Affine invariant interacting Langevin dynamics for Bayesian inference

Alfredo Garbuno Inigo[*]   Nikolas Nüsken[†]   Sebastian Reich[‡]

December 9, 2019

Abstract

We propose a computational method (with acronym ALDI) for sampling from a given target distribution based on first-order (overdamped) Langevin dynamics which satisfies the property of affine invariance. The central idea of ALDI is to run an ensemble of particles with their empirical covariance serving as a preconditioner for their underlying Langevin dynamics. ALDI does not require taking the inverse or square root of the empirical covariance matrix, which enables application to high-dimensional sampling problems. The theoretical properties of ALDI are studied in terms of non-degeneracy and ergodicity. Furthermore, we study its connections to diffusions on Riemannian manifolds and Wasserstein gradient flows.

Bayesian inference serves as a main application area for ALDI. In case of a forward problem with additive Gaussian measurement errors, ALDI allows for a gradient-free implementation in the spirit of the ensemble Kalman filter. A computational comparison between gradient-free and gradient-based ALDI is provided for a PDE constrained Bayesian inverse problem.

Keywords: Langevin dynamics, interacting particle systems, Bayesian inference, gradient flow, multiplicative noise, affine invariance, gradient-free

AMS(MOS) subject classifications: 65N21, 62F15, 65N75, 65C30, 90C56

1 Introduction

In this paper, we propose an efficient sampling method for Bayesian inference which is based on first-order (overdamped) Langevin dynamics [32] and which satisfies the property

---

[*]Computing and Mathematical Sciences, California Institute of Technology, 1200 East California Boulevard, 91125 Pasadena, United States, `agarbuno@caltech.edu`

[†]Institute of Mathematics, University of Potsdam, Karl-Liebknecht-Str. 24/25, D-14476 Potsdam, Germany, `nuesken@uni-potsdam.de`

[‡]Institute of Mathematics, University of Potsdam, Karl-Liebknecht-Str. 24/25, D-14476 Potsdam, Germany, and Department of Mathematics and Statistics, University of Reading, Reading RG6 6AX, England, `sereich@uni-potsdam.de`

of affine invariance [13]. Here affine invariance of a computational method refers to the fact that a method is invariant under an affine change of coordinates. A classical example is provided by Newton's method, while standard gradient descent is not affine invariant. The importance of affine invariance as a general guiding principle for the design of Monte Carlo sampling methods was first highlighted in the pioneering contribution [13].

Langevin dynamics based sampling methods, on the other hand, have a long history in statistical physics [36] and computational statistics [35]. An important step towards affine invariant Langevin sampling methods was taken through the introduction of Riemannian manifold Langevin Monte Carlo methods in [12] with the metric tensor given by the Fisher information matrix. However, the Fisher information matrix is typically not available in closed form and/or is difficult to approximate numerically. Instead, an alternative approach was put forward in the unpublished Master thesis [14], where an ensemble of Langevin samplers is combined to provide an empirical covariance matrix resulting in a preconditioned affine invariant MALA algorithm (see Remark 6 for more details). This methodology was put into the wider context of dynamics-based sampling methods in [21] with a focus on second-order Langevin dynamics.

An interesting link between ensembles of Langevin samplers and the ensemble Kalman filter [19, 33], both relying on ensemble based empirical covariance matrices, has been established more recently in [11] leading to a nonlinear Fokker–Planck equation for the associated mean-field equations and an associated Kalman–Wasserstein gradient flow structure in the space of probability measures. Furthermore, if applied to a Bayesian inverse problem with additive Gaussian measurement errors and nonlinear forward map, a gradient-free Langevin dynamics formulation has been proposed [11] based on ideas previously exploited in the ensemble Kalman filter literature [10].

The present paper builds upon the unpublished note [29], which identifies a statistically consistent finite ensemble size implementation of the mean-field equations put forward in [11]. More precisely, it could be established that the proposed interacting Langevin dynamics possesses the desired posterior target measure as an invariant measure provided an appropriate correction term is added, which is due to the multiplicative noise in the preconditioned Langevin system. The correction term vanishes in the mean-field limit.

In addition, we establish in this paper the affine invariance of our finite ensemble size evolution equations (with acronym ALDI) through a particular choice of the multiplicative noise term and study the ergodicity properties of ALDI. In particular, it can be shown that ALDI is ergodic if the ensemble size, N, and the dimension, D, of the underlying random variable satisfy $N > D+1$ and the empirical covariance matrix is non-degenerate at initial time. We emphasise that ALDI is straightforward to implement, does not require inversion or other matrix factorisations of the empirical covariance matrices (which is important for high-dimensional problems) and is applicable to a wide range of sampling problems.

Finally, a gradient-free formulation of ALDI in the spirit of [11] is proposed for Bayesian inverse problems with additive Gaussian measurement errors. While the invariance of the posterior distribution is lost except for Gaussian likelihood functions, affine invariance

2

is maintained. Numerical experiments are conducted for a PDE constrained Bayesian inference problem. The numerical results indicate in particular that it is entirely sufficient to implement ALDI with $N = D + 2$ particles; the minimum size required for ergodicity to hold.

The remainder of this paper is structured as follows. The subsequent Section 2 establishes the mathematical setting of the sampling problems considered in this paper and provides a unifying mathematical framework for ensemble-based Langevin dynamics equations. Given this framework, we formulate the key algorithmic requirements on the ensemble formulation proposed of this paper. We introduce the concept of affine invariance and prove affine invariance for the nonlinear Fokker–Planck equations put forward in [11]. The novel ALDI method is put forward in Section 3 and its gradient-free variant in Section 4. The affine invariance, non-degeneracy and ergodicity of ALDI are proven in Section 5, where we also put our approach into the perspective of diffusion processes on Riemannian manifolds [12, 23] and Wasserstein gradient flows [1, 38]. The importance of the correction term is demonstrated for a PDE constrained inverse problem [11]. We also compare the performance of the gradient-based and gradient-free formulations of ALDI and find that they both lead to comparable numerical results with the gradient-free formulation however much cheaper to implement. We conclude the paper with a summary section.

2 Mathematical problem formulation

We consider the computational problem of producing samples from a random variable U with values in $\mathbb{R}^D$ and given probability density function (PDF)

$$\pi^*(u) = \frac{1}{Z} \exp(-\Phi(u)) \tag{1}$$

where $\Phi : \mathbb{R}^D \to \mathbb{R}$ is an appropriate potential and

$$Z := \int_{\mathbb{R}^D} \exp(-\Phi(u))\,\mathrm{d}u < \infty \tag{2}$$

a normalisation constant.

Example 1 (Bayesian inverse problems)**.** The computational Bayesian inverse problem (BIP) of sampling a random variable U conditioned on an observation $y \in \mathbb{R}^K$, where the observations Y and the random variable U obey the forward model

$$Y = \mathcal{G}(U) + \Xi, \tag{3}$$

serves as the main motivation of this paper. Here, $\mathcal{G} : \mathbb{R}^D \to \mathbb{R}^K$ denotes some nonlinear forward map and the mean zero $\mathbb{R}^K$-valued Gaussian random variable Ξ represents measurement errors with positive definite error covariance matrix $R \in \mathbb{R}^{K \times K}$. We assume

3

that Ξ and $U \sim \pi_0$ are independent. Then, by Bayes' theorem, the distribution of the conditional random variable $U|y$ is determined by

$$\pi(\mathrm{d}u|y) = \frac{1}{Z} \exp(-l(u;y)) \, \pi_0(\mathrm{d}u), \tag{4}$$

with the least-squares misfit function

$$l(x;y) = \frac{1}{2}\|R^{-\frac{1}{2}}(y - \mathcal{G}(u))\|^2 =: \frac{1}{2}\|y - \mathcal{G}(u)\|_R^2 \tag{5}$$

and the normalisation constant

$$Z = \int_{\mathbb{R}^D} \exp(-l(u;y)) \, \pi_0(\mathrm{d}u) < \infty. \tag{6}$$

If the prior PDF π_0 is Gaussian with mean $m_0 \in \mathbb{R}^D$ and covariance matrix $P_0 \in \mathbb{R}^{D \times D}$, then the posterior is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^D$ with PDF

$$\pi^*(u) = \frac{1}{Z} \exp(-\Phi(u;y)), \tag{7}$$

where

$$\Phi(u;y) := l(u;y) + \frac{1}{2}\|u - m_0\|_{P_0}^2. \tag{8}$$

We write $\Phi(u)$ for simplicity and ignore the dependence on the data y from now on.

