2.3.1) and analyze the bias due to replacing the gradient (3) by the stochastic estimator (10) using the effective dynamics associated with the numerical method at hand (see Section 2.3.2).

2.3.1 DESCRIPTION OF THE METHOD

A popular SDE to sample from $\pi(\cdot|\mathbf{x})$ is the overdamped Langevin dynamics:

$$d\theta_t = \nabla_\theta(\log \pi(\theta_t|\mathbf{x})) dt + \sqrt{2} dW_t, \quad (15)$$




Figure 2: Histogram of $Z_{\mathbf{x}, N_{\text{data}}, n}$ for $\theta = (1, 0.2)$, when sampling is performed without replacement for the model of Section 2.5.2 with $N_{\text{data}} = 100$ data points sampled from a mixture of Gaussians with parameters $\sigma_1 = \sigma_2 = 0.4$, $w = 0.5$, $\theta_1 = 0.2$ and $\theta_2 = 1$. The reference standard Gaussian distribution is superimposed as a continuous line on the marginal distributions in θ_1, θ_2 .

where W_t is a standard d -dimensional Wiener process. Its generator is given by

$$\mathcal{L}_{\text{ovd}} = \nabla_{\theta}(\log \pi(\cdot|\mathbf{x}))^T \nabla_{\theta} + \Delta_{\theta}. \quad (16)$$

It is well known that the process (15) is irreducible and admits $\pi(\cdot|\mathbf{x})$ as a unique invariant probability measure, so that expectations can be approximated by trajectory averages (see for instance Kliemann (1987)). In general, one cannot directly simulate the overdamped Langevin dynamics (15). To numerically approximate the solution, the widely used Euler-Maruyama scheme can be considered:

$$\theta^{m+1} = \theta^m + \Delta t \nabla_{\theta}(\log \pi(\theta^m|\mathbf{x})) + \sqrt{2\Delta t} G^m, \quad (17)$$

where $\Delta t > 0$ is the step size, and $(G^m)_{m \geq 0}$ is a vector of i.i.d. standard d -dimensional Gaussian random variables. It can be shown that the bias $\mathbb{E}_{\pi_{\Delta t}}(\phi) - \mathbb{E}_{\pi}(\phi)$ for overdamped Langevin dynamics is of order $\mathcal{O}(\Delta t)$; see for instance Talay (1990) and Theorem 7.3 in Mattingly et al. (2002) for pioneering works on ergodic properties of discretization of SDEs and error estimation for globally Lipschitz vector fields.

Using the stochastic estimator (10), SGLD corresponds to the following numerical scheme:

$$\theta^{m+1} = \theta^m + \Delta t \widehat{F}_n(\theta^m) + \sqrt{2\Delta t} G^m. \quad (18)$$

We assume in the sequel that this Markov chain admits a unique invariant probability measure $\pi_{\Delta t, n}$ (we omit the dependence on \mathbf{x} to simplify the notation).

Remark 1 *It is suggested in Welling and Teh (2011) to use decreasing time steps as this allows to ultimately eliminate the bias without resorting to a Metropolis Hastings scheme. However, considering decreasing time steps means that we need more iterations for a fixed final time T , which increases the computational cost of the algorithm and results in a possibly large variance in the estimation of averages at fixed computational cost. In practice, it is more customary to use a small finite time step. As discussed in Vollmer et al. (2016), the bias arising from mini-batching dominates the one resulting from the time step discretization. Our focus in this work is therefore on the bias arising from mini-batching and not on the bias coming from the use of finite time steps.*

2.3.2 EFFECTIVE SGLD

We recall in this section how to prove that the bias between the invariant measure $\pi_{\Delta t, n}$ of SGLD (18) and $\pi(\cdot|\mathbf{x})$ is of order $\mathcal{O}((1 + \varepsilon(n))\Delta t)$ – by which we mean that error estimates such as (5) hold with Δt^s replaced by $(1 + \varepsilon(n))\Delta t$ on the right hand side, namely: for any smooth and compactly supported function ϕ , there exist $\Delta t_{\star} > 0$ and L such that

$$\left| \int_{\Theta} \phi d\pi_{\Delta t, n} - \int_{\Theta} \phi d\pi \right| \leq L(1 + \varepsilon(n))\Delta t.$$

Such error estimates were already obtained in Vollmer et al. (2016). This analysis is important since we will repeatedly rely on it to quantify the bias on the invariant probability measures of numerical schemes, in particular for Adaptive Langevin dynamics and their extensions, which is why we insist on presenting the overall strategy.

From a technical viewpoint, the analysis relies on effective dynamics, discussed at the end of Section 2.1. To derive an effective dynamics for the SGLD, we follow the general framework of backward error analysis for SDEs considered in Shardlow (2006); Zygalkis (2011); Debussche and Faou (2012); Abdulle et al. (2012), which is based on building a perturbation to the generator associated with the original dynamics (16), constructed so that SGLD evolved over one step is closer to the SDE associated with the modified generator than to the original dynamics. Recalling the definition (7) for the notation we use for remainder terms, the results in Vollmer et al. (2016) show that (see also Appendix A.1 for a short proof allowing to make precise the dependence of the remainder term on n)

$$\widehat{P}_{\Delta t, n} = e^{\Delta t(\mathcal{L}_{\text{ovd}} + \Delta t \mathcal{A}_{\text{ovd}, n})} + \mathcal{O}\left((1 + \varepsilon(n)^{3/2})\Delta t^3\right), \quad \mathcal{A}_{\text{ovd}, n} = \varepsilon(n)\mathcal{A}_{\text{mb}} + \mathcal{A}_{\text{disc}}, \quad (19)$$

where $\widehat{P}_{\Delta t, n}$ is the evolution operator associated with the SGLD scheme (18), and the operators \mathcal{A}_{mb} and $\mathcal{A}_{\text{disc}}$ respectively encode the perturbations arising from mini-batching and time discretization:

$$\begin{aligned} \mathcal{A}_{\text{mb}}\phi &= \frac{1}{2}\Sigma_{\mathbf{x}} : \nabla_{\theta}^2\phi, \\ \mathcal{A}_{\text{disc}}\phi &= -\nabla_{\theta}^2(\log \pi(\cdot|\mathbf{x})) : \nabla_{\theta}^2\phi - \frac{1}{2}\nabla_{\theta}\Delta(\log \pi(\cdot|\mathbf{x}))^T \nabla_{\theta}\phi \\ &\quad - \frac{1}{2}\nabla_{\theta}(\log \pi(\cdot|\mathbf{x}))^T \nabla_{\theta}^2(\log \pi(\cdot|\mathbf{x})) \nabla_{\theta}\phi. \end{aligned}$$

In the latter expressions, we denote by $:$ the Frobenius inner product of two square matrices, *i.e.*

$$\forall M^1, M^2 \in \mathbb{R}^{d \times d}, \quad M^1 : M^2 = \sum_{i, j=1}^d M_{i, j}^1 M_{i, j}^2.$$

In fact, as discussed in Remark 18 of Appendix A.1, the error estimate in (19) can be improved to $\mathcal{O}((1 + \varepsilon(n))\Delta t^3)$ when $\mathbb{E}[Z_{\mathbf{x}, N_{\text{data}}, n}^3] = 0$. Using results from Remark 5.5 in Lelièvre and Stoltz (2016), or from Vollmer et al. (2016), and following the general strategy of Talay–Tubaro estimates such as (9), we deduce that the bias between $\pi_{\Delta t, n}$ and π is of order $\mathcal{O}((1 + \varepsilon(n))\Delta t)$ provided that Δt and $\varepsilon(n)\Delta t$ are small enough.

Proposition 2 *Assume that the SGLD scheme (18) admits a unique invariant probability measure $\pi_{\Delta t, n}$, and that $\log \pi(\cdot|\mathbf{x})$ satisfies the assumption of Kopec (2015a), namely $\pi(\cdot|\mathbf{x})$ has moments of all orders and $\log \pi(\cdot|\mathbf{x})$ is a smooth semi-convex function which, together with its derivatives, grows at most polynomially, and satisfies the following bounds for some constants $a, b > 0$:*

$$-\theta^T \nabla_{\theta} [\log \pi(\theta|\mathbf{x})] \geq -a|\theta|^2 - b.$$

Introduce the smooth functions

$$f_{\text{disc}} = (-\mathcal{L}_{\text{ovd}})^{-1} \mathcal{A}_{\text{disc}}^* \mathbf{1}, \quad f_{\text{mb}} = (-\mathcal{L}_{\text{ovd}})^{-1} \mathcal{A}_{\text{mb}}^* \mathbf{1}. \quad (20)$$

Then, for any smooth function φ with compact support, there is $\Delta t_{\star} > 0$ and $K \in \mathbb{R}_{+}$ for which

$$\int_{\Theta} \varphi d\pi_{\Delta t, n} - \int_{\Theta} \varphi d\pi = \Delta t \int_{\Theta} \varphi (f_{\text{disc}} + \varepsilon(n)f_{\text{mb}}) d\pi + \varepsilon(n)^{3/2} \Delta t^2 R_{\varphi, n, \Delta t},$$

with

$$\sup_{1 \leq n \leq N_{\text{data}}} \sup_{\Delta t \in (0, \Delta t_{\star}]} |R_{\varphi, n, \Delta t}| \leq K.$$

The bias is determined by the norm of the correction functions $f_{\text{disc}}, f_{\text{mb}}$ in appropriate functional spaces. These functions are well defined and belong to $L^2(\pi)$ in the framework of Kopec (2015a), so that the bias term can be bounded as

$$\left| \Delta t \int_{\Theta} \varphi (f_{\text{disc}} + \varepsilon(n) f_{\text{mb}}) d\pi \right| \leq (1 + \varepsilon(n)) \Delta t (\|f_{\text{disc}}\|_{L^2(\pi)} + \|f_{\text{mb}}\|_{L^2(\pi)}) \|\varphi\|_{L^2(\pi)}. \quad (21)$$

The bias is expected to be larger when $\Sigma_{\mathbf{x}}$ is larger, because $\mathcal{A}_{\text{mb}}^* \mathbf{1}$ in (20) is larger – although this statement is not as clear cut as for Langevin dynamics since the correction function f_{mb} involves derivatives of $\Sigma_{\mathbf{x}}$, in contrast to estimates obtained for underdamped and adaptive Langevin dynamics. The dimension dependence of the correction terms is not so easy to determine, as the inverse of $-\mathcal{L}_{\text{ovd}}$ considered as an operator on $L^2(\pi)$ scales as the inverse of the Poincaré constant, and the scaling of the latter constant is not known unless the logarithm of the target distribution is (a perturbation of) a strongly log-concave function; see for instance Bakry et al. (2014).

Note that the bias can be decreased by either decreasing Δt or increasing n , hence decreasing $\varepsilon(n)$ in view of (13). However, as already discussed in Vollmer et al. (2016), the mini-batching error usually dominates the discretization error by orders of magnitude since $\varepsilon(n)$ is proportional to N_{data}^2 unless n is a fraction of N_{data} , in which case it scales as N_{data} .

Remark 3 *To reduce the mini-batching bias, it is suggested in Vollmer et al. (2016) to renormalize the magnitude of the injected noise by a quantity involving $\Sigma_{\mathbf{x}}(\theta)$ (the so-called modified SGLD scheme). Since $\Sigma_{\mathbf{x}}(\theta)$ is usually unknown, this requires estimating this matrix, which is however computationally expensive since the estimation has to be repeated for each new value of θ , and may cancel the gain provided by mini-batching in the first place. Our focus in this work is to reduce the mini-batching error without the need to estimate $\Sigma_{\mathbf{x}}(\theta)$.*

2.4 Langevin dynamics with mini-batching

It has been observed in practice that a better sampling of probability measures can be provided by Langevin dynamics, both in the literature on computational statistical physics (see for instance Cancès et al. (2007)) and more recently in the machine learning literature, see for instance Dalalyan and Riou-Durand (2020). We first present in Section 2.4.1 the numerical scheme obtained by discretizing Langevin dynamics with a splitting scheme and replacing the gradient of the log-likelihood by its estimator (10), and then analyze the bias induced on the invariant probability measure of the numerical scheme in Section 2.4.2.

2.4.1 STANDARD LANGEVIN DYNAMICS

Langevin dynamics introduces some inertia in the evolution of θ , through an extended configuration space with a momentum vector p conjugated to θ . It can be seen as a perturbation

of the Hamiltonian dynamics where some fluctuation/dissipation mechanism is added to the evolution of the momenta, and reads

$$\begin{cases} d\theta_t = p_t dt, \\ dp_t = \nabla_\theta(\log \pi(\theta_t|\mathbf{x})) dt - \Gamma p_t dt + \sqrt{2}\Gamma^{1/2} dW_t, \end{cases} \quad (22)$$

where $\Gamma \in \mathbb{R}^{d \times d}$ is a positive definite symmetric matrix. It would be possible, as in molecular dynamics, to attach a mass to each degree of freedom, but we set here for simplicity this mass to 1, the generalization to non trivial mass matrices being straightforward. The generator of Langevin dynamics (22) is given by

$$\mathcal{L}_{\text{lan}} = \nabla_\theta(\log \pi(\cdot|\mathbf{x}))^T \nabla_p + p^T \nabla_\theta - p^T \Gamma \nabla_p + \Gamma : \nabla_p^2. \quad (23)$$

The diffusion constant $\sqrt{2}\Gamma^{1/2}$ in front of the Wiener process ensures that the following probability distribution is invariant:

$$\mu(d\theta dp|\mathbf{x}) = \pi(\theta|\mathbf{x})\tau(dp) d\theta,$$

where

$$\tau(dp) = (2\pi)^{-d/2} e^{-p^2/2} dp, \quad (24)$$

see for instance Pavliotis (2014); Leimkuhler and Matthews (2015); Lelièvre and Stoltz (2016). In particular, the marginal distribution of $\mu(\cdot|\mathbf{x})$ in the θ variable is indeed the target distribution (2). It can even be shown that time averages along solutions of (22) almost surely converge to averages with respect to μ since the generator of the dynamics is hypoelliptic; see Kliemann (1987).

To numerically approximate the solution of Langevin dynamics, we use a numerical integrator based on a second order Strang splitting, the so-called OABAO scheme, encoded by the following evolution operator (although there are various other choices of orderings, see for instance Leimkuhler and Matthews (2013); Leimkuhler et al. (2016); Leimkuhler and Matthews (2015)):

$$P_{\Delta t} = e^{\Delta t \mathcal{L}_3/2} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_1} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_3/2}, \quad (25)$$

where

$$\mathcal{L}_1 = \nabla_\theta \log(\pi(\cdot|\mathbf{x}))^T \nabla_p, \quad \mathcal{L}_2 = p^T \nabla_\theta, \quad \mathcal{L}_3 = -p^T \Gamma \nabla_p + \Gamma : \nabla^2. \quad (26)$$

The elementary generators \mathcal{L}_1 and \mathcal{L}_2 are respectively associated with the elementary differential equations $dp_t = \nabla_\theta \log(\pi(\theta_t|\mathbf{x})) dt$ and $d\theta_t = p_t dt$, which can be analytically integrated. The elementary generator \mathcal{L}_3 is associated with the Ornstein–Uhlenbeck process $dp_t = -\Gamma p_t dt + \sqrt{2}\Gamma^{1/2} dW_t$, which can also be analytically integrated as:

$$p_t = e^{-\Gamma t} p_0 + \sqrt{2} \int_0^t e^{-(t-s)\Gamma} \Gamma^{1/2} dW_s \sim \mathcal{N}(\alpha_t p_0, \text{Id} - \alpha_t), \quad \alpha_t = e^{-\Gamma t}. \quad (27)$$

Finally, the numerical scheme encoded by (25) reads:

$$\begin{cases} p^{m+\frac{1}{3}} = \alpha_{\Delta t/2} p^m + (\text{Id} - \alpha_{\Delta t})^{1/2} G^m, \\ \theta^{m+\frac{1}{2}} = \theta^m + \frac{\Delta t}{2} p^{m+\frac{1}{3}}, \\ p^{m+\frac{2}{3}} = p^{m+\frac{1}{3}} + \Delta t \nabla_{\theta} \left[\log \pi \left(\theta^{m+\frac{1}{2}} \mid \mathbf{x} \right) \right], \\ \theta^{m+1} = \theta^{m+\frac{1}{2}} + \frac{\Delta t}{2} p^{m+\frac{2}{3}}, \\ p^{m+1} = \alpha_{\Delta t/2} p^{m+\frac{2}{3}} + (\text{Id} - \alpha_{\Delta t})^{1/2} G^{m+\frac{1}{2}}, \end{cases} \quad (28)$$

where $(G^m)_{m \geq 0}$ and $(G^{m+\frac{1}{2}})_{m \geq 0}$ are two independent families of i.i.d. standard d -dimensional Gaussian random variables. Moreover, it is proved in Leimkuhler et al. (2016); Abdulle et al. (2014); Durmus et al. (2021) that the Markov chain generated by (28) admits a unique invariant probability measure $\mu_{\Delta t}$ and that there exists $C \in \mathbb{R}_+$ such that, for any smooth function ϕ with compact support,

$$\left| \int_{\Theta} \phi(\theta, p) \mu_{\Delta t}(d\theta dp) - \int_{\Theta} \phi(\theta, p) \mu(d\theta dp \mid \mathbf{x}) \right| \leq C \Delta t^2.$$

The key element to prove this statement is the fact that the numerical scheme (28) is an approximation of weak order 2 of Langevin dynamics (see Leimkuhler et al. (2016); Abdulle et al. (2014))

$$P_{\Delta t} = e^{\Delta t \mathcal{L}_{\text{lan}}} + \mathcal{O}(\Delta t^3). \quad (29)$$

2.4.2 ERROR ESTIMATES FOR LANGEVIN DYNAMICS WITH MINI-BATCHING

In order to analyze the error on the posterior measure sampled by a discretization of Langevin dynamics used in conjunction with mini-batching, we derive here the effective dynamics associated with the numerical method (28) when the gradient of the log-likelihood is replaced by its stochastic estimator (10) (similar results are obtained for other Strang splittings). This corresponds to the following numerical scheme:

$$\begin{cases} p^{m+\frac{1}{3}} = \alpha_{\Delta t/2} p^m + (\text{Id} - \alpha_{\Delta t})^{1/2} G^m, \\ \theta^{m+\frac{1}{2}} = \theta^m + \frac{\Delta t}{2} p^{m+\frac{1}{3}}, \\ p^{m+\frac{2}{3}} = p^{m+\frac{1}{3}} + \Delta t \widehat{F}_n \left(\theta^{m+\frac{1}{2}} \right), \\ \theta^{m+1} = \theta^{m+\frac{1}{2}} + \frac{\Delta t}{2} p^{m+\frac{2}{3}}, \\ p^{m+1} = \alpha_{\Delta t/2} p^{m+\frac{2}{3}} + (\text{Id} - \alpha_{\Delta t})^{1/2} G^{m+\frac{1}{2}}. \end{cases} \quad (30)$$

When using mini-batching, we are in fact replacing $e^{\Delta t \mathcal{L}_1}$ in (25) by the elementary evolution operator $Q_{\Delta t}^{\mathcal{L}_1}$ acting as

$$\left(Q_{\Delta t}^{\mathcal{L}_1} \varphi \right) (\theta, p) = \mathbb{E} \left[\varphi \left(\theta, p + \Delta t \widehat{F}_n(\theta) \right) \right].$$

A simple computation provided in Appendix A.2 shows that

$$Q_{\Delta t}^{\mathcal{L}_1} = e^{\Delta t \mathcal{L}_1} + \mathcal{O}(\varepsilon(n)\Delta t^2). \quad (31)$$

Denoting by $\widehat{P}_{\Delta t, n}$ the evolution operator of the numerical scheme (30), and by $\mu_{\Delta t, n}$ its invariant probability measure (assuming that it exists and it is unique), it can then be proved that (see Appendix A.2)

$$\widehat{P}_{\Delta t, n} = e^{\Delta t \mathcal{L}_{\text{lan}}} + \mathcal{O}((\varepsilon(n) + \Delta t)\Delta t^2). \quad (32)$$

This allows to formulate the following result on the bias of $\mu_{\Delta t, n}$.