The sampling methods considered in this paper are based on stochastic processes of N interacting particles $U_t^{(i)} \in \mathbb{R}^D$, $i = 1, \dots, N$, with the property that the marginal distributions in each of the particles $U_t^{(i)}$ approximate π^* as $t \to \infty$. We treat the particle positions $U_t^{(i)}$ as random vectors with values in $\mathbb{R}^D$, which we collect into the L-dimensional random vector

$$\overline{U}_t = \left(\left(U_t^{(1)}\right)^{\mathrm{T}}, \left(U_t^{(2)}\right)^{\mathrm{T}}, \dots, \left(U_t^{(N)}\right)^{\mathrm{T}} \right)^{\mathrm{T}}, \tag{9}$$

with $L = N \cdot D$, and consider gradient-based evolution equations of the form

$$\mathrm{d}\overline{U}_t = -\mathcal{A}(\overline{U}_t) \, \nabla_{\overline{u}} \mathcal{V}(\overline{U}_t) \, \mathrm{d}t + \Gamma(\overline{U}_t) \, \mathrm{d}W_t, \tag{10}$$

$t \geq 0$, for some positive semi-definite matrix-valued $\mathcal{A}(\overline{u}) \in \mathbb{R}^{L \times L}$ and $\Gamma(\overline{u}) \in \mathbb{R}^{L \times M}$, W_t standard M-dimensional Brownian motion, and $\mathcal{V} : \mathbb{R}^L \to \mathbb{R}$ a given potential. Here $\overline{u}$ denotes an element in $\mathbb{R}^L$ of the form

$$\overline{u} = \left(\left(u^{(1)}\right)^{\mathrm{T}}, \left(u^{(2)}\right)^{\mathrm{T}}, \dots, \left(u^{(N)}\right)^{\mathrm{T}} \right)^{\mathrm{T}}, \tag{11}$$

with all $u^{(i)}$, $i = 1, \dots, N$, elements of $\mathbb{R}^D$. Furthermore, the Itô interpretation [32] of the multiplicative noise term in (10) is to be used.

The main contribution of this paper consists in developing a particular instance of (10) with the following three properties:

4

(i) The product measure

$$\overline{\pi}^*(\overline{u}) = \prod_{i=1}^{N} \pi^* \left(u^{(i)} \right) \tag{12}$$

is invariant under the dynamics. Furthermore, $\overline{\pi}^*$ is ergodic in the sense that the joint law of the process converges towards $\overline{\pi}^*$ as $t \to \infty$, in an appropriate sense and under suitable conditions on the initialisation.

(ii) The equations (10) are invariant under affine transformations

$$U = MV + b \tag{13}$$

with $M \in \mathbb{R}^{D \times D}$ an invertible matrix and $b \in \mathbb{R}^D$ a shift vector. A precise formulation of affine invariance is provided in Definition 2 below. See also [13, 14, 21].

(iii) The equations (10) are straightforward and computationally efficient to implement, that is, do not require the inversion or factorisation of D-dimensional matrices and/or higher-order derivatives of the potential $\mathcal{V}$.

Definition 2 (affine invariance). Following [13, 14, 21], a formulation (10) is called invariant under affine transformations of the form (13), that is,

$$U_t^{(i)} = MV_t^{(i)} + b, \tag{14}$$

if the resulting equations in the collective variable

$$\overline{V}_t = \left(\left(V_t^{(1)} \right)^{\mathrm{T}}, \left(V_t^{(2)} \right)^{\mathrm{T}}, \ldots, \left(V_t^{(N)} \right)^{\mathrm{T}} \right)^{\mathrm{T}} \tag{15}$$

are given by

$$\mathrm{d}\overline{V}_t = -\mathcal{A}(\overline{V}_t)\,\nabla_{\overline{v}}\widetilde{\mathcal{V}}(\overline{V}_t)\,\mathrm{d}t + \Gamma(\overline{V}_t)\,\mathrm{d}W_t, \tag{16}$$

with the potential $\widetilde{\mathcal{V}}$ defined by

$$\widetilde{\mathcal{V}}(\overline{v}) = \mathcal{V}((I_{N \times N} \otimes M)\overline{v} + 1_N \otimes b), \tag{17}$$

where $1_N \in \mathbb{R}^N$ denotes a column vector of ones, $A \otimes B$ the Kronecker product of two matrices, and $I_{N \times N}$ the N-dimensional identity matrix.

Example 3 (Langevin dynamics). The classical example of (10) is provided by the scaled first-order (overdamped) Langevin dynamics

$$\mathrm{d}U_t^{(i)} = -\mathcal{C}\,\nabla_u\Phi(U_t^{(i)})\,\mathrm{d}t + \sqrt{2}\mathcal{C}^{1/2}\mathrm{d}W_t^{(i)}, \tag{18}$$

where $W_t^{(i)}$, $i = 1, \ldots, N$, denotes independent D-dimensional Brownian motion, $\mathcal{C} \in \mathbb{R}^{D \times D}$ is a constant symmetric positive-definite matrix and $\mathcal{C}^{1/2}$ denotes its positive-definite square

5

root. In this case, the particles do not interact and $\mathcal{A} = I_{N \times N} \otimes \mathcal{C}$. Furthermore, $\Gamma = \sqrt{2} I_{N \times N} \otimes \mathcal{C}^{1/2}$ and the potential $\mathcal{V}$ is given by

$$\mathcal{V}(\overline{u}) = \sum_{i=1}^{N} \Phi(u^{(i)}). \tag{19}$$

We note that (18) satisfies items (i) and (iii) from above for any $N \geq 1$ but not (ii), in general. As pointed out in [21], the failure of (18) to be affine invariant potentially leads to inefficient sampling when Φ is poorly scaled with respect to $\mathcal{C}$. More specifically, in the case of Bayesian inverse problems with Gaussian posterior, this scenario occurs when $\mathcal{C}$ is vastly different from the target covariance.

Let $\pi_t^{(i)}$ denote the PDF of the ith particle $U_t^{(i)}$ at time $t \geq 0$ with evolution equation (18). Then these PDFs satisfy the Fokker–Planck equation

$$\partial_t \pi_t = \nabla_u \cdot \left(\pi_t \, \mathcal{C} \, \nabla_u \frac{\delta \mathrm{KL}(\pi_t | \pi^*)}{\delta \pi_t} \right), \tag{20}$$

with $\pi_t = \pi_t^i$ and the Kullback–Leibler divergence defined by

$$\mathrm{KL}(\pi | \pi^*) = \int_{\mathbb{R}^D} \log \left(\frac{\pi(u)}{\pi^*(u)} \right) \pi(\mathrm{d}u). \tag{21}$$

An important generalisation of the linear Fokker–Planck equation (20) has been proposed in [11]. It relies on making the matrix $\mathcal{C}$ dependent on the PDF π_t itself; thus leading to a nonlinear generalisation of (20). More specifically,

$$\mathcal{C}(\pi_t) = \mathbb{E}_{\pi_t} \left[(u - \mu_t)(u - \mu_t)^{\mathrm{T}} \right], \qquad \mu_t = \mathbb{E}_{\pi_t} [u]. \tag{22}$$

This choice of $\mathcal{C}$ is motivated by the gradient flow structure of the ensemble Kalman–Bucy filter, as first revealed in [4], and gives rise to the Kalman–Wasserstein gradient flow structure of the associated nonlinear Fokker–Planck equation (20) as introduced in [11]. Contrary to the classical Fokker–Planck equation (20) with constant $\mathcal{C}$, the nonlinear Fokker–Planck equation with $\mathcal{C} = \mathcal{C}(\pi_t)$ is invariant under affine transformations (13).