Proposition 4 *Assume that the Markov chain (30) admits a unique invariant probability measure $\mu_{\Delta t, n}$, and that $\log \pi(\cdot|\mathbf{x})$ satisfies the assumption of Kopec (2015b), namely it is a smooth semi-convex function which, together with its derivatives, grows at most polynomially, and satisfies the following bounds for some constants $a, b, c > 0$ with $b \in (0, 1)$:*

$$-\log \pi(\theta|\mathbf{x}) \geq a|\theta|^2, \quad -\frac{1}{2}\theta^T \nabla_{\theta} [\log \pi(\theta|\mathbf{x})] \geq -b \log \pi(\theta|\mathbf{x}) + \gamma^2 \frac{b(2-b)}{8(1-b)} |\theta|^2 - c.$$

Then, there exist smooth functions $f_{\text{disc}}, f_{\text{mb}}$ such that, for any smooth function φ with compact support, there is $\Delta t_{\star} > 0$ and $K \in \mathbb{R}_+$ for which

$$\int_{\Theta \times \mathbb{R}^d} \varphi d\mu_{\Delta t, n} - \int_{\Theta \times \mathbb{R}^d} \varphi d\mu = \Delta t \int_{\Theta \times \mathbb{R}^d} \varphi (\Delta t f_{\text{disc}} + \varepsilon(n) f_{\text{mb}}) d\mu + (\Delta t + \varepsilon(n)^{3/2}) \Delta t^2 R_{\varphi, n, \Delta t},$$

with

$$\sup_{1 \leq n \leq N_{\text{data}}} \sup_{\Delta t \in (0, \Delta t_{\star})} |R_{\varphi, n, \Delta t}| \leq K.$$

Note that the factor Δt in the term $\Delta t + \varepsilon(n)$ of the remainder is useful only when $n = N_{\text{data}}$, in which case $\varepsilon(n) = 0$. Proposition 4 makes precise the statement that the bias between μ and $\mu_{\Delta t, n}$ is of order $\mathcal{O}((\varepsilon(n) + \Delta t)\Delta t)$. Indeed, $f_{\text{disc}}, f_{\text{mb}}$ belong to $L^2(\pi)$ in the framework of Kopec (2015a) (see also the discussion in Lelièvre and Stoltz (2016)), so that the bias term can be bounded as

$$\left| \Delta t \int_{\Theta \times \mathbb{R}^d} \varphi (\Delta t f_{\text{disc}} + \varepsilon(n) f_{\text{mb}}) d\mu \right| \leq \Delta t (\Delta t \|f_{\text{disc}}\|_{L^2(\mu)} + \varepsilon(n) \|f_{\text{mb}}\|_{L^2(\mu)}) \|\varphi\|_{L^2(\mu)}. \quad (33)$$

It is clear that the error $\varepsilon(n)\Delta t$ coming from mini-batching dominates the error Δt^2 due to time discretization since $\varepsilon(n)$ is much larger than Δt unless n is very close to N_{data} . Comparing this equality to (29) highlights the fact that mini-batching degrades the consistency estimate (29) by one order in Δt with respect to Langevin dynamics (22). The bias is of the same order of magnitude as for SGLD, see (21).

Sufficient conditions for $\log \pi(\theta|\mathbf{x})$ to satisfy the assumptions in Kopec (2015b) are for instance that $-\log \pi(\theta|\mathbf{x})$ is a bounded perturbation of a strongly convex function. The dimension dependence of the correction functions $f_{\text{disc}}, f_{\text{mb}}$ is not clear since, as for overdamped Langevin dynamics, bounds on these functions depend on the norm of the inverse of the adjoint of the generator \mathcal{L}_{lan} (see (36) for f_{mb}). The scaling with respect

to the dimension for such operators has only been recently studied, see Cao et al. (2019); Bernard et al. (2022); Brigati and Stoltz (2023). As for overdamped Langevin dynamics, it depends on the Poincaré constant.

We do not provide at this stage functional estimates on the function f_{mb} which appears in the dominant term of the bias in (33). This is discussed after (35) below.

Remark 5 *To remove the bias on the invariant probability measure at dominant order in $\varepsilon(n)\Delta t$, it is suggested in Matthews and Weare (2018) to modify the integration of the Ornstein–Uhlenbeck part on the momenta. This requires however an estimation of the covariance matrix, which can again be computationally prohibitive (as discussed in Remark 3).*

2.4.3 EFFECTIVE DYNAMICS FOR LANGEVIN DYNAMICS WITH MINI-BATCHING

We now construct an effective dynamics which coincides at order 3 in Δt over one time step with the numerical scheme (30) even when using mini-batching, in order to obtain an equality similar to (29). This effective dynamics is the key building block to understand Adaptive Langevin dynamics in Section 3.1. A straightforward computation, following the lines of Matthews and Weare (2018), shows that (see Appendix A.2)

$$\widehat{P}_{\Delta t, n} = e^{\Delta t(\mathcal{L}_{\text{lan}} + \Delta t \varepsilon(n) \mathcal{A}_{\text{lan}})} + \mathcal{O}\left((1 + \varepsilon(n)^{3/2})\Delta t^3\right), \quad \mathcal{A}_{\text{lan}} = \frac{1}{2} \Sigma_{\mathbf{x}} : \nabla_p^2. \quad (34)$$

The effective dynamics associated with $\mathcal{L}_{\text{lan}} + \Delta t \varepsilon(n) \mathcal{A}_{\text{lan}}$ is

$$\begin{cases} d\tilde{\theta}_t = \tilde{p}_t dt, \\ d\tilde{p}_t = \nabla_{\theta} \left[\log \pi \left(\tilde{\theta}_t | \mathbf{x} \right) \right] dt - \Gamma \tilde{p}_t dt + \left(2\Gamma + \varepsilon(n) \Delta t \Sigma_{\mathbf{x}}(\tilde{\theta}_t) \right)^{1/2} dW_t, \end{cases} \quad (35)$$

where $2\Gamma + \varepsilon(n) \Delta t \Sigma_{\mathbf{x}}(\theta)$ is a positive definite matrix since $\Sigma_{\mathbf{x}}$ is positive (even though the latter matrix is unknown). This allows to identify the correction function f_{mb} in (33) as

$$f_{\text{mb}} = (-\mathcal{L}_{\text{lan}}^*)^{-1} \mathcal{A}_{\text{lan}}^* \mathbf{1} = \frac{1}{2} (-\mathcal{L}_{\text{lan}}^*)^{-1} \Sigma_{\mathbf{x}} : (\nabla_p^*)^2 \mathbf{1}, \quad (36)$$

where the adjoints of the operator are taken on $L^2(\mu)$. Thanks to estimates on \mathcal{L}_{lan} obtained from hypocoercive estimates (see Hérau (2006); Dolbeault et al. (2009, 2015)), we deduce that there exists a constant $C \in \mathbb{R}_+$ such that

$$\|f_{\text{mb}}\|_{L^2(\mu)} \leq C \|\Sigma_{\mathbf{x}} : (\nabla_p^*)^2 \mathbf{1}\|_{L^2(\mu)} = C \|(\nabla_p^*)^2 \mathbf{1}\|_{L^2(\tau)} \|\Sigma_{\mathbf{x}}\|_{L^2(\pi)}. \quad (37)$$

In fact, $(\nabla_p^*)^2 \mathbf{1}$ is the matrix whose entry (i, j) for $1 \leq i, j \leq d$ is $p_i p_j - \delta_{i,j}$. Here and in the sequel, we assume implicitly that $\Sigma_{\mathbf{x}}$ has all its entries in $L^2(\pi)$. The bias on the invariant probability measure sampled by the numerical scheme is therefore, at dominant order, of order $\varepsilon(n) \Delta t \|\Sigma_{\mathbf{x}}\|_{L^2(\pi)}$, in accordance with the general estimate (1) (choosing $\mathcal{S} = \{0\}$ hence $S^* = 0$).

Remark 6 *Consider the simple case when the covariance of the gradient estimator is constant, namely $\Sigma_{\mathbf{x}} = \sigma^2 \mathbf{I}_d \in \mathbb{R}^{d \times d}$, and the friction is isotropic, namely $\Gamma = \gamma \mathbf{I}_d \in \mathbb{R}^{d \times d}$,*

with $\sigma, \gamma > 0$. In this situation, the modified Langevin dynamics (35) samples the invariant probability measure with a density proportional to $\mu(\theta)^{\beta_{\text{eff}}}$, with

$$\beta_{\text{eff}} = \left(1 + \frac{\varepsilon(n)\sigma^2\Delta t}{2\gamma}\right)^{-1} < 1. \quad (38)$$

Since the exponent β_{eff} is smaller than 1, the modified Langevin dynamics samples a tempered version of μ .

2.5 Numerical illustration

We illustrate in this section the results on the bias of the posterior measure introduced by mini-batching when the elementary likelihoods are given by either a Gaussian or a mixture of Gaussians. The aim is to numerically quantify the bias in the non asymptotic regime $n = \mathcal{O}(1)$ (for which $Z_{\mathbf{x}, N_{\text{data}}, n}$ is not Gaussian), and to have a benchmark on the bias to compare with AdL.

2.5.1 GAUSSIAN POSTERIOR

We first suppose that the elements of the data set are normally distributed, namely $x_i|\theta \sim \mathcal{N}(\theta, \sigma_x^2)$, where θ is the parameter to estimate. The prior distribution on θ is a centered normal distribution with variance σ_θ^2 . In this case, simple computations show that the posterior distribution on θ is also Gaussian with mean μ_{post} and variance σ_{post}^2 , where (see for instance Vollmer et al. (2016))

$$\mu_{\text{post}} = \left(\frac{\sigma_x^2}{\sigma_\theta^2} + N_{\text{data}}\right)^{-1} \sum_{i=1}^{N_{\text{data}}} x_i, \quad \sigma_{\text{post}}^2 = \left(\frac{1}{\sigma_\theta^2} + \frac{N_{\text{data}}}{\sigma_x^2}\right)^{-1}. \quad (39)$$

Moreover,

$$\widehat{F}_n(\theta) = -\frac{\theta}{\sigma_\theta^2} + \frac{N_{\text{data}}}{n} \sum_{i \in I_n} \frac{x_i - \theta}{\sigma_x^2}.$$

The variance of $\widehat{F}_n(\theta)$ as a function of θ can therefore be analytically computed using (12) (see again Vollmer et al. (2016) for instance):

$$\Sigma_{\mathbf{x}}(\theta) = \text{var}_{\mathcal{I}} [\nabla_{\theta}(\log P_{\text{elem}}(x_{\mathcal{I}}|\theta))] = \frac{\text{var}(\mathbf{x})}{\sigma_x^4}, \quad (40)$$

where \mathcal{I} is a random variable uniformly distributed in $\{1, \dots, N_{\text{data}}\}$, and

$$\text{var}(\mathbf{x}) = \frac{1}{N_{\text{data}} - 1} \sum_{i=1}^{N_{\text{data}}} \left(x_i - \frac{1}{N_{\text{data}}} \sum_{j=1}^{N_{\text{data}}} x_j\right)^2$$

is the empirical variance of the data. In this case, the covariance of the force $\Sigma_{\mathbf{x}}$ is constant and does not depend on the parameter θ .

Let us first discuss SGLD. As shown in Section 2.1 in Vollmer et al. (2016), there is no bias on the mean of the posterior distribution $\pi_{\Delta t, n}$ of the Markov chain associated

with SGLD. The latter distribution is however not strictly Gaussian, because the random variable $Z_{\mathbf{x}, N_{\text{data}}, n}$ is not Gaussian. It can however be shown that the non Gaussianity arises from higher order terms, so that the shape of the the distribution is Gaussian at dominant order in $(1 + \varepsilon(n))\Delta t$. In order to characterize the dominant order of the shape of the distribution, we write the effective dynamics with generator $\mathcal{L}_{\text{ovd}} + \Delta t \mathcal{A}_{\text{ovd}, n}$ in (19), which turns out to be an Ornstein–Uhlenbeck process:

$$d\theta_t = \left[- \left(1 + \frac{\Delta t}{2\sigma_{\text{post}}^2} \right) \frac{\theta - \mu_{\text{post}}}{\sigma_{\text{post}}^2} \right] dt + \sqrt{2} \left(1 + \frac{\Delta t \varepsilon(n)}{2\sigma_x^4} \text{var}(\mathbf{x}) + \frac{\Delta t}{\sigma_{\text{post}}^2} \right)^{1/2} dW_t.$$

The probability measure of the latter process is a Gaussian distribution with mean 0 and variance

$$\sigma_{\text{post}}^2 \left(1 + \frac{\Delta t}{2\sigma_{\text{post}}^2} \right)^{-1} \left(1 + \frac{\varepsilon(n)\Delta t}{2\sigma_x^4} \text{var}(\mathbf{x}) + \frac{\Delta t}{\sigma_{\text{post}}^2} \right).$$

This shows that the invariant probability measure of SGLD is, at dominant order in Δt and $\varepsilon(n)\Delta t$, a Gaussian distribution with mean 0 and variance

$$\sigma_{\text{post}}^2 \left[1 + \frac{\Delta t}{2} \left(\frac{\varepsilon(n)}{\sigma_x^4} \text{var}(\mathbf{x}) + \frac{1}{\sigma_{\text{post}}^2} \right) \right] + \mathcal{O}((1 + \varepsilon(n)^2)\Delta t^2). \quad (41)$$

For Langevin dynamics with mini-batching, there is (as for SGLD) no bias on the mean of the posterior distribution $\pi_{\Delta t, n}$, as proved in Appendix B. As for SGLD, the latter distribution is also not strictly Gaussian when the random variable $Z_{\mathbf{x}, N_{\text{data}}, n}$ is not Gaussian. It is however close to a Gaussian distribution since the marginal posterior distribution in the θ variable obtained for the effective Langevin dynamics is, according to Remark 6, a Gaussian distribution with mean μ_{post} and variance $\sigma_{\text{post}}^2/\beta_{\text{eff}}$. This means that the variance for the discretization of Langevin dynamics (30) is, at dominant order in Δt and $\varepsilon(n)\Delta t$, and when $\Gamma = \gamma \mathbf{I}_d$,

$$\sigma_{\text{post}}^2 \left(1 + \frac{\varepsilon(n)\Delta t}{2\gamma\sigma_x^4} \text{var}(\mathbf{x}) \right) + \mathcal{O}((1 + \varepsilon(n)^2)\Delta t^2). \quad (42)$$

Note that this expression coincides with (41) when $\gamma = 1$ and $\varepsilon(n) \gg 1$.

To perform the numerical experiments, we generate a data set of $N_{\text{data}} = 100$ according to a Gaussian distribution with mean $\theta_0 = 0$ and variance $\sigma_x = 1$. We also set $\sigma_\theta = 1$ in the prior distribution. We run the SGLD scheme (18) and Langevin dynamics with the numerical scheme (30) with $\Gamma = 1$ for a final time $T = 10^6$ and various values of Δt (which corresponds to $N_{\text{iter}} = T/\Delta t$ time steps). We also consider various values of n , the subsampling of the data points being done with and without replacement. We report in Figure 3 the relative bias on the variance, given in the limit $N_{\text{iter}} \rightarrow +\infty$ by the ratio of

$$\left| \left[\int_{\mathbb{R}} \theta^2 \pi_{\Delta t, n}(\theta|\mathbf{x}) d\theta - \left(\int_{\mathbb{R}} \theta \pi_{\Delta t, n}(\theta|\mathbf{x}) d\theta \right)^2 \right] - \left[\int_{\mathbb{R}} \theta^2 \pi(\theta|\mathbf{x}) d\theta - \left(\int_{\mathbb{R}} \theta \pi(\theta|\mathbf{x}) d\theta \right)^2 \right] \right|,$$

and σ_{post}^2 , where $\pi_{\Delta t, n}(\cdot|\mathbf{x})$ denotes the invariant measure in the θ variable of the numerical scheme under consideration (for (30), this corresponds to the marginal measure of $\mu_{\Delta t, n}(\theta, p|\mathbf{x})$ in the θ variable).

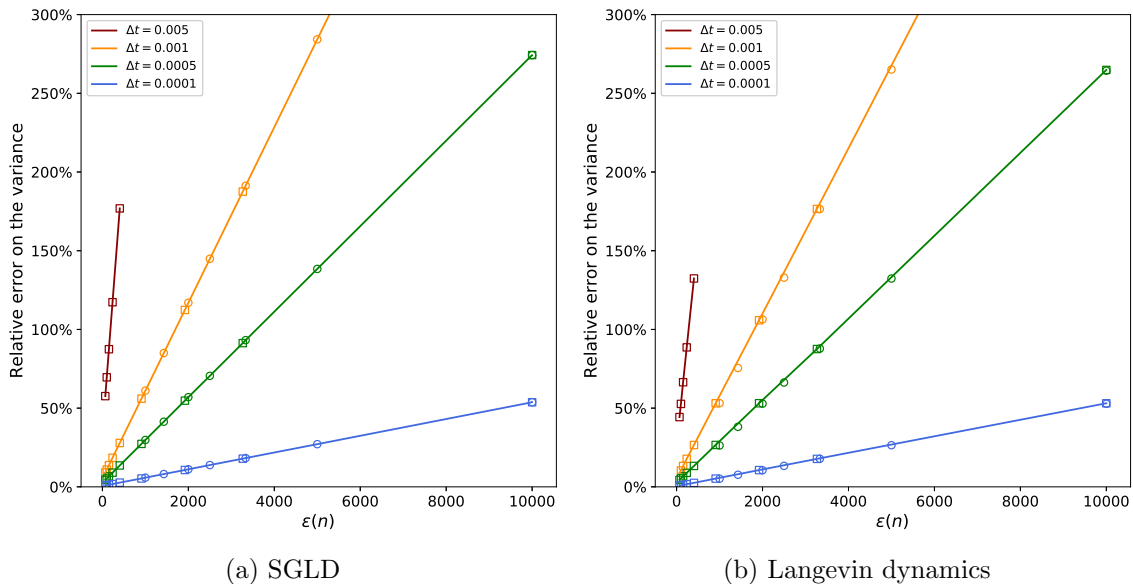


Figure 3: Relative error on the variance of the posterior distribution for various values of Δt and n when the elementary likelihood is a Gaussian distribution, when sampling with (circles) and without replacement (squares).

Several conclusions can be drawn from the results presented in Figure 3. First, note that the bias is determined by the value of $\varepsilon(n)$ irrespectively of the fact that sampling is performed with or without replacement. Second, the bias is indeed affine in $\varepsilon(n)$, with a slope which is proportional to Δt . In fact, all curves would be superimposed if we were plotting the error as a function of $\varepsilon(n)\Delta t$, which demonstrates that the error is indeed determined at dominant order by this parameter. We do not report results for $\varepsilon(n) = 0$, which corresponds to $n = N_{\text{data}}$ and sampling without replacement, since the corresponding error is anyway much smaller than the one arising from mini-batching. Finally, errors are very similar for SGLD and the scheme (30) associated with Langevin dynamics, as anticipated from the comparison of (41) and (42) when choosing $\gamma = 1$.

2.5.2 MIXTURE OF GAUSSIANS

We next consider a more realistic case where the data points are distributed according to a mixture of Gaussians:

$$P_{\text{elem}}(x_i|\theta) = \frac{w}{\sqrt{2\pi}\sigma_1} \exp\left(-\frac{(x_i - \mu_1)^2}{2\sigma_1^2}\right) + \frac{1-w}{\sqrt{2\pi}\sigma_2} \exp\left(-\frac{(x_i - \mu_2)^2}{2\sigma_2^2}\right), \quad (43)$$

where $w \in [0, 1]$, $\mu_1, \mu_2 \in \mathbb{R}$ and $\sigma_1, \sigma_2 \in (0, +\infty)$. We consider the case when the parameters to estimate are the centers of the Gaussians $\theta = (\mu_1, \mu_2)$, whereas σ_1, σ_2 and w are given. The prior distribution on the vector of parameters θ is chosen to be a centered normal distribution with covariance matrix I_2 . To perform numerical simulations, we fix $\mu_1 = 1$, $\mu_2 = 0.5$, $\sigma_1 = \sigma_2 = 0.4$, $w = 0.4$ and $N_{\text{data}} = 200$ to generate the data set according to (43). We run again the numerical schemes (18) and (30), with $\Gamma = I_2$ and an

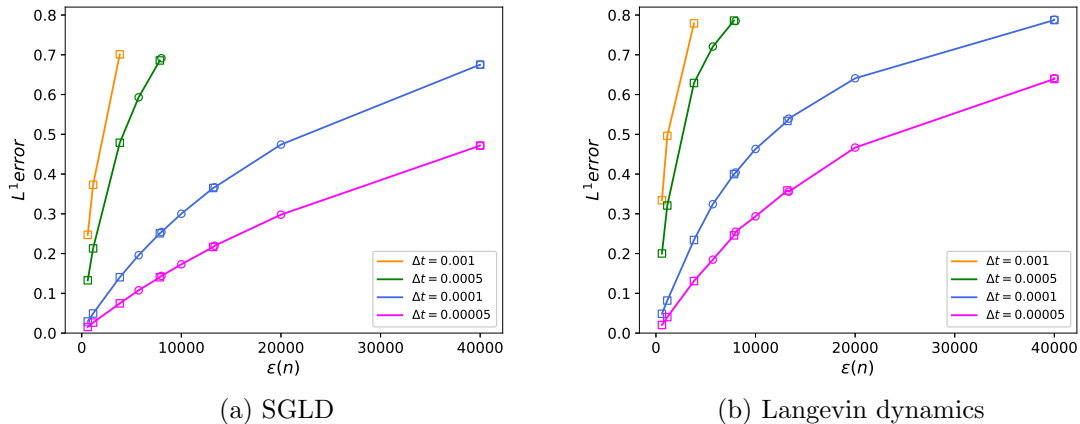


Figure 4: L^1 error on the θ_1 marginal of the posterior distribution for various values of Δt and n when the elementary likelihoods are mixtures of Gaussians, when sampling with (circles) and without replacement (squares).

integration time $T = 10^6$. We compute the L^1 error on the marginal distribution in the $\theta_1 = \mu_1$ variable, given in the limit $N_{\text{iter}} \rightarrow +\infty$ by

$$\int_{\mathbb{R}} \left| \int_{\mathbb{R}} \pi_{\Delta t, n}(\theta | \mathbf{x}) d\theta_2 - \int_{\mathbb{R}} \pi(\theta | \mathbf{x}) d\theta_2 \right| d\theta_1. \quad (44)$$

We plot this error with respect to $\varepsilon(n)$ for various values of Δt in Figure 4. We use numerical quadratures to approximate the integral with respect to θ_1 in (44) (approximating the marginals as piecewise constant functions over a grid of 500 bins over the interval $[0, 1.4]$), and also to compute the integral with respect to θ_2 for $\pi(\cdot | \mathbf{x})$.

The interpretation of the results presented in Figure 4 is quite similar to the discussion of the results of Figure 3. They confirm that the bias is determined by the value of $\varepsilon(n)$ irrespectively of the fact that sampling is performed with or without replacement. Note also that the errors are quite similar for SGLD and Langevin dynamics with $\Gamma = I_2$. Moreover, the L^1 error (44) is affine in $\varepsilon(n)$ provided $\varepsilon(n)\Delta t$ is sufficiently small so that the asymptotic analysis of the bias is indeed valid. From a quantitative viewpoint, the affine regime is observed for L^1 errors below 0.1. This regime is relevant for time steps $\Delta t \leq 10^{-4}$ for the values of n considered in our simulations, but not for simulations with larger time steps.

3. Adaptive Langevin dynamics

Using SGLD or Langevin dynamics with mini-batching introduces a bias on the posterior distribution. We recall in Section 3.1 the Adaptive Langevin (AdL) dynamics introduced in Jones and Leimkuhler (2011) and Ding et al. (2014). The aim of this dynamics is to remove, or at least substantially reduce the bias due to mini-batching. Under the key assumption that the covariance matrix $\Sigma_{\mathbf{x}}(\theta)$ defined in (12) is constant (but unknown), it can indeed be proved that Adaptive Langevin dynamics samples the target distribution (see Section 3.2). However, we demonstrate in Section 3.3 that this assumption may not hold

in practice, *e.g.* for models for which the elementary likelihood is a mixture of Gaussians, which motivates the construction of an extended version of AdL to tackle such cases. We also quantify the bias on the invariant measure due to the fact that $\Sigma_{\mathbf{x}}$ is not constant by deriving an error estimate involving (1), henceforth explaining why AdL reduces the minibatching bias compared to SGLD or Langevin dynamics. Our analysis also allows to understand why AdL with a scalar friction can provide results of almost the same quality as AdL with a genuine friction matrix. This is illustrated for a model of logistic regression for MNIST data in Section 3.4, and for Bayesian neural networks in Section 3.5.

3.1 General formulation of Adaptive Langevin dynamics

AdL was initially introduced to address the issue of the gradient of the energy not being exactly computed in molecular dynamics in Jones and Leimkuhler (2011). It was then considered for Bayesian inference in Ding et al. (2014) where it allows to remove the bias arising from mini-batching under the assumption that the covariance $\Sigma_{\mathbf{x}}$ is constant. As discussed in Section 2.4, the effect of the stochastic estimator (14) of the force in Langevin dynamics is, at dominant order, to add an unknown contribution to the diffusion coefficient in front of the Brownian motion, which we denote by $\sqrt{2}A_{\Delta t,n}(\theta)^{1/2}$, with

$$A_{\Delta t,n}(\theta) = \Gamma + \frac{\varepsilon(n)\Delta t}{2}\Sigma_{\mathbf{x}}(\theta). \quad (45)$$

Note that $A_{\Delta t,n}(\theta) \in \mathbb{R}^{d \times d}$ is an unknown positive definite symmetric matrix. The idea behind AdL is to modify the effective Langevin dynamics so that it admits $\pi(\cdot|\mathbf{x})$ as an invariant probability measure (more precisely, that it admits an invariant probability measure whose marginal distribution in the θ variable is $\pi(\cdot|\mathbf{x})$). The friction matrix, denoted here by $\xi \in \mathbb{R}^{d \times d}$, is no longer a constant, but a dynamical variable that adjusts itself to the effective noise resulting from mini-batching. The unknowns in the dynamics are therefore $(\theta, p, \xi) \in \Xi = \Theta \times \mathbb{R}^d \times \mathbb{R}^{d \times d}$. Denoting by $[\xi]_{i,j}$ the components of ξ , AdL can be written as

$$\begin{cases} d\theta_t = p_t dt, \\ dp_t = (\nabla_{\theta}(\log \pi(\theta_t|\mathbf{x}) - \xi_t p_t) dt + \sqrt{2}A_{\Delta t,n}(\theta_t)^{1/2} dW_t, \\ d[\xi_t]_{i,j} = \frac{1}{\eta}(p_{i,t}p_{j,t} - \delta_{i,j}) dt, \quad 1 \leq i, j \leq d, \end{cases} \quad (46)$$

where η is a positive scalar which sets the timescale for the evolution of the friction matrix (see Remark 8 below for more general choices). Let us emphasize once again that we consider that each degree of freedom has mass 1 for simplicity, but our analysis can be easily generalized to non trivial mass matrices, or even to non-quadratic kinetic energies as the ones considered in Stoltz and Trstanova (2018).

Remark 7 *If the initial condition ξ_0 is symmetric, then the matrix ξ_t is symmetric for all t . In any case, the quantities $\xi_{i,j,t} - \xi_{j,i,t}$ are constant. This shows that it suffices to introduce $(d+1)d/2$ new variables $([\xi]_{i,j})_{1 \leq i \leq j \leq d}$ to simulate AdL. However, for Lemma 9 below, it is more convenient to write out statements and proofs with all variables $([\xi]_{i,j})_{1 \leq i, j \leq d}$.*

Remark 8 *It is possible to formulate AdL for with different timescale parameters for the evolution of the components of the friction matrix, as defined by a symmetric matrix $\eta = \eta^T \in \mathbb{R}^{d \times d}$ with positive entries. In this case, each element i, j of ξ follows the dynamics*

$$d[\xi_t]_{i,j} = \frac{1}{\eta_{i,j}} (p_{i,t} p_{j,t} - \delta_{i,j}) dt.$$

Let us conclude this general presentation of AdL by recalling the motivation for the dynamics on the friction variable. We assume for this discussion that $d = 1$ in order to simplify the presentation. One way to understand the intuition behind AdL is to see (46) as the superposition of a Hamiltonian dynamics (which preserves the energy $-\log \pi(\theta|\mathbf{x}) + |p|^2/2$) and the following elementary dynamics:

$$\begin{cases} dp_t = -\xi_t p_t dt + \sqrt{2} A_{\Delta t, n}^{1/2} dW_t, \\ d\xi_t = \frac{1}{\eta} (p_t^2 - 1) dt. \end{cases}$$

Note that we write $A_{\Delta t, n}^{1/2}$ here instead of $A_{\Delta t, n}(\theta)^{1/2}$ since the value of θ does not change for this subdynamics. Knowing that the marginal distribution over the variable p of the invariant probability measure should be a Gaussian of variance 1 (see Ding et al. (2014)), the idea is to keep the right balance between the friction ξ and the fluctuation $A_{\Delta t, n}$. Given that the strength of the fluctuation is unknown, the friction is adjusted so that the average kinetic energy is fixed to its target value, *i.e.* $\mathbb{E}(|p|^2) = 1$. More precisely, if $|p|^2 > 1$, the kinetic energy is larger than what it should be, so the friction is increased, which ends up decreasing p . We can use the same line of argument for the opposite case. In fact, the dynamics of the additional variable ξ follows a negative feedback loop control as in the Nosé–Hoover thermostat; see Nosé (1984) and Hoover (1985).

3.2 Adaptive Langevin dynamics for gradient estimators with constant covariance

We discuss more precisely in this section the properties of AdL under the crucial assumption that the covariance of the gradient estimator is constant :

$$\Sigma_{\mathbf{x}}(\theta) = \Sigma_{\mathbf{x}}. \quad (47)$$

This implies in particular that the matrix $A_{\Delta t, n}$ defined in (45) is constant. The limitations of this assumption are discussed more precisely in Section 3.3.

We start by discussing invariant probability measures of AdL in Section 3.2.1. We next present in Section 3.2.2 numerical schemes to integrate AdL, together with error estimates on the bias of their invariant measures. These results are numerically illustrated in Section 3.2.3 for the Gaussian model of Section 2.5.1, for which the fundamental assumption (47) is satisfied. Let us emphasize here again that, in contrast to various works in the literature, we do not necessarily assume that n is much larger than 1 and that the random variable in (14) follows a Gaussian distribution.

3.2.1 INVARIANT PROBABILITY MEASURE OF ADL

The generator of the stochastic dynamics (46) can be decomposed as

$$\mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}} = \mathcal{L}_{\text{ham}} + \mathcal{L}_{\text{FD}} + \mathcal{L}_{\text{NH}},$$

where

$$\mathcal{L}_{\text{NH}} = -p^T (\xi - A_{\Delta t, n}) \nabla_p + \frac{1}{\eta} \sum_{1 \leq i, j \leq d} (p_i p_j - \delta_{i, j}) \partial_{[\xi]_{i, j}},$$

and

$$\mathcal{L}_{\text{ham}} = p^T \nabla_{\theta} + \nabla_{\theta} (\log \pi(\cdot | \mathbf{x}))^T \nabla_p, \quad \mathcal{L}_{\text{FD}} = A_{\Delta t, n} : \nabla_p^2 - p^T A_{\Delta t, n} \nabla_p.$$

We recall the following result on the invariant probability measure of (46); see Section 2 in Leimkuhler et al. (2020).

Lemma 9 *Suppose that (47) holds. Then the dynamics (46) admits the invariant probability measure*

$$\nu(d\theta dp d\xi) = \pi(\theta | \mathbf{x}) \tau(dp) \rho(d\xi) d\theta, \quad (48)$$

where $\tau(dp)$ is defined in (24) and

$$\rho(d\xi) = \prod_{1 \leq i, j \leq d} \sqrt{\frac{\eta}{2\pi}} \exp\left(-\frac{\eta}{2} \left(\xi_{i, j} - [A_{\Delta t, n}]_{i, j}\right)^2\right) d[\xi]_{i, j},$$

where $[A_{\Delta t, n}]_{i, j}$ is the (i, j) component of $A_{\Delta t, n}$.

Lemma 9 suggests that, as long as assumption (47) is satisfied, sampling a probability measure using Adaptive Langevin dynamics is not affected by the mini-batching procedure to estimate the gradient of the log-likelihood in the θ variable. The marginal distribution of (48) in the variable θ is indeed the target distribution π , whatever the value of $A_{\Delta t, n}$. This shows that AdL can indeed adjust the friction in order to compensate fluctuations of arbitrary constant magnitude. We recall the proof of Lemma 9 because we use similar computations in the proof of Theorem 17 below.

Remark 10 *Let us emphasize that Lemma 9 states that ν is invariant by AdL. However, the dynamics cannot be ergodic for this measure since $[\xi_t]_{i, j} - [\xi_t]_{j, i}$ remains constant, whereas $[\xi]_{i, j}$ and $[\xi]_{j, i}$ are independent under the probability measure (48). The dynamics can therefore at best be ergodic for the restriction of ν onto the sub-manifold*

$$\mathcal{S}(\xi_0) = \Theta \times \mathbb{R}^d \times \left\{ \xi \in \mathbb{R}^{d \times d} \mid [\xi]_{i, j} - [\xi]_{j, i} = [\xi_0]_{i, j} - [\xi_0]_{j, i} \right\},$$

which is determined by the initial condition ξ_0 . Such an ergodicity result is however not trivial at all since there are $d(d+1)$ independent degrees of freedom in the symmetric matrix ξ_t , while the noise acting on the momentum variable p is only of dimension d . The stochastic dynamics is therefore highly degenerate. In any case, the important point here is that the marginal in the θ -variable of the projected measure is $\pi(\cdot | \mathbf{x})$ whatever the distribution ρ in (48).

Proof We follow the approach in Section 2 of Leimkuhler et al. (2020). It suffices to show that, for all smooth and compactly supported functions ϕ ,

$$\int_{\Xi} \mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}} \phi \, d\nu = \int_{\Xi} \phi (\mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}}^* \mathbf{1}) \, d\nu = 0, \quad (49)$$

where adjoints are taken in $L^2(\nu)$. Simple computations based on integration by parts show that $\partial_{\theta_i}^* = -\partial_{\theta_i} - \partial_{\theta_i}(\log \pi(\theta|\mathbf{x}))$, $\partial_{p_i}^* = -\partial_{p_i} + p_i$, $\partial_{[\xi]_{i,j}}^* = -\partial_{[\xi]_{i,j}} + \eta([\xi]_{i,j} - [A_{\Delta t, n}]_{i,j})$. We can then rewrite the generators \mathcal{L}_{ham} and \mathcal{L}_{FD} as (see Leimkuhler et al. (2020)):

$$\mathcal{L}_{\text{ham}} = \sum_{i=1}^d \partial_{p_i}^* \partial_{\theta_i} - \partial_{\theta_i}^* \partial_{p_i}, \quad (50)$$

$$\mathcal{L}_{\text{FD}} = -\nabla_p^* A_{\Delta t, n} \nabla_p = - \sum_{1 \leq i, j \leq d} [A_{\Delta t, n}]_{i,j} \partial_{p_i}^* \partial_{p_j}. \quad (51)$$

Moreover, for ϕ, ψ two smooth and compactly supported functions,

$$\begin{aligned} \int_{\Xi} (\mathcal{L}_{\text{NH}} \phi) \psi \, d\nu &= \sum_{1 \leq i, j \leq d} \int_{\Xi} -p_i ([\xi]_{i,j} - [A_{\Delta t, n}]_{i,j}) (\partial_{p_j} \phi) \psi + \frac{1}{\eta} (p_i p_j - \delta_{i,j}) (\partial_{[\xi]_{i,j}} \phi) \psi \, d\nu \\ &= \sum_{1 \leq i, j \leq d} \int_{\Xi} -([\xi]_{i,j} - [A_{\Delta t, n}]_{i,j}) \partial_{p_j}^* (p_i \psi) \phi + \frac{1}{\eta} (p_i p_j - \delta_{i,j}) (\partial_{[\xi]_{i,j}}^* \psi) \phi \, d\nu \\ &= - \int_{\Xi} (\mathcal{L}_{\text{NH}} \phi) \psi \, d\nu. \end{aligned} \quad (52)$$

It is then clear that \mathcal{L}_{ham} and \mathcal{L}_{NH} are antisymmetric, while \mathcal{L}_{FD} is symmetric. Moreover, $\mathcal{L}_{\text{NH}} \mathbf{1} = \mathcal{L}_{\text{ham}} \mathbf{1} = \mathcal{L}_{\text{FD}} \mathbf{1} = 0$. The invariance of ν therefore follows from (49). \blacksquare

The mathematical properties of AdL are investigated in Herzog (2018) and Leimkuhler et al. (2020) in the case when $\Sigma_{\mathbf{x}} = \sigma^2 \mathbf{I}_d$ with σ^2 constant, and the friction is scalar valued (in which case the invariant probability measure provided by Lemma 9 is in fact the only invariant probability measure). The main contributions of Leimkuhler et al. (2020) are the following: (i) the exponential convergence of the law of the process encoded by the convergence of the semi-group $e^{t\mathcal{L}}$ is proved using the hypocoercive approach of Hérau (2006) and Dolbeault et al. (2009, 2015); (ii) a central limit theorem for time averages along one realization of the dynamics is derived with bounds on the asymptotic variance depending on the parameters η and Γ of the dynamics. If $A_{\Delta t, n} \approx \Gamma$ (which is the case when Δt is sufficiently small and n is sufficiently large), the mathematical analysis suggests to fix Γ and η of order 1.

3.2.2 NUMERICAL SCHEME

We now construct a numerical scheme for which AdL in (46) is the effective dynamics at dominant order in Δt and $\varepsilon(n)\Delta t$. Concretely, we replace the matrix $A_{\Delta t, n}$ by its expression (45) in the SDE (46) and consider the symmetric splitting scheme introduced in Leimkuhler et al. (2020), which is based on decomposing (46) into the following four

elementary SDEs (the variables which are evolved are indexed by t , while the ones which remain constant do not have any subscript):

$$d\theta_t = p dt, \quad (53)$$

$$dp_t = -\xi p_t dt + \sqrt{2}\Gamma^{1/2} dW_{1,t}, \quad (54)$$

$$dp_t = \nabla_\theta(\log \pi(\theta|\mathbf{x})) dt + \sqrt{\varepsilon(n)\Delta t}\Sigma_\mathbf{x}(\theta)^{1/2} dW_{2,t}, \quad (55)$$

$$d[\xi_t]_{i,j} = \frac{1}{\eta} (p_i p_j - \delta_{i,j}) dt, \quad 1 \leq i \leq j \leq d, \quad (56)$$

where $W_{1,t}, W_{2,t}$ are two independent standard d -dimensional Brownian motions. The elementary ordinary differential equations (53) and (56) can be trivially integrated. The elementary SDE (54) is an Ornstein–Uhlenbeck process that can be analytically integrated in law as

$$p^{m+1} = e^{-\Delta t\xi} p^m + \sigma_{\xi,\Gamma,\Delta t} G^m, \quad \sigma_{\xi,\Gamma,\Delta t}^2 = 2 \int_0^{\Delta t} e^{-s\xi} \Gamma e^{-s\xi} ds,$$

where $(G^m)_{m \geq 0}$ is a family of i.i.d. standard d -dimensional Gaussian random variables. If there exists M such that $\xi M + M\xi = \Gamma$, then $\sigma_{\xi,\Gamma,\Delta t}^2 = 2(M - e^{-\Delta t\xi} M e^{-\Delta t\xi})$. In particular, for $\Gamma = \gamma I_d$ and when ξ is invertible, one can choose $M = \gamma\xi^{-1}/2$ in which case $\sigma_{\xi,\Gamma,\Delta t}^2 = \gamma\xi^{-1}(I_d - e^{-2\Delta t\xi})$. When ξ is singular or close to singular, the latter formula has to be understood through a limiting procedure based on spectral calculus, see Leimkuhler et al. (2020).

The elementary SDE (55) is integrated as

$$p^{m+1} = p^m + \Delta t \widehat{F}_n(\theta) = p^m + \Delta t \nabla_\theta(\log \pi(\theta|\mathbf{x})) + \Delta t \sqrt{\varepsilon(n)} \Sigma_\mathbf{x}(\theta)^{1/2} Z_{\mathbf{x}, N_{\text{data}}, n}. \quad (57)$$

Since the random variable $Z_{\mathbf{x}, N_{\text{data}}, n}$ has mean 0 and identity covariance by construction, the equality (82) in Appendix A.2 shows that the numerical scheme (57) is weakly consistent with (55), with an error of order $(1 + \varepsilon(n)^{3/2})\Delta t^3$ over one time step, even if $Z_{\mathbf{x}, N_{\text{data}}, n}$ is not Gaussian. In the case where $\mathbb{E}[Z_{\mathbf{x}, N_{\text{data}}, n}^3] = 0$, the error over one time step is of order $(1 + \varepsilon(n))\Delta t^3$.