Lemma 4 (affine invariance of Kalman–Wasserstein dynamics). *The nonlinear Fokker–Planck equation (20) with $\mathcal{C} = \mathcal{C}(\pi_t)$ is invariant under affine transformations (13).*

Proof. We define the pushforward PDFs

$$\widetilde{\pi}_t(v) = |M| \, \pi_t(Mv + b), \qquad \widetilde{\pi}^*(v) = |M| \, \pi^*(Mv + b). \tag{23}$$

Then

$$\partial_t \widetilde{\pi}_t = |M| \, \partial_t \pi_t \tag{24a}$$

$$= |M| \, \nabla_u \cdot \left(\pi_t \, \mathcal{C}(\pi_t) \, \nabla_u \frac{\delta \mathrm{KL}(\pi_t | \pi^*)}{\delta \pi_t} \right) \tag{24b}$$

$$= \nabla_v \cdot \left(\widetilde{\pi}_t \, \mathcal{C}(\widetilde{\pi}_t) \, \nabla_v \frac{\delta \mathrm{KL}(\widetilde{\pi}_t | \widetilde{\pi}^*)}{\delta \widetilde{\pi}_t} \right) . \tag{24c}$$

Here we have used that

$$\mathcal{C}(\widetilde{\pi}_t) = M \, \mathcal{C}(\pi_t) \, M^{\mathrm{T}}, \tag{25}$$

as well as $\nabla_v \widetilde{f}(v) = \nabla_v f(Mv+b) = M^{\mathrm{T}} \nabla_u f(u)$ for functions $\widetilde{f}(v) = f(Mv+b) = f(u)$ and an analog statement for the divergence operator. Furthermore, the variational derivatives of the Kullback–Leibler divergences satisfy

$$\frac{\delta \mathrm{KL}(\pi_t | \pi^*)}{\delta \pi_t} = \log \left(\frac{\pi}{\pi^*} \right) = \log \left(\frac{\widetilde{\pi}}{\widetilde{\pi}^*} \right) = \frac{\delta \mathrm{KL}(\widetilde{\pi}_t | \widetilde{\pi}^*)}{\delta \widetilde{\pi}_t}. \tag{26}$$

$\square$

Building upon the affine invariance property of the nonlinear Fokker–Planck equation (20) with $\mathcal{C} = \mathcal{C}(\pi_t)$ as in (22), we demonstrate in the following section how to obtain stochastic evolution equations of the form (10) which satisfy all three properties (i)–(iii) from above.

3 Affine invariant interacting Langevin dynamics

As noted in the previous section, the nonlinear Fokker-Planck evolution (20)-(22) satisfies invariance of the target measure $\overline{\pi}^*$ (property (i)) as well as affine invariance (property (ii)). In this section, we address (iii), that is, we present an interacting particle system of the form (10) which has (20) as its mean field limit while still maintaining properties (i) and (ii) for any finite number of particles.

In order to define our interacting particle system, let us first define the empirical covariance and mean

$$\mathcal{C}_N(\overline{u}) = \frac{1}{N} \sum_{i=1}^{N} \left(u^{(i)} - \mu_N(\overline{u}) \right) \left(u^{(i)} - \mu_N(\overline{u}) \right)^{\mathrm{T}}, \qquad \mu_N(\overline{u}) = \frac{1}{N} \sum_{i=1}^{N} u^{(i)}, \tag{27}$$

that is, the particle-based estimators of the quantities defined in (22), and furthermore fix the (non-quadratic) square root

$$\mathcal{C}_N^{1/2}(\overline{u}) := \frac{1}{\sqrt{N}} \left(u^{(1)} - \mu_N(\overline{u}), u^{(2)} - \mu_N(\overline{u}), \cdots, u^{(N)} - \mu_N(\overline{u}) \right) \in \mathbb{R}^{D \times N}. \tag{28}$$

For a moment, let us assume that $\overline{u} \in \mathbb{R}^L$ is such that $\mathcal{C}_N(\overline{u})$ is invertible (we will comment on this assumption following Definition 5, see also Proposition 11) and choose the preconditioning matrix

$$\mathcal{A}(\overline{u}) = I_{N \times N} \otimes \mathcal{C}_N(\overline{u}) = \begin{pmatrix} \mathcal{C}_N(\overline{u}) & 0 & \ldots & 0 \\ 0 & \mathcal{C}_N(\overline{u}) & \ldots & \vdots \\ \vdots & & \ddots & \vdots \\ 0 & \ldots & \ldots & \mathcal{C}_N(\overline{u}) \end{pmatrix} \in \mathbb{R}^{L \times L}, \tag{29}$$

the potential

$$\mathcal{V}(\overline{u}) = \sum_{i=1}^{N} \Phi(u^{(i)}) - \frac{D+1}{2} \log |\mathcal{C}_N(\overline{u})|, \tag{30}$$

and the diffusion matrix

$$\Gamma(\overline{u}) = \sqrt{2} I_{N \times N} \otimes \mathcal{C}_N^{1/2}(\overline{u}) = \sqrt{2} \begin{pmatrix} \mathcal{C}_N^{1/2}(\overline{u}) & 0 & \ldots & 0 \\ 0 & \mathcal{C}_N^{1/2}(\overline{u}) & \ldots & \vdots \\ \vdots & & \ddots & \vdots \\ 0 & \ldots & \ldots & \mathcal{C}_N^{1/2}(\overline{u}) \end{pmatrix} \in \mathbb{R}^{L \times N^2} \tag{31}$$

in (10), that is, $M = N^2$. Using the identity

$$\mathcal{C}_N(\overline{u}) \nabla_{u^{(i)}} \log |\mathcal{C}_N(\overline{u})| = \frac{2}{N} \left(u^{(i)} - \mu_N(\overline{u}) \right), \tag{32}$$

which follows from Jacobi's formula for the derivative of determinants (see the Appendix for more details), we derive the following explicit form of the proposed interacting particle Langevin dynamics.

Definition 5 (ALDI method). The affine invariant Langevin dynamics (ALDI) is given by

$$dU_t^{(i)} = -\mathcal{C}_N(\overline{U}_t) \nabla_{u^{(i)}} \Phi(U_t^{(i)}) \, dt + \frac{D+1}{N} \left(U_t^{(i)} - \mu_N(\overline{U}_t) \right) dt + \sqrt{2} \mathcal{C}_N^{1/2}(\overline{U}_t) \, dW_t^{(i)}, \tag{33}$$

for $i = 1, \ldots, N$, where $W_t^{(i)}$ denotes N-dimensional standard Brownian motion.

We emphasise that the square root $\mathcal{C}_N^{1/2}(\overline{u})$, as defined in (28), does not require computationally expensive matrix factorisations, and hence the formulation (33) satisfies the requirement (iii). Note that although defining $\mathcal{V}$ as in (30) necessitates $N > D$ in order for the empirical covariance matrix $\mathcal{C}_N(\overline{u})$ to be non-singular, the terms in (33) are well-defined also for $N \leq D$. While a non-singular $\mathcal{C}_N(\overline{u})$ is required generically for the ALDI method

to sample from the invariant measure $\overline{\pi}^*$ (see the discussion in Section 5.1), a smaller number of particles, N, is sometimes desirable in order to reduce the computational cost for high-dimensional BIPs.

If indeed $N \leq D$, than $\mathcal{C}_N(\overline{u})$ is singular and the dynamics of the interacting particle system (33) is restricted to the linear subspace spanned by the N initial particle positions $U_0^{(i)}$, that is,

$$U_t^{(i)} = \sum_{j=1}^N m_t^{ij} U_0^{(j)}. \tag{34}$$

Stochastic differential equations in the N^2 scalar coefficients m_t^{ij} can easily be derived from (33) using the *ansatz* (34). In other words, provided that the initial samples $U_0^{(i)}$ are appropriately chosen, an implementation of (33) with $N \leq D$ can lead to a computationally efficient reduction of the BIP onto a lower dimensional linear subspace. The affine invariance of (33) holds regardless of the ensemble size and is discussed in Section 5.2 in more detail.

Remark 6 (previous work)**.** The idea of an affine invariant Monte Carlo method based on Langevin dynamics using an ensemble of particles and its empirical covariance first appeared in the unpublished Master thesis [14]. More specifically, the author proposes an affine invariant modification of the popular MALA algorithm [35, 12], where each particle $U_k^{(i)}$, $i = 1, \ldots, N$, is sequentially updated at time-step k using the proposal

$$U_{k+1}^{(i)} = U_k^{(i)} - h M_k^{(i)} \nabla_u \Phi(U_k^{(i)}) + \sqrt{2h} L_k^{(i)} \Xi_k^{(i)} \tag{35}$$

where $h > 0$ is the step-size, $M_k^{(i)}$ is an empirical covariance matrix based on a set of particles not including $U_k^{(i)}$, $L_k^{(i)}$ is a square root of $M_k^{(i)}$, that is, $M_k^{(i)} = L_k^{(i)}(L_k^{(i)})^{\mathrm{T}}$, and $\Xi_k^{(i)}$ is a D-dimensional Gaussian random variable with mean zero and covariance matrix $I_{D \times D}$. Independently of [14], a general time-continuous framework for affine invariant interacting particle formulations has been developed in [21] and affine invariant implementations of second-order Langevin dynamics using empirical covariance matrices are studied in detail.

More recently, ensemble preconditioned first-order Langevin dynamics has been revisited in [11] with an emphasis on its mean-field limit and its connection to the ensemble Kalman filter [10, 19, 33]. In fact, (33) appeared first in [11] without the correction term

$$\frac{D+1}{N} \left(U_t^{(i)} - \mu_N(\overline{U}_t) \right) \tag{36}$$

and with $\mathcal{C}_N^{1/2}(\overline{u})$ being replaced by the symmetric matrix square root of the covariance matrix $\mathcal{C}_N(\overline{u})$. The correction term (36) in (33) is, however, needed for $\overline{\pi}^*$ to be an invariant distribution under the resulting interacting particle system (10) and first appeared in the unpublished note [29]. The invariance of $\overline{\pi}^*$ under (33) is proven in Section 5.1.

The correction term (36) vanishes as $N \to \infty$ for D fixed which justifies the nonlinear Fokker–Planck equation (20) with $\mathcal{C} = \mathcal{C}(\pi_t)$ in this mean-field limit. See [11] for more details.

We note that a general discussion on necessary correction terms for Langevin dynamics with multiplicative noise can, for example, be found in [35, 12] from the perspective of Riemannian Brownian motion. We also note that general conditions on diffusion processes that guarantee invariance of a given target distribution have been investigated in [8, Section 2.2] and [24, 21].

The specific form of the multiplicative noise term in (33) avoids matrix factorisations and ensures the affine invariance of ALDI as demonstrated in Section 5.2.

4 Approximate gradient-free sampling

A central idea put forward in [11] (see also [31]) in the context of BIPs described in Example 1 is to combine the preconditioned Langevin dynamics with gradient-free formulations of the ensemble Kalman filter. Recalling the forward map $\mathcal{G}$ from (3), the empirical cross-correlation matrix $\mathcal{D}_N(\overline{u}) \in \mathbb{R}^{D \times K}$ is defined via

$$\mathcal{D}_N(\overline{u}) = \frac{1}{N} \sum_{i=1}^{N} \left(u^{(i)} - \mu_N(\overline{u}) \right) \left(\mathcal{G}(u^{(i)}) - \mu_N(\overline{\mathcal{G}}) \right)^{\mathrm{T}}, \qquad \mu_N(\overline{\mathcal{G}}) = \frac{1}{N} \sum_{i=1}^{N} \mathcal{G}(u^{(i)}). \quad (37)$$

We now make the approximation $\mathcal{C}_N \nabla_u \mathcal{G} \approx \mathcal{D}_N$, motivated by the fact that this is exact for affine forward maps, $\mathcal{G}(u) = Gu + c$. For background on this approach, we refer to [10, Appendix A.1]. In terms of the ALDI formulation (33) and the potential $\Phi(u)$, given by (8), we obtain:

Definition 7 (gradient-free ALDI). Given a potential $\Phi(u)$ of the form (8), the gradient-free ALDI formulation is given by

$$\mathrm{d}U_t^{(i)} = - \left\{ \mathcal{D}_N(\overline{U}_t) R^{-1} \left(\mathcal{G}(U_t^{(i)}) - y \right) + \mathcal{C}_N(\overline{U}_t) P_0^{-1}(U_t^{(i)} - m_0) \right\} \mathrm{d}t \qquad (38a)$$

$$+ \frac{D+1}{N} \left(U_t^{(i)} - \mu_N(\overline{U}_t) \right) \mathrm{d}t + \sqrt{2} \mathcal{C}_N^{1/2}(\overline{U}_t) \, \mathrm{d}W_t^{(i)}, \qquad (38b)$$

for $i = 1, \ldots, N$, where $W_t^{(i)}$ denote independent N-dimensional standard Brownian motions.

While the invariance of $\overline{\pi}^*$ is lost under the gradient-free formulation (38) (except, of course, for affine forward operators), affine invariance of the equations of motions is maintained, see Section 5.2. Gradient-free formulations have been found to work well for unimodal posterior distributions in [11], but fail for multi-modal distributions as demonstrated in [34]. A localised covariance formulation of ALDI has been proposed in [34] to overcome this limitation. Localised covariance matrices were already considered in [21]; but not in the context of gradient-free formulations.

10

5   Properties of ALDI

The aim of this section is to analyse some of the properties of the dynamics (33), in particular verifying the conditions (i) and (ii) outlined in Section 2. The key observation (crucially depending on the correction term $\frac{D+1}{2}\log|\mathcal{C}_N(\overline{u})|$ to the potential $\mathcal{V}$ in (30)) is that the corresponding Fokker-Planck equation has the same mathematical structure as its counterpart (20) for the mean-field regime:

Proposition 8 (linear Fokker–Planck equation). *Let $\overline{U}_t$ satisfy the stochastic evolution equations (33) and assume that the time-marginal PDF $\overline{\pi}_t$ of $\overline{U}_t$ is smooth.   Then $\overline{\pi}_t$ satisfies the linear Fokker-Planck equation*

$$\partial_t \overline{\pi}_t = \nabla_{\overline{u}} \cdot \left(\overline{\pi}_t \, \mathcal{A} \, \nabla_{\overline{u}} \frac{\delta \mathrm{KL}(\overline{\pi}_t | \overline{\pi}^*)}{\delta \overline{\pi}_t} \right). \tag{39}$$

Proof. The proof can be found in the Appendix. See also the technical report [29].   □

Note that the PDF $\overline{\pi}_t$ in (39) is defined on the extended space $\mathbb{R}^L$, $L = N \cdot D$, whereas π_t in (20) is defined on $\mathbb{R}^D$. In contrast to (39), the mean field equation (20) is nonlinear since $\mathcal{C}$, as defined in (22), depends on the solution π_t itself.