The numerical scheme we consider is finally obtained by a Strang splitting where (55) is updated in the central step of the algorithm, in order to compute the force only once per time step and avoid its storage. The order in which the remaining elementary dynamics are integrated is unimportant for our purposes, although some orderings may be better than others, in particular in some limiting regimes where one of the parameters goes to 0 or infinity; see Leimkuhler and Shang (2016) for an extensive discussion in the context of

AdL. Fixing $\Gamma = \gamma \text{Id}$, the numerical scheme reads as follows:

$$\left\{ \begin{array}{l} p^{m+\frac{1}{2}} = e^{-\Delta t \xi^m / 2} p^m + \left[\gamma (\xi^m)^{-1} \left(\text{Id} - e^{-\Delta t \xi^m} \right) \right]^{1/2} G^m, \\ \xi^{m+\frac{1}{2}} = \xi^m + \frac{\Delta t}{2\eta} \left(p^{m+\frac{1}{2}} \left(p^{m+\frac{1}{2}} \right)^T - \text{Id} \right), \\ \theta^{m+\frac{1}{2}} = \theta^m + \frac{\Delta t}{2} p^{m+\frac{1}{2}}, \\ \widehat{p}^{m+\frac{1}{2}} = p^{m+\frac{1}{2}} + \Delta t \widehat{F}_n \left(\theta^{m+\frac{1}{2}} \right), \\ \theta^{m+1} = \theta^{m+\frac{1}{2}} + \frac{\Delta t}{2} \widehat{p}^{m+\frac{1}{2}}, \\ \xi^{m+1} = \xi^{m+\frac{1}{2}} + \frac{\Delta t}{2\eta} \left(\widehat{p}^{m+\frac{1}{2}} \left(\widehat{p}^{m+\frac{1}{2}} \right)^T - \text{Id} \right), \\ p^{m+1} = e^{-\Delta t \xi^{m+1} / 2} \widehat{p}^{m+\frac{1}{2}} + \left[\gamma (\xi^{m+1})^{-1} \left(\text{Id} - e^{-\Delta t \xi^{m+1}} \right) \right]^{1/2} G^{m+\frac{1}{2}}, \end{array} \right. \quad (58)$$

where $(G^m)_{m \geq 0}$ and $(G^{m+\frac{1}{2}})_{m \geq 0}$ are two independent families of i.i.d. standard Gaussian random variables. Note that it is possible to work only with the additional variables $(\xi_{i,j})_{1 \leq i \leq j \leq d}$ since the updates of $\xi_{i,j}$ and $\xi_{j,i}$ in (46) are the same. The initial conditions for ξ^0 is γId , while the components of θ^0, p^0 are initialized to 0. Of course, more educated choices can be considered depending on the system at hand (*e.g.* restarting from values sampled by SGLD or some Langevin-like dynamics).

Remark 11 *Other versions of the numerical scheme (58) can be considered, in particular the one used in Leimkuhler et al. (2020), for which $\xi \in \mathbb{R}$. In this case, (56) should be replaced by*

$$d\xi_t = \frac{1}{\eta} (p^T p - d) dt. \quad (59)$$

Another option to reduce the number of additional variables is to consider the variable ξ as a diagonal matrix with entries ξ_i . In this case, (56) has to be replaced with

$$d[\xi_t]_i = \frac{1}{\eta} (p_i^2 - 1) dt, \quad 1 \leq i \leq d. \quad (60)$$

The numerical schemes for these two cases can be obtained by an obvious modification of the right hand side of the update formulas for $\xi^{m+\frac{1}{2}}$ and ξ^{m+1} in (58).

Using the Baker–Campbell–Hausdorff formula (see for instance Section III.4.2 in Hairer et al. (2006)), as done in Leimkuhler et al. (2016) and Leimkuhler and Shang (2016) for instance, one can prove that the evolution operator $\widehat{P}_{\Delta t, n}$ of the numerical scheme (58) satisfies (see Appendix A.3)

$$\widehat{P}_{\Delta t, n} = e^{\Delta t \mathcal{L}_{\text{AdL}, \Sigma_x}} + \mathcal{O} \left((1 + \varepsilon(n)^{3/2}) \Delta t^3 \right). \quad (61)$$

Then, by following the strategy of proof of Theorem 3.3 in Lelièvre and Stoltz (2016), the following error estimate holds.

Proposition 12 *Assume that (46) and (58) both admit a unique invariant probability measure, respectively denoted by ν_{ξ_0} and $\nu_{\xi_0, \Delta t, n}$. Then, for any smooth function φ with compact support, there exist $C \in \mathbb{R}_+$ and $\Delta t_\star > 0$ such that, for any $0 < \Delta t \leq \Delta t_\star$,*

$$\left| \int_{\Xi} (\mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}} \varphi)(\theta, p, \xi) \nu_{\xi_0, \Delta t, n}(d\theta dp d\xi) - \int_{\Xi} (\mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}} \varphi)(\theta, p, \xi) \nu_{\xi_0}(d\theta dp d\xi) \right| \leq C \left(1 + \varepsilon(n)^{3/2}\right) \Delta t^2. \quad (62)$$

In view of (62), the order of magnitude of the bias on the invariant probability measure is of order $\varepsilon(n)^{3/2} \Delta t^2$, which is smaller than the bias $\varepsilon(n) \Delta t$ obtained with SGLD or Langevin (recall Propositions 2 and 4). Even if the estimate is stated for averages of functions $\mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}} \varphi$, this bound allows to uniquely characterize the bias on the invariant measure. To state it for any smooth function φ (and not just functions in the range of the operator $\mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}}$), one would need results ensuring that the inverse of $\mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}}$ is well defined, and that it stabilizes spaces of smooth functions with polynomial growth and whose derivatives also grow at most polynomially. This would be the counterpart of similar estimates for Langevin like dynamics (see Kopec (2015a,b)). Such results are however not available, and are presumably rather difficult to prove. At the moment, as discussed at the end of Section 3.2.1, the only available ergodicity results on AdL concern the exponential convergence of the semigroup, which ensures that the inverse of the generator is well defined on certain functional spaces, such as $L^2(\nu)$ (see Herzog (2018); Leimkuhler et al. (2020)).

Note also that the error estimate (62) holds even when $Z_{\mathbf{x}, N_{\text{data}}, n}$ is not Gaussian (namely for small values of n). When $\mathbb{E}[Z_{\mathbf{x}, N_{\text{data}}, n}^3] = 0$, the error estimate (62) holds with $C(1 + \varepsilon(n)) \Delta t^2$ on the right hand side. When $Z_{\mathbf{x}, N_{\text{data}}, n}$ is Gaussian, the integration (57) is in fact exact in law. In any case, the analysis we made crucially depends on the fact that (47) holds. When it does not hold, there is an additional bias on the invariant measure due to mini-batching of order $\varepsilon(n) \Delta t$, much larger than the error of order $(1 + \varepsilon(n)^{3/2}) \Delta t^2$ in (62). This is quantified in Section 3.3.

3.2.3 NUMERICAL ILLUSTRATION FOR GAUSSIAN LIKELIHOODS

We consider the same setting as in Section 2.5.1. The covariance $\Sigma_{\mathbf{x}}$, given by (40), is constant for this model, so that assumption (47) indeed holds. We run the numerical scheme (58) with $\gamma = 1$ for an integration time $T = 10^6$. Similar computations have already been performed in Ding et al. (2014) and Leimkuhler and Shang (2016). In Figure 5, we plot the error on the variance of the marginal posterior distribution in the θ variable with respect to $\varepsilon(n)$. First, we note that, for Δt small enough (namely $\Delta t \leq 0.005$), the bias is not affected by the value of $\varepsilon(n)$ (even for the very large value N_{data}^2 obtained when $n = 1$), and hence by the size of the mini-batch. The bias is therefore simply due to the discretization time step. For large values of Δt (namely $\Delta t \geq 0.008$ in our experiments), the error is affected by the value of n . This is probably due to the fact that this rather large value of Δt is out of the regime where the asymptotic analysis for the bias holds. In any case, a comparison with Figure 3 shows that the error obtained with AdL is much smaller than the one obtained with SGLD (18) or the numerical scheme (30) (which corresponds to Langevin dynamics with mini-batching), even with much larger time steps.




Figure 5: Relative error on the variance of the posterior distribution for various values of n using AdL when the elementary likelihood is Gaussian, and sampling is performed without replacement.

3.3 Impact of a non constant covariance matrix

The bias analysis in the previous section does not hold in the case of a non constant covariance matrix $\Sigma_{\mathbf{x}}$. We first quantify in Section 3.3.1 the bias due to mini-batching when assumption (47) is not satisfied. We illustrate our analysis in Section 3.3.2 with a numerical example where the elementary likelihood is a mixture of Gaussians. In this case, $\Sigma_{\mathbf{x}}(\theta)$ genuinely depends on θ .

3.3.1 MINI-BATCHING BIAS FOR ADAPTIVE LANGEVIN DYNAMICS AND NON CONSTANT COVARIANCE

We consider the situation where $\Sigma_{\mathbf{x}}(\theta)$ genuinely depends on θ , and prove that the bias on the invariant measure is of order $\varepsilon(n)\Delta t + \Delta t^2$. This is done in two steps. We first characterize the bias of the continuous Adaptive Langevin dynamics with covariance matrix $\Sigma_{\mathbf{x}}(\theta)$ instead of the constant covariance matrix

$$\bar{\Sigma}_{\mathbf{x}} = \int_{\Theta} \Sigma_{\mathbf{x}}(\theta) \pi(\theta|\mathbf{x}) d\theta \quad (63)$$

obtained by averaging over the target probability distribution. We then combine this estimate with (62).

The motivation for introducing the particular constant matrix (63) is discussed after Lemma 14 below. The generator of (46) can be written in terms of $\bar{\Sigma}_{\mathbf{x}}$ as

$$\mathcal{L}_{\text{AdL}, \Sigma_{\mathbf{x}}} = \mathcal{L}_{\text{AdL}, \bar{\Sigma}_{\mathbf{x}}} + \varepsilon(n)\Delta t \tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}}, \quad \tilde{\mathcal{L}}_{\mathcal{M}} = \mathcal{M} : \nabla_p^2. \quad (64)$$

Since $\bar{\Sigma}_{\mathbf{x}}$ does not depend on θ , the probability measure ν defined in (48) is an invariant probability measure of the dynamics with generator $\mathcal{L}_{\text{AdL}, \bar{\Sigma}_{\mathbf{x}}}$. It is then easy to prove that

the extra bias due to mini-batching on the invariant measure of (46) when $\Sigma_{\mathbf{x}}$ is not constant is of order $\varepsilon(n)\Delta t$.

Proposition 13 *Assume that the continuous dynamics (46) for $\Sigma_{\mathbf{x}}$ depending on the θ variable, and the same dynamics (46) with constant covariance set to $\bar{\Sigma}_{\mathbf{x}}$, both admit unique invariant probability measures, respectively denoted by ν_{ξ_0} and $\bar{\nu}_{\xi_0}$; and that the resolvent $\mathcal{L}_{\text{AdL},\bar{\Sigma}_{\mathbf{x}}}^{-1}$ is bounded on the subspace of functions in $L^2(\bar{\nu}_{\xi_0})$ with average 0 with respect to $\bar{\nu}_{\xi_0}$. Then, for any smooth function φ with compact support,*

$$\begin{aligned} & \int_{\Xi} (\mathcal{L}_{\text{AdL},\Sigma_{\mathbf{x}}}\varphi)(\theta, p, \xi) \nu_{\xi_0}(d\theta dp d\xi) - \int_{\Xi} (\mathcal{L}_{\text{AdL},\Sigma_{\mathbf{x}}}\varphi)(\theta, p, \xi) \bar{\nu}_{\xi_0}(d\theta dp d\xi) \\ &= -\varepsilon(n) \int_{\Xi} (\mathcal{L}_{\text{AdL},\Sigma_{\mathbf{x}}}\varphi)(\theta, p, \xi) f_{\text{AdL},\Sigma_{\mathbf{x}}}(\theta, p, \xi) \bar{\nu}_{\xi_0}(d\theta dp d\xi), \end{aligned} \quad (65)$$

where (adjoints being taken on $L^2(\bar{\nu}_{\xi_0})$)

$$f_{\text{AdL},\Sigma_{\mathbf{x}}} = \left(-\mathcal{L}_{\text{AdL},\bar{\Sigma}_{\mathbf{x}}}^* \right)^{-1} \tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}}-\bar{\Sigma}_{\mathbf{x}}}^* \mathbf{1}. \quad (66)$$

The ergodicity assumption on the dynamics are not trivial statements, as discussed in Remark 10. The boundedness of the resolvent has been proved when ξ is scalar valued, see Leimkuhler et al. (2020). If this result could be extended to $\mathcal{L}_{\text{AdL},\Sigma_{\mathbf{x}}}$, with further estimates on the derivatives (similary to the ones discussed after Proposition 12), it would be possible to replace $\mathcal{L}_{\text{AdL},\Sigma_{\mathbf{x}}}\varphi$ by φ in (65). In contrast to (overdamped) Langevin dynamics, the dimension dependence of (66) is currently not known (see however partial elements in (Bernard et al., 2022, Section 3.3)).

Proof Note first that, by the assumed invariance of ν_{ξ_0} by (46), it holds

$$\int_{\Xi} (\mathcal{L}_{\text{AdL},\Sigma_{\mathbf{x}}}\varphi)(\theta, p, \xi) \nu_{\xi_0}(d\theta dp d\xi) = 0.$$

Moreover, using the invariance of $\bar{\nu}_{\xi_0}$ under the dynamics (46) with constant covariance set to $\bar{\Sigma}_{\mathbf{x}}$,

$$\begin{aligned} \int_{\Xi} (\mathcal{L}_{\text{AdL},\Sigma_{\mathbf{x}}}\varphi)(\theta, p, \xi) \bar{\nu}_{\xi_0}(d\theta dp d\xi) &= \int_{\Xi} \left(\mathcal{L}_{\text{AdL},\bar{\Sigma}_{\mathbf{x}}}\varphi \right)(\theta, p, \xi) \bar{\nu}_{\xi_0}(d\theta dp d\xi) \\ &\quad + \varepsilon(n)\Delta t \int_{\Xi} \left(\tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}}-\bar{\Sigma}_{\mathbf{x}}}\varphi \right)(\theta, p, \xi) \bar{\nu}_{\xi_0}(d\theta dp d\xi) \\ &= \varepsilon(n)\Delta t \int_{\Xi} \varphi \tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}}-\bar{\Sigma}_{\mathbf{x}}}^* \mathbf{1} d\bar{\nu}_{\xi_0} \\ &= -\varepsilon(n)\Delta t \int_{\Xi} \left[\mathcal{L}_{\text{AdL},\bar{\Sigma}_{\mathbf{x}}}\varphi \right] \left[\left(-\mathcal{L}_{\text{AdL},\bar{\Sigma}_{\mathbf{x}}}^* \right)^{-1} \tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}}-\bar{\Sigma}_{\mathbf{x}}}^* \mathbf{1} \right] d\bar{\nu}_{\xi_0}, \end{aligned}$$

where leads to the claimed equality in view of the definition of $f_{\text{AdL},\Sigma_{\mathbf{x}}}$. \blacksquare

The estimate (65) shows that the difference between ν_{ξ_0} and $\bar{\nu}_{\xi_0}$ is of order $\varepsilon(n)\Delta t$, with a magnitude related to the norm of $f_{\text{AdL},\Sigma_{\mathbf{x}}}$ in $L^2(\bar{\nu}_{\xi_0})$. The latter quantity, estimated in the following lemma, is typically smaller than similar quantities for Langevin like dynamics such as (37).

Lemma 14 *Suppose that the same assumptions as in Proposition 13 hold true, and that the marginal distributions of $\bar{\nu}_{\varepsilon_0}$ in the θ and p variables are respectively $\pi(\cdot|\mathbf{x})$ and τ . Then, there exists $C \in \mathbb{R}_+$ such that*

$$\|f_{\text{AdL}, \Sigma_{\mathbf{x}}}\|_{L^2(\bar{\nu}_{\varepsilon_0})} \leq C \|\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}\|_{L^2(\pi)}. \quad (67)$$

Lemma 14 provides a motivation for the choice of $\bar{\Sigma}_{\mathbf{x}}$. Indeed, the error estimate in Lemma 14 would be true with $\bar{\Sigma}_{\mathbf{x}}$ replaced by any constant, positive, symmetric matrix $M \in \mathbb{R}^{d \times d}$. The choice leading to the sharpest upper bound is the $L^2(\pi)$ -orthogonal projection $\bar{\Sigma}_{\mathbf{x}}$ of $\Sigma_{\mathbf{x}}$ onto constant matrices.

Proof Note first that $\tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}}^* \mathbf{1} = (\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}) : (\nabla_p^2)^* \mathbf{1}$. The (i, j) -component of $(\nabla_p^2)^* \mathbf{1}$ is $\partial_{p_i}^* \partial_{p_j}^* \mathbf{1} = p_i p_j - \delta_{i,j}$. Since

$$\int_{\mathbb{R}^d} (p_i p_j - \delta_{i,j})^2 \tau(p) dp = 1 + \delta_{i,j},$$

we obtain, by a Cauchy–Schwarz inequality,

$$\|\tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}}^* \mathbf{1}\|_{L^2(\bar{\nu}_{\varepsilon_0})}^2 \leq \|\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}\|_{L^2(\pi)}^2 \sum_{i,j=1}^d \int_{\mathbb{R}^d} (p_i p_j - \delta_{i,j})^2 \tau(p) dp = d(d+1) \|\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}\|_{L^2(\pi)}^2.$$

The conclusion then follows from the definition (66) and the assumed boundedness of the resolvent. \blacksquare

We finally deduce from (62) and (65) that the total error between the invariant probability measures for (46) and (58) is of order $\varepsilon(n)\Delta t + \Delta t^2$. However, as motivated by Lemma 14, the prefactor for the error term $\varepsilon(n)\Delta t$ is much smaller than the one for SGLD and Langevin dynamics, and depends only on the deviation of the covariance of the gradient from its average.

Remark 15 *Using other versions of the numerical scheme (58) (see Remark 11) affects the prefactor for the error term $\varepsilon(n)\Delta t$. If we consider the case where the variable ξ is a scalar, in other words, if we consider (59) for the numerical scheme, $\bar{\Sigma}_{\mathbf{x}}$ in (63) should be replaced by $s^* \text{Id}_d$, with*

$$s^* = \frac{1}{d} \int_{\Theta} \text{Tr}(\Sigma_{\mathbf{x}}(\theta)) \pi(\theta|\mathbf{x}) d\theta.$$

In this case, the prefactor in front of the minibatching error term is larger than $\|\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}\|_{L^2(\pi)}$, as it is given by

$$\|\Sigma_{\mathbf{x}} - s^* \text{Id}_d\|_{L^2(\pi)} = \min_{s \in \mathbb{R}} \|\Sigma_{\mathbf{x}} - s \text{Id}_d\|_{L^2(\pi)}.$$

The latter quantity corresponds to the projection error onto the class of isotropic matrices.

If we consider the variable ξ to be a diagonal matrix, i.e. we consider (60) for the numerical scheme, $\bar{\Sigma}_{\mathbf{x}}$ in (63) should be replaced by a diagonal matrix with entries

$$D_i^* = \int_{\Theta} [\Sigma_{\mathbf{x}}(\theta)]_{i,i} \pi(\theta|\mathbf{x}) d\theta, \quad 1 \leq i \leq d.$$

The upper bound in Lemma 14 then corresponds to the projection error of $\Sigma_{\mathbf{x}}$ onto constant diagonal matrices:

$$\|\Sigma_{\mathbf{x}} - D^*\|_{L^2(\pi)} = \min_{(D_1, \dots, D_d) \in \mathbb{R}^d} \|\Sigma_{\mathbf{x}} - \text{diag}(D_1, \dots, D_d)\|_{L^2(\pi)}.$$

By interpreting the various $L^2(\pi)$ -norms of the difference between $\Sigma_{\mathbf{x}}$ and a constant matrix as the distance to finite dimensional subspaces of $L^2(\pi)$ included in each other, one can conclude that

$$\|\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}\|_{L^2(\pi)} \leq \|\Sigma_{\mathbf{x}} - D^*\|_{L^2(\pi)} \leq \|\Sigma_{\mathbf{x}} - s^* \mathbf{I}_d\|_{L^2(\pi)},$$

i.e. the projection error obtained for isotropic matrices is the largest, followed by the projection error for diagonal matrices, the smallest projection error being obtained for genuine matrices.

3.3.2 MIXTURE OF GAUSSIANS

We numerically illustrate here that the bias on the posterior distribution is indeed of order $\varepsilon(n)\Delta t + \Delta t^2$ when (47) does not hold. We consider to this end the model described in Section 2.5.1 with the same parameters as in that section. Assuming that the data points x_i are distributed according to $P_{\text{elem}}(\cdot|\theta_0)$ for some θ_0 , the covariance $\Sigma_{\mathbf{x}}(\theta)$ has the following limit when the number of data points N_{data} is large:

$$\begin{aligned} \Sigma(\theta) &:= \lim_{N_{\text{data}} \rightarrow \infty} \Sigma_{\mathbf{x}}(\theta) \\ &= \int_{\mathcal{X}} \nabla_{\theta}(\log P_{\text{elem}}(x|\theta)) \nabla_{\theta}(\log P_{\text{elem}}(x|\theta))^T P_{\text{elem}}(x|\theta) dx \\ &\quad - \left(\int_{\mathcal{X}} \nabla_{\theta}(\log P_{\text{elem}}(x|\theta)) P_{\text{elem}}(x|\theta) dx \right) \left(\int_{\mathcal{X}} \nabla_{\theta}(\log P_{\text{elem}}(x|\theta)) P_{\text{elem}}(x|\theta) dx \right)^T. \end{aligned} \tag{68}$$

This limit is well defined provided $\nabla_{\theta}(\log P_{\text{elem}}(\cdot|\theta)) \in L^2(P_{\text{elem}}(\cdot|\theta_0))$ for all $\theta \in \Theta$. Computing $\Sigma(\theta)$ in (68) is intractable in general for two reasons: (i) θ_0 is unknown and anyway the data points may even not be distributed according to $P_{\text{elem}}(\cdot|\theta_0)$; (ii) the integral over x is possibly a high dimensional one or requires an evaluation cost of $\mathcal{O}(N_{\text{data}})$ to approximate it. We nonetheless use the formula (68) to compute the elements of $\Sigma(\theta)$ in the low-dimensional case considered here to better understand the issues at stake. We plot in Figure 6 the components $\Sigma_{11}(\theta)$, $\Sigma_{12}(\theta)$ and $\Sigma_{22}(\theta)$ of the symmetric matrix $\Sigma(\theta)$ in (68). It is clear that the variance genuinely depends on the value of the parameter $\theta = (\theta_1, \theta_2)$, so that assumption (47) fails in this case.

To check whether the failure of assumption (47) has an impact on the properties of AdL, we use the data set described in Section 2.5.2. We run the numerical scheme (58) for the following two cases: when ξ is a $d \times d$ matrix or when it is a scalar. We fix $T = 10^6$, $\gamma = 1$ and $\eta = 0.1$. This last choice is motivated by metastability issues we noticed while using AdL for larger values of η . We report in Figure 7 the L^1 error given by (44). Note first that using AdL greatly reduces the L^1 error compared to the results

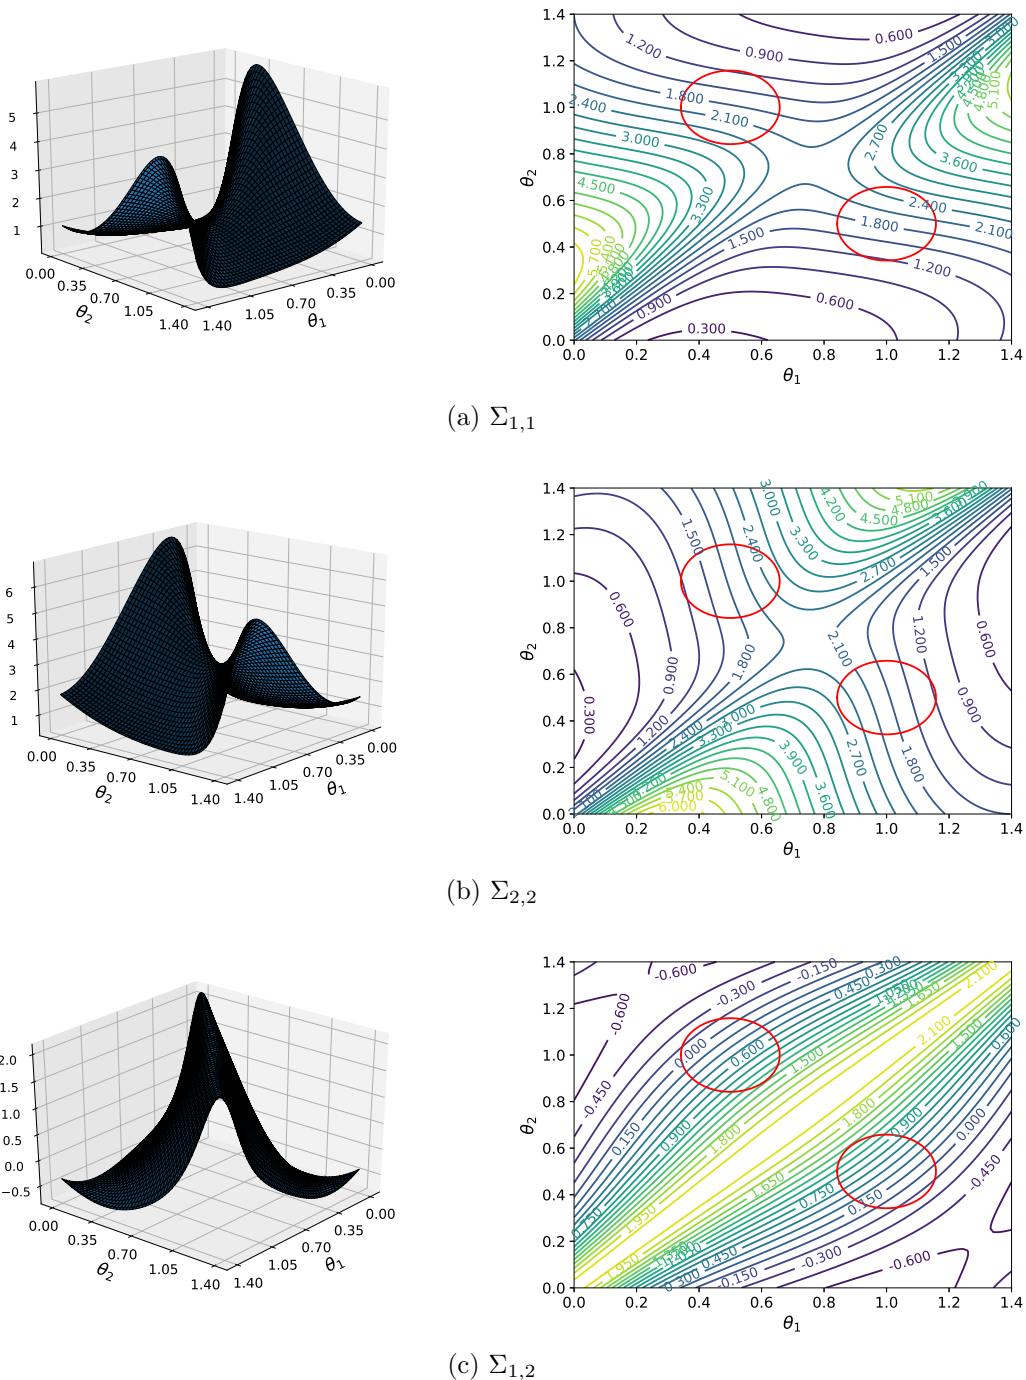


Figure 6: The three components of the elementary covariance matrix Σ in (68) when the elementary likelihood are mixture of Gaussians: surface (left) and contour (right) plots. The modes of the posterior distribution are indicated by orange ellipses.




Figure 7: L^1 error on the posterior distribution for various values of Δt and n when sampling from the posterior distribution for the mixture of Gaussians case using AdL, with mini-batching performed without replacement.

of Figure 4, obtained with SGLD (18) or the numerical scheme (30) (which corresponds to Langevin dynamics with mini-batching). This allows to run simulations with much larger timesteps Δt . However, AdL, in both cases, fails to completely correct the bias due to mini-batching. Consistently with the analysis of Section 3.3.1, we observe that the bias seems to scale linearly with $\varepsilon(n)\Delta t \|\Sigma_{\mathbf{x}}(\theta) - S^*\|_{L^2(\pi)}$ (where S^* is the $L^2(\pi)$ -projection of $\Sigma_{\mathbf{x}}$ onto the set of admissible matrices) when this quantity is sufficiently small. We numerically compute

$$\left\| \Sigma_{\mathbf{x}}(\theta) - \frac{1}{d} \int_{\Theta} \text{Tr}(\Sigma_{\mathbf{x}}(\theta')) \pi(\theta'|\mathbf{x}) d\theta' \right\|_{L^2(\pi)} \approx 1.75$$

for the scalar case, and

$$\left\| \Sigma_{\mathbf{x}}(\theta) - \int_{\Theta} \Sigma_{\mathbf{x}}(\theta') \pi(\theta'|\mathbf{x}) d\theta' \right\|_{L^2(\pi)} \approx 0.57$$

for the full matrix case. This is consistent with the numerical results for $\varepsilon(n)\Delta t \leq 100$, where the error on the posterior distribution in the scalar case is roughly 3 times larger than the error in the full matrix case.

3.4 Lack of reduction of the bias for a model of logistic regression

The estimate of Lemma 14 suggests that the reduction in the minibatching error is directly related to the improved quality of the approximation of the covariance matrix. We demonstrate in this section on the example of a Bayesian logistic regression model that the reduction may not be substantial in certain cases when passing from a scalar friction to a genuinely matricial one. This motivates using scalar AdL as this method has a smaller computational cost, and explains why previous numerical results obtained with scalar AdL were of good quality.

Presentation of the model. We consider a subset of the MNIST data set containing the digits 7 and 9, and which have been pre-processed by a principal component analysis, as described in Section 4.3 of Leimkuhler et al. (2020). More precisely, there are $N_{\text{data}} = 12,251$ scores on the 100 first principal component $\mathbf{z} = (z_1, z_2, \dots, z_{N_{\text{data}}}) \in \mathbb{R}^d$ (with $d = 100$), labeled by $\mathbf{y} = (y_1, y_2, \dots, y_{N_{\text{data}}}) \in \{0, 1\}$ (the labels 0 and 1 correspond respectively to digits 7 and 9). We assume that the elementary likelihood on the data point $x_i = (z_i, y_i)$ is

$$P_{\text{elem}}(x_i|\theta) = \sigma(\theta^T z_i)^{y_i} (1 - \sigma(\theta^T z_i))^{1-y_i},$$

where

$$\sigma(z) = \frac{\exp(z)}{1 + \exp(z)}, \quad (69)$$

and $\theta \in \mathbb{R}^d$ is the vector of parameters we want to estimate. Assuming that the prior over the vector of parameter θ is a centered standard normal distribution, a simple computation based on the identity $\sigma'(z) = \sigma(z)(1 - \sigma(z))$ shows that

$$\widehat{F}_n(\theta) = \theta + \frac{N_{\text{data}}}{n} \sum_{i \in I_n} (y_i - \sigma(\theta^T z_i)) z_i.$$

Numerical results. We use AdL to sample from the posterior distribution, with $\gamma = \eta = 1$, $T = 10^3$ and $\Delta t = 10^{-3}$. We run the numerical scheme (58) for the following three cases: ξ is a scalar, a diagonal matrix or a full matrix. We plot the marginal distributions of some parameters θ_i in Figure 8. The results suggest that the posterior distribution of the parameters is close to a Gaussian distribution. As pointed out in Section 3.2.3, AdL is sufficient in this case to remove the minibatching error.

According to Lemma 9, the posterior distribution of the variable ξ is centered on $A_{\Delta t, n}$ given by (45), when the latter matrix is constant. An estimate of the average covariance matrix can then be obtained with the estimator

$$\frac{1}{\varepsilon(n)\Delta t} \left(\frac{1}{N_{\text{iter}}} \sum_{m=1}^{N_{\text{iter}}} \xi^m - \gamma I_d \right),$$

where γI_d is replaced by γ in the scalar case. We plot in Figure 9 the diagonal components of the estimate of the average of the covariance matrix of the stochastic gradient for the following three cases considered here (ξ scalar, diagonal of full matrix). We can see that $\overline{\Sigma}_{\mathbf{x}}$ has a dominant part which is equal to $s^* I_d$ (although there are many non zero off diagonal terms). In view of Remark 15, we expect AdL to give comparable results whether the friction variable is a scalar, a diagonal matrix and a full matrix. In order to investigate this, we plot in Figure 10 the mean relative error on the diagonal entries of the covariance matrix of the vector of parameters θ with respect to $\varepsilon(n)$, given by

$$\frac{1}{d} \sum_{i=1}^d \left| \frac{[\text{cov}(\theta)_n]_{i,i} - [\text{cov}(\theta)_{N_{\text{data}}}]_{i,i}}{[\text{cov}(\theta)_{N_{\text{data}}}]_{i,i}} \right|,$$

where $\text{cov}(\theta)_n$ and $\text{cov}(\theta)_{N_{\text{data}}}$ denotes respectively the covariance of the vector of parameters θ along trajectories of AdL with a minibatch of size n , and AdL with a computation

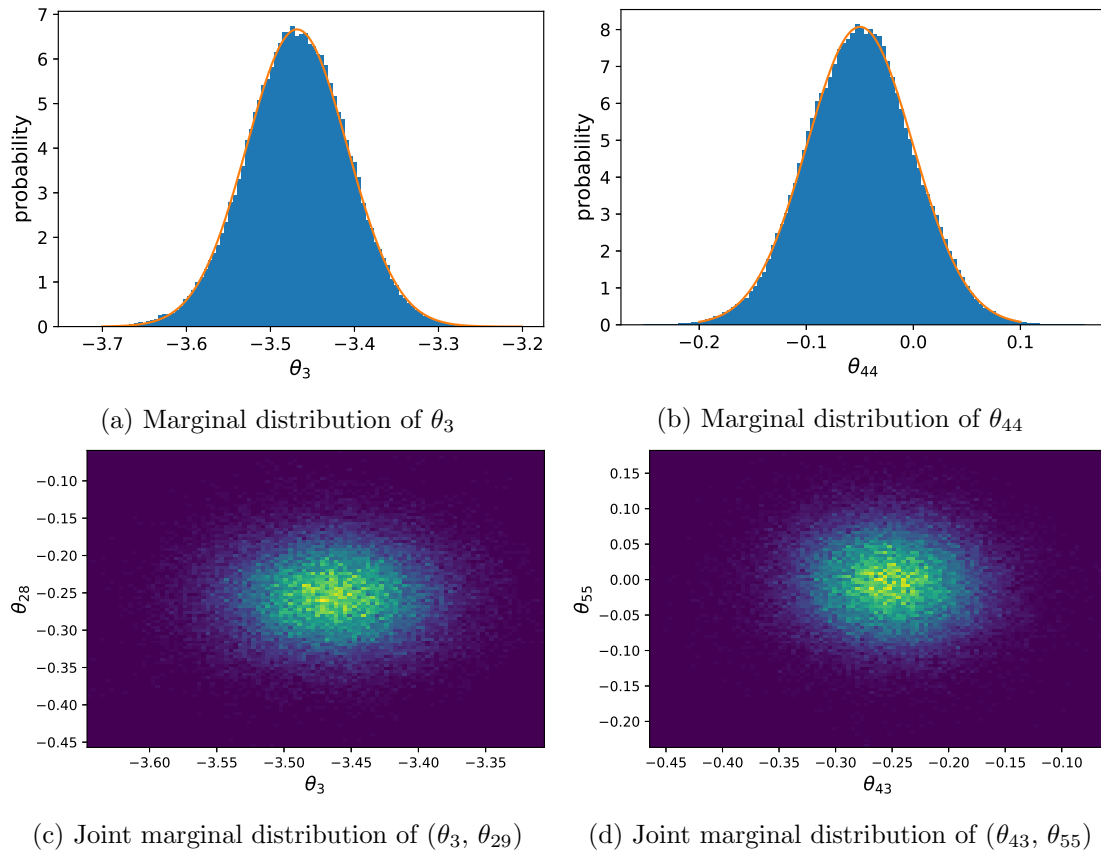


Figure 8: Marginal posterior distributions of θ_3 and θ_{44} (top), and of (θ_3, θ_{28}) and $(\theta_{43}, \theta_{55})$ (bottom).




Figure 9: Right: Estimated covariance matrix of θ . Left: Diagonal components of the estimate of the average covariance matrix of the stochastic gradient as a function of the index $1 \leq i \leq d$.




Figure 10: Error on the covariance matrix of the parameters using AdL for ξ scalar, diagonal matrix or full matrix.

of the exact gradient using the N_{data} data points. The error on the covariance matrix is some measure of the bias on the posterior distribution effectively sampled by the numerical scheme. The results confirm that AdL with scalar ξ already leads to acceptable results in this example and performs as well as AdL with diagonal ξ .

3.5 Bayesian neural networks

In this section, we give another numerical example which demonstrates that, when the covariance matrix is not isotropic, there is a reduction of bias when passing from scalar to diagonal AdL. We consider to this end Bayesian neural networks (BNNs) used for binary classification problems. BNNs sample parameters of neural networks and therefore provide a distribution of likelihoods for points outside of the training set. The fact that a distribution




Figure 11: Data set for the classification problem tackled with BNNs. The solid lines are a guide to the eye to indicate possible decision boundaries. The cross represents the point $x = (1, 1.5)$ at which the parameters of the BNN, and hence its outputs, are sampled.

of likelihoods has a large variance is an indication that the classification of points outside the training set is uncertain. We consider here such a situation. We assess how the distributions of likelihoods is impacted by the minibatching error, and how this error can be reduced by diagonal versions of AdL.

Presentation of the model. We use a synthetic data set created by generating points using a mixture of Gaussians. Points in the class 0 (resp. 1) are generated by considering realizations of a random variable $X_0 \in \mathbb{R}^2$ (resp. X_1)

$$X_0 = x_{0,\mathcal{I}} + cG_0, \quad X_1 = x_{1,\mathcal{J}} + cG_1,$$

where G_0, G_1 are two independent 2-dimensional standard Gaussian random variables, while $\mathcal{I} \sim \mathcal{U}\{1, 2, 3\}$, $\mathcal{J} \sim \mathcal{U}\{1, 2, 3, 4\}$ and

$$\begin{cases} x_{0,1} = (1, 1), \\ x_{0,2} = (0, 2), \\ x_{0,3} = (1, 3). \end{cases} \quad \begin{cases} x_{1,1} = (2, 0.5), \\ x_{1,2} = (1, 2), \\ x_{1,3} = (2, 3), \\ x_{1,4} = (1, 4), \end{cases}$$

with $c^2 = 0.03$. The first class is therefore the superposition of 3 Gaussian distributions, while the second one is the superposition of 4 Gaussian distributions. The total number of data points is $N_{\text{data}} = 500$ (250 points in each class); see Figure 11. The probability that a point x has label 1 is estimated with a 2-layer neural network, with a hidden layer of size d_1 . More precisely,

$$N_\theta(x) = \sigma(W^2 r(W^1 x + b)),$$

where $W^1 \in \mathbb{R}^{d_1 \times 2}$, $b \in \mathbb{R}^{d_1}$ and $W^2 \in \mathbb{R}^{d_1 \times 1}$. The activation functions r is the ReLU function $r(z) = \max(0, z)$ which acts componentwise, while σ is the sigmoid function (69). The vector of parameters is therefore $\theta = (W^1, b, W^2) \in \mathbb{R}^d$ with $d = 4d_1$, where the matrices W^1 and W^2 are reshaped into vectors. For the numerical simulations, we set $d_1 = 8$.

The loss function is the binary cross entropy loss

$$L_f(x, y) = y \log(f(x)) + (1 - y) \log(1 - f(x)), \quad (70)$$

where $x \in \mathbb{R}^2$ is a point from the data set and $y \in \{0, 1\}$ its corresponding label. The likelihood distribution can be written in terms of the loss function as

$$P_{\text{elem}}(x, y|\theta) \propto \exp(-L_{N_\theta}(x, y)).$$

The prior distribution on the parameters θ is a standard d -dimensional Gaussian distribution.

Shape of the friction matrix in Adaptive Langevin dynamics. We restrict ourselves to diagonal friction matrices, and do not consider the full matricial case since it is computationally too expensive and cannot be considered as a practical sampling algorithm in the BNN framework. It is convenient to reshape ξ as a vector of dimension $4d_1$. As in (60), the first $2d_1$ entries of ξ correspond to the scalar coefficients associated with the sampling of the entries of W^1 , the following d_1 components to the components of b , and the final d_1 components to those of W^2 . This corresponds more formally to the choice

$$\xi = \left(\begin{array}{c|c|c} \xi_1 & 0 & 0 \\ \hline 0 & \xi_2 & 0 \\ \hline 0 & 0 & \xi_3 \end{array} \right),$$

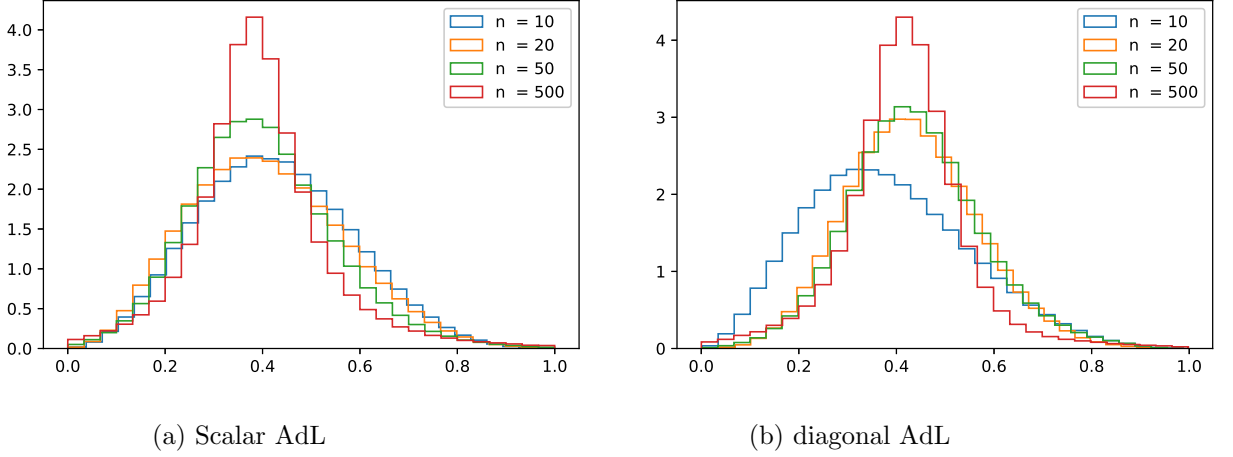

where ξ_1, ξ_2, ξ_3 are diagonal matrices of respective sizes $2d_1, d_1$ and d_1 . The number of coefficients can be reduced by considering a single scalar variable for W^1, b, W^2 , which corresponds to the choice

$$\xi = \left(\begin{array}{c|c|c} \xi_1 \mathbf{I}_{2d_1} & 0 & 0 \\ \hline 0 & \xi_2 \mathbf{I}_{d_1} & 0 \\ \hline 0 & 0 & \xi_3 \mathbf{I}_{d_1} \end{array} \right),$$

with $\xi_1, \xi_2, \xi_3 \in \mathbb{R}$.