5.1   Non-degeneracy and ergodicity

As a first result, we have that property (i) is satisfied for the extended target measure (12) on the joint state space $\mathbb{R}^L$. This follows directly from Proposition 8:

Corollary 9 (invariance of the posterior measure). *The extended target measure (12) is invariant for (33), that is, if $\overline{U}_0 \sim \overline{\pi}^*$, then $\overline{U}_t \sim \overline{\pi}^*$ for all $t \geq 0$.*

Proof. Observe that $\mathrm{KL}(\overline{\pi}|\overline{\pi}^*)$ is minimised for $\overline{\pi} = \overline{\pi}^*$, and hence

$$\frac{\delta \mathrm{KL}(\overline{\pi}|\overline{\pi}^*)}{\delta \overline{\pi}} \bigg|_{\overline{\pi}=\overline{\pi}^*} = 0.$$

Using (39), we immediately see that $\partial_t \overline{\pi}^* = 0$, implying the claimed result.   □

Note that $\overline{\pi}^*$ is not the unique invariant measure for the dynamics (33). For instance, if $\overline{U} = (u^{(1)}, \ldots, u^{(N)})$ with $u^{(1)} = u^{(2)} = \ldots = u^{(N)}$, then $\mathcal{C}_N(\overline{U}) = 0$ and $u^i - \mu_N(\overline{U}) = 0$, and hence $\delta_{\overline{U}}$ (the Dirac measure centred at $\overline{U}$) is invariant. To ensure favourable ergodic properties, we need to prove that $\overline{\pi}^*$ is also the unique invariant measure that is reachable by the dynamics from an appropriate set of initial conditions. First, we shall make the following assumption on the potential Φ:

11

Assumption 10 (regularity and growth conditions on the potential Φ). *Assume that* $\Phi \in C^2(\mathbb{R}^D) \cap L^1(\pi^*)$. *Furthermore, assume that there exists a compact set $K \subset \mathbb{R}^D$ and constants $c, C > 0$ such that*

$$c|u|^2 \leq \Phi(u) \leq C|u|^2, \tag{40a}$$

$$c|u| \leq |\nabla\Phi(u)| \leq C|u|, \tag{40b}$$

$$cI_{D\times D} \leq \text{Hess}\,\Phi(u) \leq CI_{D\times D}, \tag{40c}$$

for all $u \in \mathbb{R}^D \setminus K$.

The bound (40c) is to be understood in the sense of quadratic forms. Assumption 10 is satisfied for target measures with Gaussian tails. Indeed, $\Phi = \Phi_0 + \Phi_1$ is admissible, where $\Phi_0(u) = \frac{1}{2}u \cdot Su$ is quadratic (with S strictly positive definite), and $\Phi_1 \in C_c^\infty(\mathbb{R}^D)$ is a smooth perturbation with compact support. We would like to emphasise that Assumption 10 can be relaxed with minimal effort, but we refrain from doing so for ease of exposition.

Due to the fact that $\mathcal{C}_N(\overline{u})$ is not uniformly bounded from below on $\mathbb{R}^L$, the associated Fokker-Planck operator is not uniformly elliptic and standard ergodicity results are not applicable. However, we have the following non-degeneracy result.

Proposition 11 (non-degeneracy of the empirical covariance matrix). *Let Assumption 10 be satisfied and assume that $\mathcal{C}_N(\overline{U}_0)$ is strictly positive definite. Then (33) admits a unique global strong solution, and $\mathcal{C}_N(\overline{U}_t)$ stays strictly positive definite for all $t \geq 0$, almost surely.*

Proof. The proof rests on the identity (32) so that (33) can be written in the form

$$dU_t^{(i)} = -\mathcal{C}_N(\overline{U}_t)\nabla_{u^{(i)}}\mathcal{V}(\overline{U}_t)\,dt + \sqrt{2}\mathcal{C}_N^{1/2}(\overline{U}_t)\,dW_t^{(i)}, \quad i = 1,\ldots,N, \tag{41}$$

with the potential $\mathcal{V}$ given by (30), making use of the repulsive effect of the term

$$-\frac{D+1}{2}\log|\mathcal{C}_N(\overline{u})|.$$

Details can be found in the Appendix. $\qquad\square$

With Proposition 11 in place, the proof of the following ergodicity result is relatively straightforward:

Proposition 12 (ergodicity). *Assume the conditions from Proposition 11, and furthermore that $N > D+1$. Then the dynamics is ergodic, that is, $\overline{\pi}_t \to \overline{\pi}^*$ as $t \to \infty$ in total variation distance.*

Proof. The proof can be found in the Appendix. $\qquad\square$

Remark 13. In the case when $N \leq D$, ergodicity will not hold, since the dynamics is constrained to a subspace according to the discussion following Definition 5. In the case when $N = D+1$ one can show that the set E defined in (45) has two connected components. The dynamics will then be ergodic with respect to $\overline{\pi}^*$ restricted to one of these, depending on the initial condition. This is acceptable from an algorithmic viewpoint, but we do not treat this case separately for simplicity.

5.2 Affine invariance

We show that (33) and its gradient-free variant (38) are affine-invariant, in the terminology introduced in [13, 14] and summarised in Definition 2.

Lemma 14 (affine invariance of ALDI). *The Fokker–Planck equation (39), its associated interacting particle system (33) as well as its gradient-free formulation (38) are all invariant under affine transformations of the form (13).*

Proof. We follow the proof of Lemma 4. Since $\mathcal{C}_N(\overline{u}) = M\,\mathcal{C}_N(\overline{v})M^{\mathrm{T}}$ we also have

$$\mathcal{A}(\overline{u}) = (I_{N\times N} \otimes M)\mathcal{A}(\overline{v})(I_{N\times N} \otimes M)^{\mathrm{T}}. \tag{42}$$

Furthermore,

$$\nabla_{\overline{v}}\widetilde{f}(\overline{v}) = \nabla_{\overline{v}}f((I_{N\times N} \otimes M)\overline{v} + 1_N \otimes b) = (I_{N\times N} \otimes M)^{\mathrm{T}}\nabla_{\overline{u}}f(\overline{u}) \tag{43}$$

for functions $\widetilde{f}(\overline{v}) = f((I_{N\times N} \otimes M)\overline{v} + 1_N \otimes b) = f(\overline{u})$, and an analogous statement holds for the divergence operator. Finally, equality (26) also holds for the Kullback–Leibler divergences over extended state space. Along the same lines, the affine invariance can also be checked directly at the level of the stochastic differential equations (33). In particular, it holds that $\mathcal{C}_N^{1/2}(\overline{u}) = M\mathcal{C}_N^{1/2}(\overline{v})$. Furthermore,

$$\mathcal{D}_N(\overline{u}) = M\mathcal{D}_N(\overline{v}) \tag{44}$$

with $\widetilde{\mathcal{G}}(v) = \mathcal{G}(Mu + b)$ and $\mathcal{D}_N(\overline{v})$ the empirical covariance matrix between v and $\widetilde{\mathcal{G}}(v)$. This implies the affine invariance of the gradient-free formulation (38). □

Remark 15 (path-wise versus distributional affine invariance). Definition 2 is based on path-wise affine invariance at the level of the SDE (10). Path-wise invariance implies affine invariance of the associated time-marginal distributions $\overline{\pi}_t$, that is, affine invariance of the implied Fokker–Planck equation. The converse is not true, in general.

5.3 Geometric properties and gradient flow structure

In this section, we place the dynamics (33) in a geometric context, viewing (a suitable subset of) $\mathbb{R}^L$ as a Riemannian manifold when equipped with an appropriate metric tensor. This approach has been pioneered in [12]; we also recommend the review paper [23]. Leveraging this perspective, we show that the evolution induced by (33) on the set of smooth PDFs can be interpreted as a gradient flow in the sense of [16]. In the limit as $N \to \infty$ we formally recover the Kalman-Wasserstein geometry introduced in [11].

We restrict our attention to the case $N > D + 1$ in this section, when the dynamics (33) is ergodic on the set

$$E = \left\{\overline{u} \in \mathbb{R}^L : \quad \mathcal{C}_N(\overline{u}) \text{ is invertible}\right\} \tag{45}$$

13

according to the results in Section 5.1. Extending the framework to the case when $N \leq D + 1$ is subject of ongoing work. We now turn E into an L-dimensional Riemannian manifold by introducing the metric tensor

$$g = (\mathcal{A}^{-1})_{IJ} \, \mathrm{d}u^I \otimes \mathrm{d}u^J, \tag{46}$$

where $\mathcal{A}$ has been defined in (29). Furthermore, here and in the rest of this section, Einstein's summation convention is used.

Remark 16 (definition of metric tensor). To avoid notational clutter, we have suppressed the coordinates of the individual particles in the expression (46). Denoting the γ-th coordinate of the i-th particle by $u^{(i,\gamma)}$ and using a similar convention for $\mathcal{A}$, we can write more explicitly (not using the Einstein convention for clarity)

$$g = \sum_{i,j=1}^{N} \sum_{\gamma,\sigma=1}^{D} (\mathcal{A}^{-1})_{(i,\gamma),(j,\sigma)} \mathrm{d}u^{(i,\gamma)} \, \mathrm{d}u^{(j,\sigma)} = \sum_{i=1}^{N} \sum_{\gamma,\sigma=1}^{D} (\mathcal{C}_N^{-1})_{\gamma\sigma} \mathrm{d}u^{(i,\gamma)} \, \mathrm{d}u^{(j,\sigma)}, \tag{47}$$

where the second identity rests on the block-diagonal structure of $\mathcal{A}$.