Numerical results. We fix $\gamma = 1$ and $\eta = 200$, and use the numerical scheme (58) for AdL both in the diagonal and scalar cases to sample from the posterior distribution on θ . The numerical scheme is run for 4×10^5 epochs (each epoch being composed of N_{data}/n iterations of the numerical method), for various sizes of minibatch and various time steps. We plot in Figure 12 the distribution of $N_\theta(x)$ for various size of minibatch in the case of diagonal and scalar AdL; and in Figure 13 the L^1 error on the distribution of $N_\theta(x)$ for a given point $x = (1, 1.5)$ outside of the data set. The exact posterior distribution is considered to be the one corresponding to AdL without minibatching.

The results show that the diagonal version of AdL reduces the minibatching error compared to the scalar version of AdL for values of $\varepsilon(n)\Delta t$ sufficiently small. To confirm the

Figure 12: Distribution of $N_\theta(x)$ for various size of minibatch for $\Delta t = 10^{-2}$.Figure 13: L^1 error on the posterior distribution of $N_\theta(x)$ for various minibatching sizes and two values of the timestep Δt .

effect of the non isotropic matrix on the reduction of the bias, we approximate the mean and the variance of the diagonal element of the covariance matrix $\Sigma_{\mathbf{x}}$ using ξ for the diagonal AdL and compare them to the mean of ξ for the scalar case. We confirm that

$$\frac{1}{N_{\text{iter}}} \sum_{i=1}^{N_{\text{iter}}} \frac{1}{d} \sum_{j=1}^d [\xi^{\text{diagonal},i}]_{j,j} \approx \frac{1}{N_{\text{iter}}} \sum_{i=1}^{N_{\text{iter}}} \xi^{\text{scalar},i},$$

where ξ_{diagonal} and ξ_{scalar} are respectively the ξ variables for the diagonal and scalar AdL. The variance on the diagonal elements is about 2 (for each set of NN parameters), while the average of these elements is of the order of $\gamma = 1$.

4. Extended Adaptive Langevin Dynamics

AdL corrects for the bias due to mini-batching when the covariance of the stochastic gradient is constant. This is unfortunately not always the case as we demonstrated in Section 3.3. In this section, we suggest an extension of AdL which allows to remove the bias when the covariance is not constant but can be decomposed on a finite basis of functions. This approach in fact generally allows to reduce the minibatching bias even if the covariance matrix cannot be decomposed in a finite basis, as the residual minibatching bias, still of order $\varepsilon(n)\Delta t$, has a prefactor related to some projection error of the covariance matrix onto the approximation space under consideration, see (76) below.

We start by presenting the modified AdL dynamics in Section 4.1, under the assumption that the covariance of the force estimator can be decomposed on a finite basis of functions. This key assumption guarantees that the marginal distribution in the θ variable is the target posterior (2). We next propose a numerical scheme for extended AdL in Section 4.2, where we also quantify the bias on the invariant measure arising from the use of finite time steps, and possibly from mini-batching in situations where the covariance of the force estimator cannot be decomposed as a finite sum for the chosen basis of functions. We then discuss in Section 4.3 a crucial point of the method, namely the choice of basis functions. We finally present in Section 4.4 numerical results for the example of Section 2.5.2 where the elementary likelihoods are given by a mixture of Gaussians, and for which AdL fails to fully remove the bias because the covariance matrix is not constant. We demonstrate that with a reduced number function basis, the bias is significantly reduced even for the smallest values of n , including $n = 1$.

4.1 Presentation of the dynamics

Using the same notation as in Section 3, let us consider the case where $A_{\Delta t, n}(\theta)$ in (45) genuinely depends on θ in the following manner.

Assumption 16 *The matrix-valued function $A_{\Delta t, n}$ can be decomposed on a finite basis of functions f_0, \dots, f_K as*

$$A_{\Delta t, n}(\theta) = \sum_{k=0}^K A_{\Delta t, n}^k f_k(\theta), \quad (71)$$

where $A_{\Delta t, n}^0, \dots, A_{\Delta t, n}^K \in \mathbb{R}^{d \times d}$ are symmetric matrices. Moreover, $A_{\Delta t, n}(\theta)$ is a symmetric positive matrix for any $\theta \in \Theta$.

The choice of the basis of function is of great importance for the numerical performance of the method and is discussed more precisely in Section 4.3. The situation when (71) does not hold is considered at the end of Section 4.2. Note that the matrices $A_{\Delta t, n}^0, \dots, A_{\Delta t, n}^K$ need not be positive as long as $A_{\Delta t, n}(\theta)$ is. In accordance with (71), we choose the variable $\xi_t(\theta)$ to be of the following form:

$$\xi_t(\theta) = \sum_{k=0}^K \xi_{k, t} f_k(\theta), \quad \xi_{k, t} \in \mathbb{R}^{d \times d}. \quad (72)$$

If $K = 0$ and $f_0 = \mathbf{1}$, then (71) coincides with assumption (47), in which case AdL is sufficient to remove the bias due to mini-batching. In practice, the expression (71) can be thought of as a truncation of the expansion of the function $A_{\Delta t, n}(\theta)$ on a complete basis.

We are now in position to write the following extended Adaptive Langevin dynamics (eAdL), for which we introduce $K + 1$ additional (matrix) equations on the coefficients $\xi_{k, t}$ in (72):

$$\begin{cases} d\theta_t = p_t dt, \\ dp_t = \nabla_{\theta}(\log \pi(\theta_t | \mathbf{x})) dt - \xi_t(\theta_t) p_t dt + \sqrt{2} A_{\Delta t, n}(\theta_t)^{1/2} dW_t, \\ d[\xi_{k, t}]_{i, j} = \frac{f_k(\theta_t)}{\eta_k} (p_{i, t} p_{j, t} - \delta_{i, j}), \quad 1 \leq i, j \leq d, \quad 0 \leq k \leq K, \end{cases} \quad (73)$$

where $A_{\Delta t, n}$ is given by (45), $[\xi_{k, t}]_{i, j}$ is the (i, j) component of $\xi_{k, t} \in \mathbb{R}^{d \times d}$, and η_k are positive scalars for $0 \leq k \leq K$. The interest of eAdL is the following consistency result on the existence of an invariant probability measure with the correct marginal distribution in the θ variable.

Theorem 17 *Suppose that Assumption 16 holds. Then, the eAdL dynamics (73) admits the invariant probability measure*

$$\nu_K(d\theta dp d\xi_0 \dots d\xi_K) = \pi(\theta | \mathbf{x}) \tau(dp) \rho_K(d\xi) d\theta, \quad (74)$$

where $\tau(dp)$ is defined in (24), and

$$\rho_K(d\xi_0 \dots d\xi_K) = \prod_{k=0}^K \prod_{i, j=1}^d \sqrt{\frac{\eta_k}{2\pi}} \exp\left(-\frac{\eta_k}{2} \left([\xi_k]_{i, j} - [A_{\Delta t, n}^k]_{i, j}\right)^2\right) d[\xi_k]_{i, j},$$

with $[A_{\Delta t, n}^k]_{i, j}$ the (i, j) component of $A_{\Delta t, n}^k \in \mathbb{R}^{d \times d}$.

As for AdL (see the discussion following Lemma 9), Theorem 17 suggests that we recover the target posterior distribution $\pi(\cdot | \mathbf{x})$ whatever the extra noise due to mini-batching, since the marginal in the variable θ of the probability measure ν_K in (74) is $\pi(\cdot | \mathbf{x})$. However, from a discussion similar to the one in Remark 10, the dynamics cannot be ergodic for the extended measure ν_K .

Proof We follow the same approach as for the proof of Lemma 9. The generator of the dynamics reads

$$\mathcal{L}_{\text{eAdL}, \Sigma_{\mathbf{x}}} = \mathcal{L}_{\text{ham}} + \mathcal{L}_{\text{FD}} + \tilde{\mathcal{L}}_{\text{NH}},$$

where \mathcal{L}_{ham} and \mathcal{L}_{FD} are the same operators as in (50) and (51) respectively (even if $A_{\Delta t, n}$ depends on θ), while

$$\tilde{\mathcal{L}}_{\text{NH}} = \sum_{k=0}^K f_k \mathcal{L}_{\text{NH}, k},$$

where

$$\mathcal{L}_{\text{NH}, k} = - \sum_{i, j=1}^d p_i \left([\xi_k]_{i, j} - [A_{\Delta t, n}^k]_{i, j}\right) \partial_{p_j}^* + \frac{1}{\eta_k} (p_i p_j - \delta_{i, j}) \partial_{[\xi_k]_{i, j}}.$$

A computation similar to the one performed in the proof of Lemma 9 shows that $\mathcal{L}_{\text{NH},k}^* = -\mathcal{L}_{\text{NH},k}$, with adjoints taken with respect to $L^2(\nu_K)$; while \mathcal{L}_{ham} and \mathcal{L}_{FD} are respectively antisymmetric and symmetric on $L^2(\nu_K)$. This implies that $\mathcal{L}_{\text{eAdL},\Sigma_x}^* \mathbf{1} = 0$, from which the claimed invariance of ν_K follows. \blacksquare

4.2 Numerical scheme and estimates on the bias

We present in this section a numerical integrator for the eAdL dynamics (73) based on a Strang splitting similar to the one considered in Section 3.2.2 for AdL. We consider the same elementary SDEs (53)-(56) as for the discretization of AdL, except that there are now $K + 1$ elementary SDEs in (56), indexed by $0 \leq k \leq K$. The associated numerical scheme obtained for $\Gamma = \gamma \text{I}_d$ reads

$$\left\{ \begin{array}{l} \xi^m = \sum_{k=0}^K \xi_k^m f_k(\theta^m), \\ p^{m+\frac{1}{2}} = e^{-\Delta t \xi^m / 2} p^m + \left[\gamma (\xi^m)^{-1} \left(\text{I}_d - e^{-\Delta t \xi^m} \right) \right]^{1/2} G^m, \\ \xi_k^{m+\frac{1}{2}} = \xi_k^m + \frac{\Delta t}{2\eta_k} f_k(\theta^m) \left[\left(p^{m+\frac{1}{2}} \right) \left(p^{m+\frac{1}{2}} \right)^T - \text{I}_d \right], \quad k = 0, \dots, K, \\ \theta^{m+\frac{1}{2}} = \theta^m + \frac{\Delta t}{2} p^{m+\frac{1}{2}}, \\ \tilde{p}^{m+\frac{1}{2}} = p^{m+\frac{1}{2}} + \Delta t \widehat{F}_n \left(\theta^{m+\frac{1}{2}} \right), \\ \theta^{m+1} = \theta^{m+\frac{1}{2}} + \frac{\Delta t}{2} \tilde{p}^{m+\frac{1}{2}}, \\ \xi_k^{m+1} = \xi_k^{m+\frac{1}{2}} + \frac{\Delta t}{2\eta_k} f_k(\theta^{m+1}) \left[\left(\tilde{p}^{m+\frac{1}{2}} \right) \left(\tilde{p}^{m+\frac{1}{2}} \right)^T - \text{I}_d \right], \quad k = 0, \dots, K, \\ \xi^{m+1} = \sum_{k=0}^K \xi_k^{m+1} f_k(\theta^{m+1}), \\ p^{m+1} = e^{-\Delta t \xi^{m+1} / 2} \tilde{p}^{m+\frac{1}{2}} + \left[\gamma (\xi^{m+1})^{-1} \left(\text{I}_d - e^{-\Delta t \xi^{m+1}} \right) \right]^{1/2} G^{m+\frac{1}{2}}, \end{array} \right. \quad (75)$$

where $(G^m)_{m \geq 0}$ and $(G^{m+\frac{1}{2}})_{m \geq 0}$ are two independent families of i.i.d. standard d -dimensional Gaussian random variables.

Characterization of the bias. We assume as before that the numerical scheme (75) admits a unique invariant probability measure, which may depend on the initial condition $(\xi_0^0, \dots, \xi_K^0)$. Error estimates on this invariant probability measure can be obtained by following the approach of in Section 3.3.1. Given the high level of similarity between the estimates in Section 3.3.1 and the ones obtained here, we only make precise what needs to be changed.

The main conclusion is that the bias is still of order $\Delta t^2 + \varepsilon(n)\Delta t$, but with a smaller prefactor for the dominant term $\varepsilon(n)\Delta t$. Let us emphasize that we do not require Assumption 16 to hold for this analysis. To make this precise, the first step is to write the generator

as

$$\mathcal{L}_{\text{eAdL}, \Sigma_{\mathbf{x}}} = \mathcal{L}_{\text{eAdL}, \bar{\Sigma}_{\mathbf{x}}^K} + \varepsilon(n) \Delta t \tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}^K},$$

where $\tilde{\mathcal{L}}_{\mathcal{M}}$ is defined in (64), and $\bar{\Sigma}_{\mathbf{x}}^K$ is a symmetric positive matrix which belongs to the vector space generated by f_0, \dots, f_K . The specific choice of this matrix for the theoretical analysis relies again on an orthogonal projection and the associated projection error, see (76) below. One can then state an inequality similar to (65), where $\bar{\nu}_{\xi^0}$ is replaced by the invariant probability measure $\bar{\nu}_{K, (\xi_0^0, \dots, \xi_K^0)}$ of the eAdL dynamics associated with the generator $\mathcal{L}_{\text{eAdL}, \bar{\Sigma}_{\mathbf{x}}^K}$, and ν_{ξ^0} is replaced by the invariant probability measure $\nu_{K, (\xi_0^0, \dots, \xi_K^0)}$ of the eAdL dynamics with generator $\mathcal{L}_{\text{eAdL}, \Sigma_{\mathbf{x}}}$. The prefactor of the dominant error term, proportional to $\varepsilon(n) \Delta t$, depends on the function

$$f_{\text{eAdL}, \Sigma_{\mathbf{x}}, \bar{\Sigma}_{\mathbf{x}}^K} = \left(-\mathcal{L}_{\text{eAdL}, \bar{\Sigma}_{\mathbf{x}}^K}^* \right)^{-1} \tilde{\mathcal{L}}_{\Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}^K}^* \mathbf{1},$$

where adjoints are taken on $L^2(\bar{\nu}_{K, (\xi_0^0, \dots, \xi_K^0)})$. Under appropriate conditions on the structure of $\bar{\nu}_{K, (\xi_0^0, \dots, \xi_K^0)}$ and resolvent bounds for $\mathcal{L}_{\text{eAdL}, \bar{\Sigma}_{\mathbf{x}}^K}$, it is possible to upper bound the norm of $f_{\text{eAdL}, \Sigma_{\mathbf{x}}}$ in $L^2(\bar{\nu}_{K, (\xi_0^0, \dots, \xi_K^0)})$ by $\left\| \Sigma_{\mathbf{x}} - \bar{\Sigma}_{\mathbf{x}}^K \right\|_{L^2(\pi)}$. By optimizing upon the matrix valued function $\bar{\Sigma}_{\mathbf{x}}^K$, one can conclude that the prefactor of the error term proportional to $\varepsilon(n) \Delta t$ is bounded, up to a constant, by

$$\min_{\mathcal{M}_0, \dots, \mathcal{M}_K \in \mathbb{R}^{d \times d}} \left\| \Sigma_{\mathbf{x}} - \sum_{k=0}^K \mathcal{M}_k f_k \right\|_{L^2(\pi)}, \quad (76)$$

which corresponds to the $L^2(\pi)$ projection error of $\Sigma_{\mathbf{x}}$ onto the vector space of symmetric matrices generated by the basis. This shows that the better the approximation of the covariance $\Sigma_{\mathbf{x}}$ is for the chosen basis (f_0, \dots, f_K) , the smaller the bias is.

Discussion on the computation cost. As for the numerical discretization of AdL (58) (see Remark 11), we can also use the numerical scheme (75) while considering each variable ξ_k as a scalar or a diagonal matrix. Observations similar to the ones made in Remark 15 apply here as well. More precisely, the prefactor of the error term proportional to $\varepsilon(n) \Delta t$ is bounded by the $L^2(\pi)$ -projection of $\Sigma_{\mathbf{x}}$ onto: (i) the vector space generated by functions with values in the vector space of isotropic matrices, *i.e.* the variables ξ_k are scalars for each k , so that $\mathcal{M}_k = \sigma_k \mathbf{I}_d$ where $\sigma_k \in \mathbb{R}$ in (76); (ii) the vector space of diagonal matrices generated by the basis when the variables ξ_k are diagonal matrices for each k , meaning that $\mathcal{M}_0, \dots, \mathcal{M}_K$ are diagonal matrices in (76). For the numerical simulations reported in Section 4.4, we consider the variables ξ_k as full $d \times d$ symmetric matrices.

The extra computational cost associated with extended Adaptive Langevin dynamics depends on the choice made for the form of the friction matrix. When ξ is a scalar, the extra computational cost compared to standard Langevin like dynamics is of order K , while it is of order Kd when ξ is a diagonal matrix. When ξ is a matrix, the computational cost is much higher because of the matrix exponential which appears in the first and last steps of the numerical scheme (75). The computation of such exponentials involves in general a matrix diagonalization, which has a cost of order d^3 . Additionally, one needs to update

the components of K matrices of size $d \times d$, which has a cost of order Kd^2 . Overall, the extra computational cost of extended Adaptive Langevin is therefore of order $(K + d)d^2$ in the matrix case, to be compared to the cost of order d^3 of usual Adaptive Langevin with constant friction matrices. This cost could however be mitigated to CKd^2 upon using a midpoint scheme to integrate the Ornstein–Uhlenbeck process, and solving the resulting matrix problem using iterative methods. In any case, when the dimension d of the parameters is large, it is usually too expensive from a computational viewpoint to use matrix versions of Adaptive Langevin. This may also be unnecessary, as all the d^2 entries of the covariance matrix are not necessarily relevant; recall the example of Section 3.4 and see the discussion in Section 5.

4.3 Choice the basis functions

The choice of the basis functions f_0, \dots, f_K is a key point in eAdL, since there is a trade-off between a good approximation of the matrix $A_{\Delta t, n}(\theta)$ such that the projection error (76) is small, and the number of additional degrees of freedom (which is equal to the number of functions K introduced multiplied by $d(d+1)/2$ if the unknown are matrices). One typically prefers to introduce only a limited number of degrees of freedom. One option towards this goal is for instance to consider only isotropic matrices $\xi_k = \sigma_k \mathbf{I}_d$ with $\sigma_k \in \mathbb{R}$. Such a choice is advocated in Ding et al. (2014) for $K = 0$ and $f_0 = \mathbf{1}$.

Spatial decomposition. Numerical experiments suggest that the covariance matrix of the estimator of the gradient may change rapidly in certain regions of the parameter space (see in particular Figure 6). A convenient approach to approximating this matrix is to partition the domain Θ into $K + 1$ subdomains, denoted by $\mathcal{D}_0, \dots, \mathcal{D}_K$, and to consider the functions f_k to be indicator functions of these domains, *i.e.* $f_k = \mathbf{1}_{\mathcal{D}_k}$. This corresponds to a piecewise constant approximation of the covariance matrix. If the domain is simple and the dimension d is sufficiently small, one can think of simple geometric decompositions, using *e.g.* rectangles if Θ is itself rectangular, or rings around the various modes of the posterior distribution. For more complicated domains in possibly high dimension, one can decompose the parameter space with a Voronoi tessellation, where the centers of the Voronoi tessellation are for instance the local maxima of the posterior distribution. Such points can be located by performing preliminary SGLD or AdL runs. This amounts to considering a constant friction matrix in each mode of the probability distribution.

Polynomial approximation. If the domain Θ is sufficiently simple from a geometric viewpoint, a more sophisticated method is to couple a spatial decomposition with some spectral approximation, by introducing basis functions on each subdomain. For instance, for a rectangular decomposition of the domain, a polynomial basis can be defined on each subdomain

$$\mathcal{D}_k = [M_{k,1}^-, M_{k,1}^+] \times \dots \times [M_{k,d}^-, M_{k,d}^+], \quad (77)$$

by tensorization of monomial functions of the form

$$e_{k,j,i}(\theta_j) = \left(\frac{\theta_j - M_{k,j}^-}{M_{k,j}^+ - M_{k,j}^-} \right)^i. \quad (78)$$

The integer k indexes the domain under consideration, j characterizes the degree of freedom under consideration, and i is the power of the monomial. The normalization we choose ensures that each of the term in the tensor product defining the basis function has values in $[0, 1]$. Defining for instance the degree of the tensor product \mathbf{deg} to be the maximum of the degree of the individual polynomials, the total number of degrees of freedom is given by $K(\mathbf{deg} + 1)^d$. When \mathbf{deg} is large, the elementary polynomials (78) assume non negligible values only close to the boundary of the domain, which can result in numerical instabilities. We therefore rescale the basis functions so that their $L^2(\pi)$ norm is 1 (as the normalization factor being estimated using preliminary AdL runs).

4.4 Numerical illustration on a mixture of Gaussians

We present in this section numerical results for the example introduced in Section 2.5.2. The covariance matrix in this case is not constant, and AdL fails to correct for the extra noise due to mini-batching; see Section 3.3.2 and Figure 7. We consider the same parameters as in Section 3.3.2. To define the basis functions, we first consider a symmetric rectangular partition on the domain, with 4 meshes of the form (77), with $M_{0,1}^- = M_{0,2}^- = M_{1,2}^- = M_{2,1}^- = 0$, $M_{0,1}^+ = M_{0,2}^+ = M_{1,1}^- = M_{1,2}^+ = M_{2,1}^+ = M_{2,2}^- = M_{3,1}^- = M_{3,2}^- = 0.7$ and $M_{1,1}^+ = M_{2,2}^+ = M_{3,1}^+ = M_{3,2}^+ = 1.4$. We define polynomials on each mesh as discussed in Section 4.3. We run the numerical scheme (75) for eAdL with $\eta_k = 1$ for all values of $0 \leq k \leq K$, for a final time $T = 10^6$, various values of Δt , and various degrees of polynomials. The case $K = 0$ corresponds to standard AdL, while $K = 3$ corresponds to an estimation of the covariance matrix by a piecewise constant matrix valued function on the 4 domains under consideration, $K = 15$ to products of affine functions in each degree of freedom on the 4 domains, and $K = 35$ to products of second order polynomials in each variable on each domain. We fix the size of the minibatch n to 5 and 15, which corresponds respectively to $\varepsilon(n) = 7800$ and $\varepsilon(n) = 2466.67$. We report in Figure 14 the L^1 error on the marginal distribution over θ_1 of the posterior distribution sampled by the numerical scheme for various values of K . Let us first emphasize that the bias decreases compared to AdL even for the smallest value $K = 3$ (compare with Figure 7). A dramatic decrease is observed for $K = 15$. A small bias remains for $K \geq 15$, probably mostly due to the time step error. We plot in Figure 15 the quantity $\|\Sigma_{\mathbf{x}} - S_K\|_{L^2(\pi)}$ where S_K is the $L^2(\pi)$ projection of $\Sigma_{\mathbf{x}}$ onto the vector space of symmetric matrices generated by the basis for each value of K . The key observation here is that the projection error, which quantifies the quality of the approximation of $\Sigma_{\mathbf{x}}$, decays similarly to the bias on the invariant measure when K is increased. This therefore confirms the bias analysis of Section 4.2, which predicts that the bias should be controlled by the projection error.

5. Discussion and perspectives

We quantified in this work the bias on the posterior distribution arising from mini-batching for (kinetic or underdamped) Langevin dynamics, both for standard and adaptive Langevin dynamics. As highlighted in (1), the bias is dictated by the quality of the approximation of the covariance matrix $\Sigma_{\mathbf{x}}(\theta)$ of the gradient estimator by the average (at θ fixed) of the extra friction variable ξ introduced in AdL. The latter average is independent of the pa-

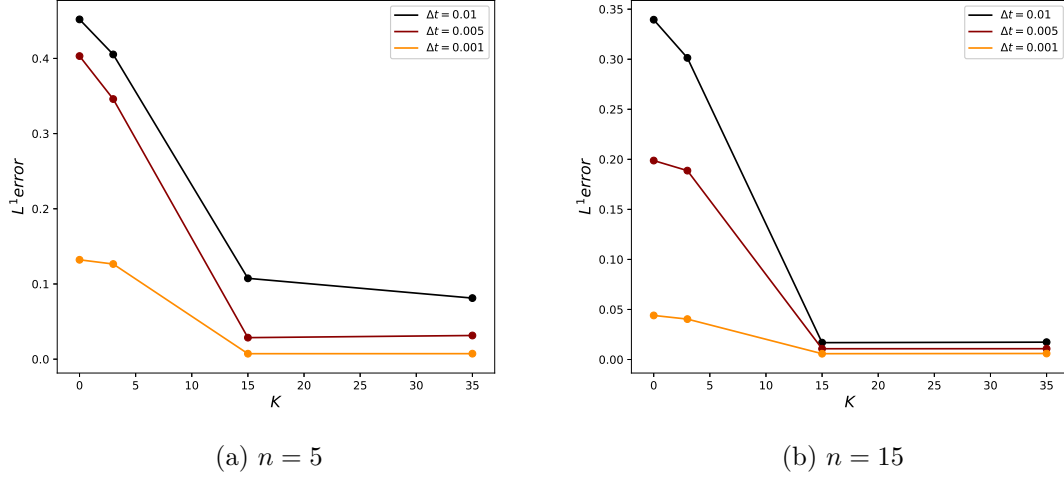


Figure 14: L^1 error on the the marginal over θ_1 of the posterior distribution with eAdL for various values of K and Δt ($K = 0$ corresponds to AdL results).




Figure 15: Projection error $\|\Sigma_{\mathbf{x}} - S_K\|_{L^2(\pi)}$, where S_K is the $L^2(\pi)$ projection of $\Sigma_{\mathbf{x}}$ onto the vector space of symmetric matrices generated by the basis for each value of K .

parameter θ in the original formulation of AdL, and this is sufficient to remove the bias if the target distribution is Gaussian, and allows to strongly reduce the bias for posteriors close to Gaussian (as for the MNIST example considered in Section 3.4). The latter situation is typical when the number of data points N_{data} is large, as the posterior distribution concentrates in a Gaussian manner around one mode (up to symmetries) due to the Bernstein–von Mises theorem (see Section 10.2 in van der Vaart (1998)). In this case, it is likely that AdL already captures most of the mini-batching bias. Our analysis however allows to go one step further, and make precise when a reduction to diagonal, or even isotropic friction matrices, does not degrade the quality of the sampling.

Standard AdL is however not sufficient in various situations, in particular when $\Sigma_{\mathbf{x}}$ genuinely depends on θ , which is the case when N_{data} is of moderate size. In this case, the posterior distribution can strongly deviate from a Gaussian distribution. Extended AdL provides a framework to systematically reduce the mini-batching bias in such situations, by adding extra degrees of freedom to better tune the friction variable. Striking reductions in the bias can be observed already for not too many additional degrees of freedom, even for the extreme situation where the batch size is set to its minimal value $n = 1$. In any case, eAdL always leads to a smaller bias than AdL.

Let us also emphasize here that AdL and eAdL can be used in conjunction with other techniques to reduce the variance of the stochastic estimator of the force, such as control variate or extrapolation techniques (see for instance Nagapetyan et al. (2017); Brosse et al. (2018); Baker et al. (2019); Durmus et al. (2016)). AdL and eAdL should therefore not be seen as alternatives to these techniques, but as complements.

The choice of the dimensionality of the vector space for the friction variable is a key element for standard AdL and its extension, as this directly impacts the computational cost of the method; see the discussion at the end of Section 4.2. One would like the number of degrees of freedom to be as small as possible. From a theoretical perspective, this motivates further characterizing the structure of the covariance matrix $\Sigma_{\mathbf{x}}$, which have been observed to be low rank in certain situations related to neural network training; see Chaudhari and Soatto (2018). The low rank structure suggests that most of the variance $\Sigma_{\mathbf{x}}(\theta)$ can be captured by relying on frictions $\xi(\theta)$ parametrized by a few degrees of freedom only, potentially much smaller than d . Better characterizing the structure of $\Sigma_{\mathbf{x}}$ could also shed some light on the current active research efforts on understanding the so-called implicit bias of neural network training.

Acknowledgments

The authors thank Ben Leimkuhler and Tiffany Vlaar for stimulating discussions on AdL and Matthias Sachs for kindly providing the code used to perform numerical experiments of Section 3.4. I. Sekkat gratefully acknowledges financial support from Université Mohammed VI Polytechnique. The work of G. Stoltz is funded in part by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 810367), and by the Agence Nationale de la Recherche, under grants ANR-19-CE40-0010-01 (QuAMProcs) and ANR-21-CE40-0006 (SINEQ).

Appendix A. Consistency estimates

We provide in this appendix the algebraic manipulations for various consistency estimates. To simplify the notation, we simply write Z in all this section for the random variable which appears in (14). The analysis we perform is asymptotic since it relies on the assumption that $\varepsilon(n) \geq 1$ and $\Delta t \leq 1$ are such that $\varepsilon(n)\Delta t \leq 1$ is sufficiently small. Let us also recall that the notation $\mathcal{O}(\Delta t^{s+1})$ in equalities such as (6) is made precise in (7). We explicitly write the first estimates in Section A.1 using the full notation from (7) to give a clear meaning to all equalities, but then return to the notation $\mathcal{O}(\Delta t^{s+1})$ for the sake of readability, hoping that the first estimates and (7) allow the reader to see how to give a precise meaning to all the equalities we write.

A.1 Proof of (19)

We prove here that the evolution operator for SGLD satisfies (19) for Δt and $\varepsilon(n)\Delta t$ small. The equality (19) should be understood as follows: for any given smooth function $\theta \mapsto \varphi(\theta)$ with compact support, there exist $\Delta t_\star > 0$ and $K \in \mathbb{R}_+$ such that, for any $\Delta t \in (0, \Delta t_\star]$, there is a function $R_{\varphi,n,\Delta t}$ for which

$$\left(\widehat{P}_{\Delta t,n}\varphi\right)(\theta) = \left(e^{\Delta t(\mathcal{L}_{\text{ovd}} + \Delta t \mathcal{A}_{\text{ovd},n})}\varphi\right)(\theta) + (1 + \varepsilon(n)^{3/2})\Delta t^3 R_{\varphi,n,\Delta t}(\theta),$$

with

$$\sup_{\Delta t \in (0, \Delta t_\star]} \sup_{1 \leq n \leq N_{\text{data}}} \sup_{\theta \in \Theta} |R_{\varphi,n,\Delta t}(\theta)| \leq K.$$

We rewrite to this end the SGLD scheme (18) as

$$\theta^{m+1} = \Phi_{\Delta t,n}(\theta^m, Z^m, G^m),$$

with

$$\Phi_{\Delta t,n}(\theta, Z, G) = \theta + \sqrt{\varepsilon(n)}\Delta t \Sigma_{\mathbf{x}}^{1/2}(\theta)Z + \sqrt{2\Delta t}G + \Delta t \nabla_{\theta}(\log \pi(\theta|\mathbf{x})).$$

Consider a given smooth function $\theta \mapsto \varphi(\theta)$ with compact support. Then, there exists $C \in \mathbb{R}_+$ such that

$$\begin{aligned} \varphi(\Phi_{\Delta t,n}(\theta^m, Z^m, G^m)) &= \varphi\left(\Phi_{\Delta t,n}(\theta^m, 0, G^m) + \sqrt{\varepsilon(n)}\Delta t \Sigma_{\mathbf{x}}^{1/2}(\theta^m)Z^m\right) \\ &= \varphi(\Phi_{\Delta t,n}(\theta^m, 0, G^m)) + \sqrt{\varepsilon(n)}\Delta t \Sigma_{\mathbf{x}}^{1/2}(\theta^m)Z^m \cdot (\nabla_{\theta}\varphi)(\Phi_{\Delta t,n}(\theta^m, 0, G^m)) \\ &\quad + \frac{1}{2}\varepsilon(n)\Delta t^2(\nabla_{\theta}^2\varphi)(\Phi_{\Delta t,n}(\theta^m, 0, G^m)) : \Sigma_{\mathbf{x}}^{1/2}(\theta^m)Z^m \otimes \Sigma_{\mathbf{x}}^{1/2}(\theta^m)Z^m \\ &\quad + \frac{1}{6}\varepsilon(n)^{3/2}\Delta t^3 D^3\varphi(\Phi_{\Delta t,n}(\theta^m, 0, G^m)) : \left(\Sigma_{\mathbf{x}}^{1/2}(\theta^m)Z^m\right)^{\otimes 3} \\ &\quad + \varepsilon(n)^2\Delta t^4 \mathcal{R}_{\varphi,n,\Delta t}(\theta^m, G^m, Z^m), \end{aligned}$$

with

$$\sup_{\Delta t \in (0, \Delta t_\star]} \sup_{1 \leq n \leq N_{\text{data}}} \sup_{\theta \in \Theta} |\mathcal{R}_{\varphi,n,\Delta t}(\theta, G, Z)| \leq C(1 + |Z|^4).$$

In the above expansion in powers of Δt , we use the notation

$$D^3\varphi(\theta) : v_1 \otimes v_2 \otimes v_3 = \sum_{i,j,k=1}^d [v_1]_i [v_2]_j [v_3]_k (\partial_{\theta_i, \theta_j, \theta_k}^3 \varphi)(\theta),$$

and $v^{\otimes 3} = v \otimes v \otimes v$. Using classical results on the evolution operator of the Euler–Maruyama scheme to compute the expectation of $\varphi(\Phi_{\Delta t, n}(\theta^m, 0, G^m))$ with respect to G^m (see for example Vollmer et al. (2016) for details), we can compute the expectation of $\varphi(\Phi_{\Delta t, n}(\theta^m, Z^m, G^m))$ with respect to the variables G^m and Z^m as

$$\begin{aligned} \left(\widehat{P}_{\Delta t, n} \varphi \right) (\theta^m) &= \text{I}_d + \Delta t \mathcal{L}_{\text{ovd}} \phi(\theta^m) + \Delta t^2 \left(\frac{1}{2} \mathcal{L}_{\text{ovd}}^2 + \mathcal{A}_{\text{disc}} \right) \varphi(\theta^m) \\ &\quad + \Delta t^2 \varepsilon(n) \mathbb{E}_G [(\mathcal{A}_{\text{mb}} \varphi)(\Phi_{\Delta t, n}(\theta^m, 0, G))] + \left(1 + \varepsilon(n)^{3/2} \right) \Delta t^3 \mathcal{R}_{\varphi, n, \Delta t}(\theta^m), \end{aligned}$$

with

$$\sup_{\Delta t \in (0, \Delta t_*]} \sup_{1 \leq n \leq N_{\text{data}}} \sup_{\theta \in \Theta} |\mathcal{R}_{\varphi, n, \Delta t}(\theta)| < +\infty.$$

For the end of the proof, for the sake of readability, we no longer make explicit that the operator equalities should be understood as applied to a smooth and compactly supported functions φ , with remainder terms multiplying functions of θ which are uniformly bounded with respect to $1 \leq n \leq N_{\text{data}}$, $\Delta t \leq \Delta t_*$ and $\theta \in \Theta$. Since $\mathbb{E}_G [\mathcal{A}_{\text{mb}} \varphi(\Phi_{\Delta t, n}(\theta^m, 0, G))] = \mathcal{A}_{\text{mb}} \varphi(\theta^m) + \mathcal{O}(\Delta t)$, we obtain on the one hand that

$$\widehat{P}_{\Delta t, n} = \text{I}_d + \Delta t \mathcal{L}_{\text{ovd}} + \Delta t^2 \left(\frac{1}{2} \mathcal{L}_{\text{ovd}}^2 + \mathcal{A}_{\text{disc}} + \varepsilon(n) \mathcal{A}_{\text{mb}} \right) + \mathcal{O} \left((1 + \varepsilon(n)^{3/2}) \Delta t^3 \right). \quad (79)$$

On the other hand, notice that

$$\begin{aligned} e^{\Delta t (\mathcal{L}_{\text{ovd}} + \varepsilon(n) \Delta t \mathcal{A}_{\text{mb}} + \Delta t \mathcal{A}_{\text{disc}})} &= \text{I}_d + \Delta t \mathcal{L}_{\text{ovd}} + \Delta t^2 \left(\frac{1}{2} \mathcal{L}_{\text{ovd}}^2 + \mathcal{A}_{\text{disc}} + \varepsilon(n) \mathcal{A}_{\text{mb}} \right) \\ &\quad + \mathcal{O}((1 + \varepsilon(n)) \Delta t^3). \end{aligned} \quad (80)$$

Using (79) alongside with (80), we deduce the desired result (19).

Remark 18 *Note that when $\mathbb{E}[Z^3] = 0$, the remainder term is in fact of order $\mathcal{O}((1 + \varepsilon(n)^2 \Delta t) \Delta t^3)$ in (79). This leads to a remainder of order $\mathcal{O}((1 + \varepsilon(n)) \Delta t^3)$ in (19).*

A.2 Proof of (31), (32) and (34)

To prove (31), we consider a smooth function $(\theta, p) \mapsto \varphi(\theta, p)$ with compact support, and rewrite the operator as

$$\begin{aligned} \left(Q_{\Delta t}^{\mathcal{L}_1} \varphi \right) (\theta^m, p^m) &= \mathbb{E} \left[\varphi \left(\theta^m, p^m + \Delta t \widehat{F}_n(\theta^m) \right) \right] \\ &= \mathbb{E} \left[\varphi \left(\theta^m, p^m + \Delta t \nabla_{\theta} (\log \pi(\theta^m | \mathbf{x})) + \sqrt{\varepsilon(n)} \Delta t \Sigma_{\mathbf{x}}^{1/2}(\theta^m) Z^m \right) \right]. \end{aligned}$$

We then have

$$\begin{aligned}
 \left(Q_{\Delta t}^{\mathcal{L}_1}\varphi\right)(\theta^m, p^m) &= \mathbb{E}[\varphi(\theta^m, p^m + \Delta t \nabla_{\theta}(\log \pi(\theta^m | \mathbf{x})))] \\
 &+ \sqrt{\varepsilon(n)} \Delta t \mathbb{E}\left[\Sigma_{\mathbf{x}}^{1/2}(\theta^m) Z^m \cdot \nabla_p \varphi(\theta^m, p^m + \nabla_{\theta}(\log \pi(\theta^m | \mathbf{x}))\right] \\
 &+ \frac{1}{2} \varepsilon(n) \Delta t^2 \mathbb{E}\left[\nabla_p^2 \varphi(\theta^m, p^m + \Delta t \nabla_{\theta}(\log \pi(\theta^m | \mathbf{x})) : \Sigma_{\mathbf{x}}^{1/2}(\theta^m) Z^m \otimes \Sigma_{\mathbf{x}}^{1/2}(\theta^m) Z^m\right] \\
 &+ \mathcal{O}\left(\varepsilon(n)^{3/2} \Delta t^3\right),
 \end{aligned}$$

so that

$$\begin{aligned}
 \left(Q_{\Delta t}^{\mathcal{L}_1}\varphi\right)(\theta^m, p^m) &= (e^{\Delta t \mathcal{L}_1} \varphi)(\theta^m, p^m) + \varepsilon(n) \Delta t^2 (\mathcal{A}_{\text{lan}} \varphi)(\theta^m, p^m + \Delta t \nabla_{\theta}(\log \pi(\theta^m | \mathbf{x}))) \\
 &+ \mathcal{O}\left(\varepsilon(n)^{3/2} \Delta t^3\right) \\
 &= (e^{\Delta t \mathcal{L}_1} \varphi)(\theta^m, p^m) + \varepsilon(n) \Delta t^2 (\mathcal{A}_{\text{lan}} \varphi)(\theta^m, p^m) + \mathcal{O}\left(\varepsilon(n)^{3/2} \Delta t^3\right),
 \end{aligned} \tag{81}$$

from which the result directly follows.