In what follows, we will denote by dvol_g the Riemannian volume, by ∇_g the Riemannian gradient, by $(W_t^g)_{t \geq 0}$ Riemannian Brownian motion and by d_g the geodesic distance on (E, g). For more details, we refer to [15, 20] and, in the context of computational statistics, to [23]. Using these objects induced by g, both the SDE (33) and the corresponding Fokker-Planck equation (39) admit a compact formulation:

Proposition 17 (Riemmanian interpretation of ALDI). *Let $\overline{\pi}^g$ denote the density of $\overline{\pi}$ with respect to the Riemannian volume, that is, $\overline{\pi}^g \, \mathrm{dvol}_g = \overline{\pi} \, \mathrm{d}\overline{u}$ and $\overline{\pi}_*^g \, \mathrm{dvol}_g = \overline{\pi}^* \, \mathrm{d}\overline{u}$. Then the dynamics (33) can be written in the form*

$$\mathrm{d}\overline{U}_t = \nabla_g \log \overline{\pi}_*^g(\overline{U}_t) \, \mathrm{d}t + \sqrt{2} \, \mathrm{d}W_t^g, \tag{48}$$

and the Fokker-Planck equation (39) can be written in the form

$$\partial_t \overline{\pi}_t^g = \nabla_g \cdot \left(\overline{\pi}_t^g \, \nabla_g \frac{\delta \mathrm{KL}(\overline{\pi}_t | \overline{\pi}^*)}{\delta \overline{\pi}_t} \right). \tag{49}$$

Remark 18. Note that the Kullback-Leibler divergence and its functional derivative depend on the measures but not on the respective densities, in contrast to the *Onsager operator* [25, 28, 30] $\phi \mapsto -\nabla_g \cdot (\overline{\pi}_t^g \, \nabla_g \phi)$.

Proof. Using the results from [23], in particular the equations (46)-(47), the proof of the first statement reduces to verifying that

$$\partial_J g^{IJ} = \frac{D+1}{N} \left(u^{(i)} - \mu(\overline{u}) \right)_\gamma, \tag{50}$$

14

where g^{IJ} stands for the components of the inverse of g, and we have used the notation $I = (i, \gamma)$ from Remark 16. This follows directly from the definition of g and the identity (71) giving rise to the drift correction (36). Indeed, together with the coordinate expressions

$$\nabla_g \cdot f = \frac{1}{\sqrt{|g|}} \partial_I \left(\sqrt{|g|} f^I \right), \quad (\nabla_g V)^I = g^{IJ} \partial_J V \tag{51}$$

for vector-valued functions f and scalar-valued V, the result follows by direct substitution.

For the second statement, note that $\mathrm{dvol}_g = \sqrt{|g|}\, \mathrm{d}\overline{u}$, and hence $\overline{\pi}^g = |g|^{-1/2} \overline{\pi}$. $\qquad\square$

To exhibit the gradient flow structure, we recall that the natural quadratic Wasserstein distance on (E, g) is given by

$$\mathcal{W}_g^2(\overline{\mu}, \overline{\nu}) = \inf_{\gamma \in \Pi(\overline{\mu}, \overline{\nu})} \int_{\mathbb{R}^L \times \mathbb{R}^L} d_g^2(\overline{u}, \overline{v})\, \mathrm{d}\gamma(\overline{u}, \overline{v}), \tag{52}$$

where $\Pi(\overline{\mu}, \overline{\nu})$ denotes the set of probability measures on $\mathbb{R}^L \times \mathbb{R}^L$ with marginals $\overline{\mu}$ and $\overline{\nu}$. It is well-known that the evolution (49) can be interpreted as gradient flow dynamics of the Kullback–Leibler divergence on the set of probability measures equipped with the distance (52), see for instance [38, Chapter 15] or [22]. By the Benamou-Brenier formula [3], we have the representation

$$\mathcal{W}_g^2(\overline{\mu}, \overline{\nu}) = \inf_{(\overline{\rho}, \overline{\Phi})} \left\{ \int_0^1 \int_{\mathbb{R}^L} g(\nabla_g \overline{\Phi}_t, \nabla_g \overline{\Phi}_t)\, \mathrm{d}\overline{\rho}_t\, \mathrm{d}t : \right. \tag{53a}$$

$$\left. \partial_t \overline{\rho}_t^g + \nabla_g \cdot (\overline{\rho}_t^g \nabla_g \overline{\Phi}_t) = 0, \quad \overline{\rho}_0 = \overline{\mu},\ \overline{\rho}_1 = \overline{\nu} \right\}, \tag{53b}$$

where the constraining continuity equation in (53b) is to be interpreted in a weak form. In standard coordinates (using the definition (46) as well as the formulas (51)) we see that

$$\mathcal{W}_g^2(\overline{\mu}, \overline{\nu}) = \inf_{(\overline{\rho}, \overline{\Phi})} \left\{ \int_0^1 \int_{\mathbb{R}^L} \nabla \overline{\Phi}_t \cdot \mathcal{A} \nabla \overline{\Phi}_t\, \mathrm{d}\overline{\rho}_t \mathrm{d}t : \partial_t \overline{\rho}_t + \nabla \cdot (\overline{\rho}_t \mathcal{A} \nabla \overline{\Phi}_t) = 0,\ \overline{\rho}_0 = \overline{\mu},\ \overline{\rho}_1 = \overline{\nu} \right\}, \tag{54}$$

revealing a close similarity with the Kalman-Wasserstein distance (here denoted by $\mathcal{W}_{\mathrm{Kalman}}$) introduced in [11]. Indeed, denoting by $\otimes_{i=1}^N \mu_i$ and $\otimes_{i=1}^N \nu_i$ the product measures on $\mathbb{R}^L$ associated to $\mu, \nu \in \mathcal{P}(\mathbb{R}^D)$, we formally expect that

$$\frac{1}{N} \mathcal{W}_g^2(\otimes_{i=1}^N \mu_i, \otimes_{i=1}^N \nu_i) \xrightarrow{N \to \infty} \mathcal{W}_{\mathrm{Kalman}}(\mu, \nu), \tag{55}$$

assuming that $\mathcal{C}_N \approx \mathcal{C}(\rho)$ for sufficiently large N, where $\mathcal{C}(\rho)$ was defined in (22). A rigorous passage from $\mathcal{W}_g$ to the Kalman–Wasserstein distance might be a rewarding direction for future research; we note that a similar analysis (relating the gradient flow structures associated to a finite particle system and its mean-field limit) has been carried out recently in [5].

6 Numerical experiment: A PDE constrained inverse problem

We consider the inverse problem of determining the periodic permeability field $a(x) > 0$ in the elliptic partial differential equation (PDE)

$$- \partial_x(a(x)\partial_x p(x)) = f(x), \qquad x \in \Omega = [0, 2\pi), \tag{56}$$

from $N_y = 10$ observed grid values

$$y_j = p(\mathrm{x}_j) + \eta_j, \qquad \mathrm{x}_j = \frac{2\pi j}{N_y}, \tag{57}$$

$j = 0, \ldots, N_y - 1$, of the pressure field p for a given forcing f. Both p and f are assumed to be periodic over the domain Ω and the measurement errors η_j in (57) are i.i.d. Gaussian with mean zero and variance $\sigma_R = 0.0001$. A related 2-dimensional Darcy flow problem has been studied in [11]. In this paper, we restrict the simulations to the 1-dimensional formulation (56) for computational simplicity.

This infinite-dimensional problem is made finite-dimensional by introducing a computational grid

$$x_i = \frac{2\pi i}{D}, \qquad i = 0, \ldots, D-1, \tag{58}$$

with $D = 50$ grid points. Hence (56) gets replaced by the finite-difference approximation

$$\frac{a_{i+1/2}(p_{i+1} - p_i) - a_{i-1/2}(p_i - p_{i-1})}{h^2} = f_i, \tag{59}$$

$i = 1, \ldots, D$. Here $h = 2\pi/D$ denotes the mesh size and $p_j \approx p(x_j)$, etc. We also make use of the periodicity and set $p_D = p_0$ as well as $f_D = f_0$.

Since the permeability field should be non-negative, we set

$$a_{i-1/2} = e^{u_i} \tag{60}$$

for $i = 1, \ldots, D$. The forward problem is now given by the solution $\{p_i\}_{i=0}^{D-1}$ to (59) for given $\{f_i\}_{i=0}^{D-1}$ and $\{u_i\}_{i=1}^{D}$ and its restriction to the observation grid $\{\mathrm{x}_j\}_{j=0}^{N_y-1}$. We denote this map by $\mathcal{G}(u)$, suppressing the dependence on the forcing given by

$$f_i = \exp\left(-\frac{(2x_i - L)^2}{40}\right). \tag{61}$$

The measurement error covariance matrix is given by $R = \sigma_R I_{N_y \times N_y}$. This completes the description of our forward model (3).

The prior distribution on u is assumed to be Gaussian with mean zero and covariance matrix

$$P_0 = \frac{h}{4}\Delta_h^2 \tag{62}$$

where Δ_h denotes the standard second-order finite-difference operator over Ω with mesh-size h and periodic boundary conditions, that is, the operator defined by the left-hand side of (59) with $a_{i\pm1/2} = 1$. The observations (57) are generated numerically by solving (59) for the reference permeability field

$$a_{i+1/2} = \exp\left(\frac{1}{2}\sin((i+1/2)h)\right) \tag{63}$$

and setting

$$y_j = p_i + \eta_j, \qquad i = (j-1)\frac{D}{N_y}, \qquad \eta_j \sim \mathcal{N}(0, \sigma_R), \tag{64}$$

$j = 1, \ldots, N_y$. From (63) one can read off the reference vector $u^\dagger$ with entries

$$u_i^\dagger = \frac{1}{2}\sin((i-1/2)h), \tag{65}$$

$i = 1, \ldots, D$.

We implemented the gradient-based ALDI formulation (33) as well as the gradient-free ALDI formulation (38) using the Euler–Maruyama method with step-size $\Delta t = 0.01$ over a time interval $t \in [0, 5]$. We also implemented both methods with and without the correction term (36). The ensemble sizes were taken as $N = 25, 52, 100, 200$. Except for the smallest ensemble size, all other choices resulted in non-singular empirical covariance matrices $\mathcal{C}_N(\overline{U}_t)$.

We compare the simulation results based on the estimation bias

$$\text{BIAS} = \frac{h}{T}\int_\tau^{\tau+T} \|\mu_N(\overline{U}_t) - u^\dagger\|^2 \mathrm{d}t \tag{66}$$

and the ensemble spread

$$\text{SPREAD} = \frac{h}{T}\int_\tau^{\tau+T} \text{trace}\left(\mathcal{C}_N(\overline{U}_t)\right) \tag{67}$$

computed along numerical solutions for $\tau = 4$ and $T = 2$. Each experiment was repeated ten times to reduce the impact of random effects. The results can be found in Tables 1 and 2, respectively. It can be seen that the correction term has a profound impact on both the bias as well as the ensemble spread for the smallest ensemble size $N = 25$. This effect is largely diminished for the largest ensemble size of $N = 200$. We also find that the gradient-free formulation yields results which are essentially indistinguishable from those including the exact gradient while being computationally much more efficient. Finally, the

17

N	gf & w/o	gf & w	g & w/o	g & w
25	0.5286	0.4185	0.5196	0.4155
52	0.4109	0.3133	0.4032	0.3048
100	0.3614	0.3229	0.3578	0.3186
200	0.3145	0.3072	0.3097	0.3018

Table 1: Computed estimation bias (66) for ensemble sizes $N \in \{25, 52, 100, 200\}$ and implementations of ALDI with (w) and without (w/o) correction term (36) as well as with exact gradient (g) and gradient-free (gf).

N	gf & w/o	gf & w	g & w/o	g & w
25	0.0109	0.0720	0.0109	0.0728
52	0.0152	0.0474	0.0150	0.0472
100	0.0228	0.0441	0.0227	0.0438
200	0.0332	0.0449	0.0331	0.0448

Table 2: As in Table 1, but reporting the results for the ensemble spread (67).

results for ALDI indicate that it is entirely sufficient to implement it with $N = D + 2$ particles; the minimum size required for ergodicity to hold.

In order to provide a better insight into the impact of the correction term (36) on the final ensemble distributions we display results for $M = 25$ and $M = 200$ in Figures 1 and 2, respectively.

We conclude from this simple experiment that the correction term (36) is required for implementations of ALDI whenever the ensemble size is of the order of the dimension of the parameter space or less. The experiments also confirm that gradient-free implementations can offer a computationally attractive alternative to gradient-based implementations of ALDI.