Remark 19 *If $\mathbb{E}[Z^3] = 0$, the error term $\mathcal{O}\left(\varepsilon(n)^{3/2} \Delta t^3\right)$ can be replaced by $\mathcal{O}\left(\varepsilon(n) \Delta t^3\right)$.*

To prove (32), we use the Baker–Campbell–Hausdorff formula on the operator of the numerical scheme (30):

$$\begin{aligned}
 \widehat{P}_{\Delta t, n} &= e^{\Delta t \mathcal{L}_3/2} e^{\Delta t \mathcal{L}_2/2} Q_{\Delta t}^{\mathcal{L}_1} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_3/2} \\
 &= e^{\Delta t \mathcal{L}_3/2} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_1} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_3/2} + \varepsilon(n) \Delta t^2 e^{\Delta t \mathcal{L}_3/2} e^{\Delta t \mathcal{L}_2/2} \mathcal{A}_{\text{lan}} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_3/2} \\
 &+ \mathcal{O}\left(\varepsilon(n)^{3/2} \Delta t^3\right) \\
 &= e^{\Delta t \mathcal{L}_{\text{lan}}} + \mathcal{O}\left((\Delta t + \varepsilon(n)) \Delta t^2\right).
 \end{aligned}$$

Using computations similar to the ones leading to (19), we can prove the following estimate on $Q_{\Delta t}^{\mathcal{L}_1}$:

$$e^{\Delta t (\mathcal{L}_1 + \varepsilon(n) \Delta t \mathcal{A}_{\text{lan}})} = Q_{\Delta t}^{\mathcal{L}_1} + \mathcal{O}\left((1 + \varepsilon(n)^{3/2}) \Delta t^3\right). \tag{82}$$

The desired result (34) follows by using the BCH formula.

A.3 Proof of (61)

The evolution operator of the numerical scheme of AdL is given by

$$\widehat{P}_{\Delta t, n} = e^{\Delta t \mathcal{L}_4/2} e^{\Delta t \mathcal{L}_3/2} e^{\Delta t \mathcal{L}_2/2} Q_{\Delta t}^{\mathcal{L}_1} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_3/2} e^{\Delta t \mathcal{L}_4/2},$$

where $\mathcal{L}_1 = \nabla_{\theta} \log(\pi(\cdot | \mathbf{x}))^T \nabla_p$ and \mathcal{L}_2 are respectively the generators of (55) when $\varepsilon(n) = 0$ and (53) (which coincide with the operators in (26)), while \mathcal{L}_3 and \mathcal{L}_4 are the generators of (56) and (54). Replacing $Q_{\Delta t}^{\mathcal{L}_1}$ by its expression in (81) and using the BCH formula, we

obtain, by computations similar to the ones leading to (82):

$$\begin{aligned}
\widehat{P}_{\Delta t, n} \varphi &= e^{\Delta t \mathcal{L}_4/2} e^{\Delta t \mathcal{L}_3/2} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_1} e^{\Delta t \mathcal{L}_2/2} e^{\Delta t \mathcal{L}_3/2} e^{\Delta t \mathcal{L}_3/2} + \varepsilon(n) \Delta t^2 \mathcal{A}_{\text{lan}} \varphi + \mathcal{O}\left(\varepsilon(n)^{3/2} \Delta t^3\right) \\
&= e^{\Delta t (\mathcal{L}_1 + \mathcal{L}_2 + \mathcal{L}_3 + \mathcal{L}_4)} \varphi + \varepsilon(n) \Delta t^2 \mathcal{A}_{\text{lan}} \varphi + \mathcal{O}\left((1 + \varepsilon(n)^{3/2}) \Delta t^3\right) \\
&= e^{\Delta t (\mathcal{L}_1 + \mathcal{L}_2 + \mathcal{L}_3 + \mathcal{L}_4 + \varepsilon(n) \Delta t \mathcal{A}_{\text{lan}})} \varphi + \mathcal{O}\left(\varepsilon(n) \Delta t^3\right) + \mathcal{O}\left((1 + \varepsilon(n)^{3/2}) \Delta t^3\right) \\
&= e^{\Delta t \mathcal{L}_{\text{AdL}, \Sigma_x}} \varphi + \mathcal{O}\left((1 + \varepsilon(n)^{3/2}) \Delta t^3\right).
\end{aligned}$$

When $\mathbb{E}[Z^3] = 0$, the error term becomes $\mathcal{O}\left((1 + \varepsilon(n)) \Delta t^3\right)$.

Appendix B. Unbiasedness of the mean of Gaussian distributions for Langevin dynamics

We prove that, in the case of one dimensional Gaussian likelihoods, there is no bias on the mean of the posterior distribution when using the discretization of Langevin dynamics (30) (see Section 2.5.1). We start by rewriting the numerical scheme as

$$\begin{pmatrix} \theta^{m+1} \\ p^{m+1} \end{pmatrix} = M_1 \begin{pmatrix} \theta^m \\ p^m \end{pmatrix} + (1 - \alpha_{\Delta t})^{1/2} M_2 \begin{pmatrix} G^m \\ G^{m+1/2} \end{pmatrix} + V_3^m,$$

where

$$M_1 = \begin{pmatrix} 1 - a \frac{\Delta t^2}{2} & \alpha_{\Delta t/2} \left(\Delta t - a \frac{\Delta t^3}{4} \right) \\ -a \Delta t \alpha_{\Delta t/2} & \alpha_{\Delta t} \left(1 - a \frac{\Delta t^2}{2} \right) \end{pmatrix}, \quad M_2 = \begin{pmatrix} \Delta t \left(1 - a \frac{\Delta t^2}{4} \right) & 0 \\ \alpha_{\Delta t/2} \left(1 - a \frac{\Delta t^2}{2} \right) & 1 \end{pmatrix},$$

$$V_3^m = \begin{pmatrix} \frac{\Delta t^2}{2} b^m \\ \alpha_{\Delta t/2} \Delta t b^m \end{pmatrix},$$

and

$$a = \frac{1}{\sigma_\theta^2} + \frac{N_{\text{data}}}{\sigma_x^2}, \quad b^m = \frac{N_{\text{data}}}{n} \sum_{i \in I_n^m} \frac{x_i}{\sigma_x^2}.$$

Since

$$\mathbb{E}[b^m] = \frac{1}{\sigma_x^2} \sum_{i=1}^{N_{\text{data}}} x_i := b,$$

we obtain, by first taking expectations with respect to G^m and $G^{m+1/2}$, and then with respect to realizations of I_n^m in b^m :

$$\mathbb{E} \left[\begin{pmatrix} \theta^m \\ p^m \end{pmatrix} \right] = M_1^m \mathbb{E} \left[\begin{pmatrix} \theta^0 \\ p^0 \end{pmatrix} \right] + \sum_{j=0}^{m-1} M_1^j V_3^j, \tag{83}$$

with $V_2 = (\Delta t^2 b/2, \alpha_{\Delta t/2} \Delta t b)^T$. We note that (recalling that, for the one dimensional case considered here, the friction $\Gamma > 0$ is a scalar)

$$M_1 = I_d - \Delta t \begin{pmatrix} 0 & -1 \\ a & \Gamma \end{pmatrix} + \mathcal{O}(\Delta t^2),$$

which shows that the eigenvalues of M_1 have real parts which are strictly smaller than 1 provided Δt is sufficiently small. Therefore, $M_1^m \rightarrow 0$ as $m \rightarrow +\infty$, and the matrix $I_d - M_1$ is invertible. Moreover, a simple computation shows that

$$(I_d - M_1) \begin{pmatrix} b/a \\ 0 \end{pmatrix} = V_3,$$

so that

$$(I_d - M_1)^{-1} V_3 = \begin{pmatrix} b/a \\ 0 \end{pmatrix}.$$

By letting m go to infinity in (83), and using the above equality, we finally conclude that

$$\lim_{m \rightarrow +\infty} \mathbb{E} \left[\begin{pmatrix} \theta^m \\ p^m \end{pmatrix} \right] = \begin{pmatrix} b/a \\ 0 \end{pmatrix} = \begin{pmatrix} \mu_{\text{post}} \\ 0 \end{pmatrix},$$

which shows that the mean posterior distribution is unbiased.

References

- A. Abdulle, D. Cohen, G. Vilmart, and K. C. Zygalakis. High weak order methods for stochastic differential equations based on modified equations. *SIAM J. Sci. Comput.*, 34(3):A1800–A1823, 2012.
- A. Abdulle, G. Vilmart, and K. C. Zygalakis. High order numerical approximation of the invariant measure of ergodic SDEs. *SIAM J. Numer. Anal.*, 52(4):1600–1622, 2014.
- S. Ahn, A Korattikara, and M. Welling. Bayesian posterior sampling via stochastic gradient Fisher scoring. ICML’12, page 1771–1778, Madison, WI, USA, 2012. Omnipress.
- M. P. Allen and D. J. Tildesley. *Computer Simulation of Liquids*. Oxford University Press, Inc., 2nd edition, 2017.
- J. Baker, P. Fearnhead, E. B. Fox, and C. Nemeth. Control variates for stochastic gradient MCMC. *Stat. Comput.*, 29(3):599–615, 2019.
- D. Bakry, I. Gentil, and M. Ledoux. *Analysis and geometry of Markov diffusion operators*. Springer, 2014.
- R. Bardenet, A. Doucet, and C. Holmes. Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In *International Conference on Machine Learning (ICML)*, Proceedings of the 31st International Conference on Machine Learning (ICML), pages 405–413, Beijing, China, 2014.

- E. Bernard, M. Fathi, A. Levitt, and G. Stoltz. Hypocoercivity with Schur complements. *Annales Henri Lebesgue*, 5:523–557, 2022.
- J. Bierkens, P. Fearnhead, and G. Roberts. The zig-zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Statist.*, 47(3):1288–1320, 2019.
- N. Bou-Rabee and H. Owhadi. Long-run accuracy of variational integrators in the stochastic context. *SIAM J. Numer. Anal.*, 48(1):278–297, 2010.
- G. Brigati and G. Stoltz. How to construct decay rates for kinetic Fokker–Planck equations? *arXiv preprint*, 2302.14506, 2023.
- S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, editors. *Handbook of Markov Chain Monte Carlo*. CRC Press, Boca Raton, FL, 2011.
- N. Brosse, A. Durmus, and E. Moulines. The promises and pitfalls of stochastic gradient Langevin dynamics. In *NeurIPS 2018 (Advances in Neural Information Processing Systems 2018)*, Montreal, France, 2018.
- E. Cancès, F. Legoll, and G. Stoltz. Theoretical and numerical comparison of some sampling methods for molecular dynamics. *ESAIM: Mathematical Modelling and Numerical Analysis*, 41(2):351–389, 2007.
- Y. Cao, J. Lu, and L. Wang. On explicit l^2 -convergence rate estimate for underdamped Langevin dynamics. *arXiv preprint*, 1908.04746, 2019.
- P. Chaudhari and S. Soatto. Stochastic gradient descent performs variational inference converges to limit cycles for deep networks. In *International Conference on Learning Representations*, pages 1–10, 2018.
- T. Chen, E. Fox, and C. Guestrin. Stochastic Gradient Hamiltonian Monte Carlo. In E. P. Xing and T. Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1683–1691, Beijing, China, 22–24 Jun 2014. PMLR.
- N. Chopin, T. Lelièvre, and G. Stoltz. Free energy methods for Bayesian inference: efficient exploration of univariate Gaussian mixture posteriors. *Stat. Comput.*, 22(4):897–916, 2012.
- A. S. Dalalyan and L. Riou-Durand. On sampling from a log-concave density using kinetic Langevin diffusions. *Bernoulli*, 26(3):1956–1988, 2020.
- A. Debussche and E. Faou. Weak backward error analysis for SDEs. *SIAM J. Numer. Anal.*, 50(3):1735–1752, 2012.
- A. Dieuleveut, A. Durmus, and F. Bach. Bridging the gap between constant step size stochastic gradient descent and Markov chains. *The Annals of Statistics*, 48(3):1348–1382, 2020.

- N. Ding, Y. Fang, R. Babbush, C. Chen, R. D. Skeel, and H. Neven. Bayesian sampling using stochastic gradient thermostats. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3203–3211. Curran Associates, Inc., 2014.
- J. Dolbeault, C. Mouhot, and C. Schmeiser. Hypocoercivity for kinetic equations with linear relaxation terms. *C. R. Math. Acad. Sci. Paris*, 347(9-10):511–516, 2009.
- J. Dolbeault, C. Mouhot, and C. Schmeiser. Hypocoercivity for linear kinetic equations conserving mass. *Trans. Amer. Math. Soc.*, 367(6):3807–3828, 2015.
- A. Durmus, U. Simsekli, E. Moulines, R. Badeau, and G. Richard. Stochastic gradient Richardson–Romberg Markov chain Monte Carlo. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- A. Durmus, A. Enfroy, E. Moulines, and G. Stoltz. Uniform minorization condition and convergence bounds for discretizations of kinetic Langevin dynamics. *arXiv preprint*, 2107.14542, 2021.
- D. Frenkel and B. Smit. *Understanding Molecular Simulation, From Algorithms to Applications*. Academic Press, 2nd edition, 2002.
- M. Gürbüzbalaban, X. Gao, Y. Hu, and L. Zhu. Decentralized stochastic gradient Langevin dynamics and Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 22(239):1–69, 2021.
- E. Hairer, C. Lubich, and G. Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, volume 31 of *Springer Series in Computational Mathematics*. Springer-Verlag, 2006.
- W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57:97–109, 1970.
- F. Hérau. Hypocoercivity and exponential time decay for the linear inhomogeneous relaxation Boltzmann equation. *Asymptot. Anal.*, 46(3-4):349–359, 2006.
- D. P. Herzog. Exponential relaxation of the Nosé-Hoover equation under Brownian heating. *Communications in Mathematical Sciences*, 16(8):2231–2260, 2018.
- W. G. Hoover. Canonical dynamics - Equilibrium phase-space distributions. *Phys. Rev. A*, 31(3):1695–1697, 1985.
- P. Izmailov, S. Vikram, M. D. Hoffman, and A. G. Wilson. What are Bayesian neural network posteriors really like? In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4629–4640. PMLR, 2021.
- A. Jones and B. Leimkuhler. Adaptive stochastic methods for sampling driven molecular systems. *J. Chem. Phys.*, 135(8):084125, 2011.

- W. Kliemann. Recurrence and invariant measures for degenerate diffusions. *Ann. Probab.*, 15(2):690–707, 1987.
- M. Kopec. Weak backward error analysis for overdamped Langevin processes. *IMA J. Numer. Anal.*, 35(2):583–614, 2015a.
- M. Kopec. Weak backward error analysis for Langevin process. *BIT Numerical Mathematics*, 55(4):1057–1103, 2015b.
- A. Korattikara, Y. Chen, and M. Welling. Austerity in MCMC land: Cutting the Metropolis–Hastings budget. In E. P. Xing and T. Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 181–189, Beijing, China, 22–24 Jun 2014. PMLR.
- B. Leimkuhler and C. Matthews. Rational construction of stochastic numerical methods for molecular sampling. *Appl. Math. Res. Express*, pages 34–56, 2013.
- B. Leimkuhler and C. Matthews. *Molecular Dynamics: With Deterministic and Stochastic Numerical Methods*. Interdisciplinary Applied Mathematics. Springer, 2015.
- B. Leimkuhler and X. Shang. Adaptive thermostats for noisy gradient systems. *SIAM J. Sci. Comput.*, 38(2):A712–A736, 2016.
- B. Leimkuhler, C. Matthews, and G. Stoltz. The computation of averages from equilibrium and nonequilibrium Langevin molecular dynamics. *IMA J. Numer. Anal.*, 36(1):13–79, 2016.
- B. Leimkuhler, C. Matthews, and T. Vlaar. Partitioned integrators for thermodynamic parameterization of neural networks. *Foundations of Data Science*, 1(4):457–489, 2019.
- B. Leimkuhler, M. Sachs, and G. Stoltz. Hypocoercivity properties of adaptive Langevin dynamics. *SIAM J. Appl. Math.*, 80(3):1197–1222, 2020.
- T. Lelièvre and G. Stoltz. Partial differential equations and stochastic methods in molecular dynamics. *Acta Numerica*, 25:681–880, 2016.
- Y.-A. Ma, Y. Chen, C. Jin, N. Flammarion, and M. I. Jordan. Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, 116(42):20881–20885, 2019.
- C. Matthews and J. Weare. Langevin Markov Chain Monte Carlo with stochastic gradients. *arXiv preprint*, 1805.08863, 2018.
- J. C. Mattingly, A. M. Stuart, and D. J. Higham. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic Process. Appl.*, 101(2):185–232, 2002.
- J. C. Mattingly, A. M. Stuart, and M. V. Tretyakov. Convergence of numerical time-averaging and stationary measures via Poisson equations. *SIAM Journal on Numerical Analysis*, 48(2):552–577, 2010.

- N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6):1087–1091, 1953.
- G. N. Milstein and M. V. Tretyakov. *Stochastic Numerics for Mathematical Physics*. Scientific Computation. Springer Berlin Heidelberg, 2013.
- W. Mou, Y.-A. Ma, M. J. Wainwright, P. L. Bartlett, and M. I. Jordan. High-order Langevin diffusion yields an accelerated MCMC algorithm. *Journal of Machine Learning Research*, 22(42):1–41, 2021.
- T. Nagapetyan, A. Duncan, L. Hasenclever, S. Vollmer, L. Szpruch, and K. Zygalakis. The true cost of Stochastic Gradient Langevin Dynamics. *arXiv preprint*, 1706.02692, 2017.
- S. Nosé. A unified formulation of the constant temperature molecular-dynamics methods. *J. Chem. Phys.*, 81(1):511–519, 1984.
- A. Pakman, D. Gilboa, D. Carlson, and L. Paninski. Stochastic bouncy particle sampler. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2741–2750. PMLR, 2017.
- G. A. Pavliotis. *Stochastic Processes and Applications: Diffusion Processes, the Fokker-Planck and Langevin Equations*. Texts in Applied Mathematics. Springer New York, 2014. ISBN 9781493913220.
- M. Quiroz, R. Kohn, M. Villani, and M. Tran. Speeding up MCMC by efficient data subsampling. *J. Amer. Statist. Assoc.*, 114(526):831–843, 2019.
- H. Robbins and S. Monro. A stochastic approximation method. *Ann. Math. Statistics*, 22:400–407, 1951.
- C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer-Verlag, New York, second edition, 2004.
- G. O. Roberts and J. S. Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 60(1):255–268, 1998.
- G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341 – 363, 1996.
- P. J. Rossky, J. D. Doll, and H. L. Friedman. Brownian dynamics as smart Monte Carlo simulation. *J. Chem. Phys.*, 69(10):4628–4633, 1978.
- X. Shang, Z. Zhu, B. Leimkuhler, and A. J. Storkey. Covariance-controlled adaptive Langevin thermostat for large-scale Bayesian sampling. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28, pages 37–45. Curran Associates, Inc., 2015.
- T. Shardlow. Modified equations for stochastic differential equations. *BIT Numerical Mathematics*, 46(1):111–125, 2006.

- G. Stoltz and Z. Trstanova. Stable and accurate schemes for Langevin dynamics with general kinetic energies. *Multiscale Model. Sim.*, 16(2):777–806, 2018.
- D. Talay. Second-order discretization schemes of stochastic differential systems for the computation of the invariant law. *Stochastics and Stochastic Reports*, 29(1):13–36, 1990.
- D. Talay. Stochastic Hamiltonian dissipative systems: Exponential convergence to the invariant measure, and discretization by the implicit Euler scheme. *Markov Proc. Rel. Fields*, 8:163–198, 2002.
- D. Talay and L. Tubaro. Expansion of the global error for numerical schemes solving stochastic differential equations. *Stochastic Anal. Appl.*, 8(4):483–509, 1990.
- Y. W. Teh, A. H. Thiery, and S. J. Vollmer. Consistency and fluctuations for stochastic gradient Langevin dynamics. *Journal of Machine Learning Research*, 17(7):1–33, 2016.
- M. Tuckerman. *Statistical Mechanics: Theory and Molecular Simulation*. Oxford University Press, 2010.
- A. W. van der Vaart. *Asymptotic Statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998.
- S. J. Vollmer, K. C. Zygalakis, and Y. W. Teh. Exploration of the (non-)asymptotic bias and variance of stochastic gradient Langevin dynamics. *Journal of Machine Learning Research*, 17(159):1–48, 2016.
- H. Wang and D.-Y. Yeung. A survey on Bayesian deep learning. *ACM Comput. Surv.*, 53(5), 2020.
- M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML’11, pages 681–688, USA, 2011. Omnipress. ISBN 978-1-4503-0619-5.
- Z. Zhu, J. Wu, B. Yu, L. Wu, and J. Ma. The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 7654–7663. PMLR, 2019.
- K. C. Zygalakis. On the existence and the applications of modified equations for stochastic differential equations. *SIAM J. Sci. Comput.*, 33(1):102–130, 2011.