7   Conclusions

We have proposed a finite ensemble size implementation of the Kalman–Wasserstein gradient flow formalism put forward in [11], which requires the inclusion of a correction term (36) due to the multiplicative nature of the noise in the Langevin equations (33) [29]. In addition to sampling from the desired target distribution, it has also been demonstrated that the equations of motion are affine invariant. While ALDI can be used with $N \leq D$ ensemble members, effectively leading to a linear subspace sampling method, it has also been proven that $N > D+1$ and a non-singular initial empirical covariance matrix $\mathcal{C}_N(\overline{U}_0)$ ensure that $|\mathcal{C}_N(\overline{U}_t)| \neq 0$ for all times and that the equations of motion (33) are ergodic
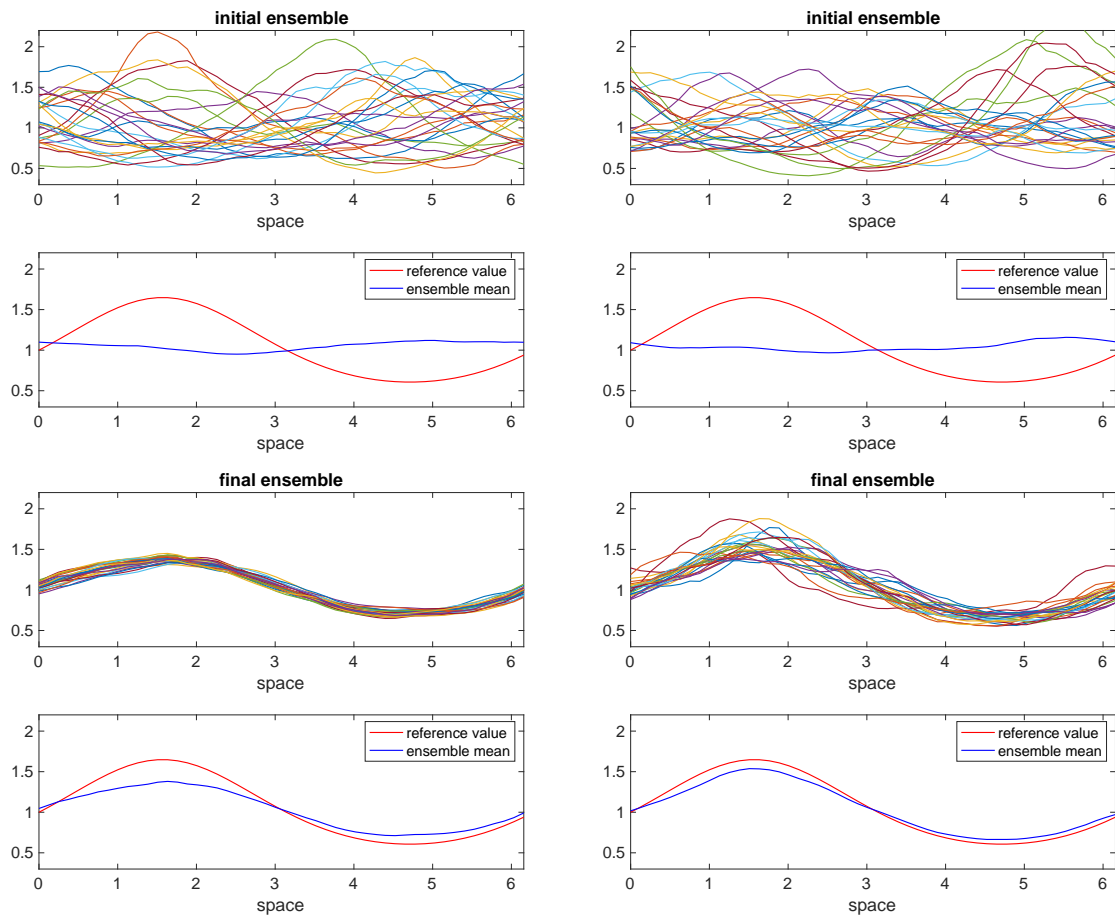
Figure 1: Displayed are the initial (top row) and final (bottom row) ensembles of the permeability fields $a(x) = \exp(u(x))$ for $N = 25$. The left column is from ALDI without correction term (36) while the right column includes the correction term.
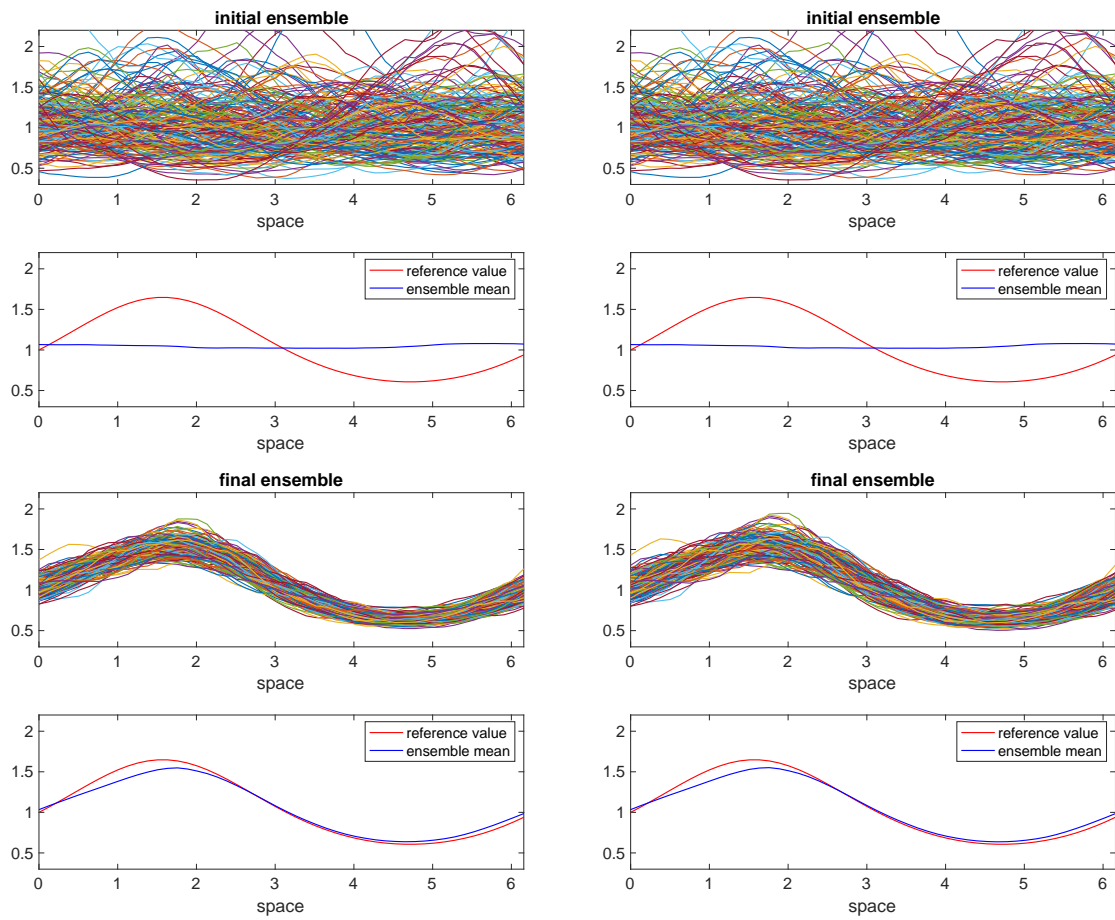
Figure 2: As in Figure 1, except for ensemble size $N = 200$.

with invariant measure $\overline{\pi}^*$. Further computational savings can be achieved through the gradient-free implementation (38) for BIPs as introduced in Example 1. The effectiveness of gradient-free affine invariant sampling methods has been demonstrated for a Darcy flow inversion problem. This example has also demonstrated the significance of the correction term for reducing estimation errors both for $N < D$ as well as for $N = \mathcal{O}(D)$ implementations of the ALDI method (33).

A numerical issue which has not been studied in this paper is the choice of an efficient time-stepping method for ALDI. In particular, adaptive and semi-implicit time-stepping methods might be necessary whenever the initial distribution π_0 is not close to the target measure π^*. This issue has been studied for the related continuous-time ensemble Kalman–Bucy filter in [2]. Wer also reemphasise that multi-model target distributions might require localised empirical covariance matrices in (33) as first suggested in [21] and further explored in [34].

While this paper has focused on a computational implementation of the Kalman–Wasserstein gradient flows proposed in [11], we wish to point out that this gradient flow structure has also become the focus of theoretical studies. We mention in particular [7], which provides a rigorous mean field limit with rates in Wasserstein-2 for the linear case, and [6], which studies the decay for the mean field limit in Wasserstein-2 in the linear case using explicitly the dynamics of the covariance matrix.

Appendix A: Proofs for non-degeneracy and ergodicity

Proof of Proposition 8. The Fokker-Planck equation is given by

$$\partial_t \overline{\pi}_t = \mathcal{L}^\dagger \overline{\pi}_t, \tag{68}$$

where $\mathcal{L}$ denotes the infinitesimal generator of (33) and $\mathcal{L}^\dagger$ refers to its adjoint in $L^2(\mathbb{R}^L)$, given by

$$\left(\mathcal{L}^\dagger \overline{\pi}\right)(\overline{u}) = \sum_{i=1}^N \nabla_{u^{(i)}} \cdot \left(\overline{\pi}(\overline{u}) \left\{ \mathcal{C}_N(\overline{u}) \nabla_{u^{(i)}} \Phi(u^{(i)}) - \frac{D+1}{N}(u^{(i)} - \mu_N(\overline{u})) \right\} \right) \tag{69a}$$

$$+ \sum_{i=1}^N \nabla_{u^{(i)}} \cdot \left\{ \overline{\pi}(\overline{u}) \nabla_{u^{(i)}} \cdot \mathcal{C}_N(\overline{u}) + \mathcal{C}_N(\overline{u}) \nabla_{u^{(i)}} \overline{\pi}(\overline{u}) \right\}. \tag{69b}$$

See [32, Chapter 4] and [21]. Here the divergence of the matrix-valued $\mathcal{C}_N(\overline{u})$ is component-wise given by (in terms of the notation introduced in Remark 16)

$$\{(\nabla_{u^{(i)}} \cdot \mathcal{C}_N(\overline{u}))\}_k = \sum_{\gamma=1}^{D} \frac{\partial}{\partial u^{(i,\gamma)}} \{\mathcal{C}_N(\overline{u})\}_{k\gamma}, \qquad k = 1, \ldots, D. \tag{70}$$

An explicit calculation [29] reveals that

$$\nabla_{u^{(i)}} \cdot \mathcal{C}_N(\overline{u}) = \frac{D+1}{N}(u^{(i)} - \mu_N(\overline{u})) \tag{71}$$

and the Fokker–Planck operator $\mathcal{L}^\dagger$ reduces to

$$\mathcal{L}^\dagger \overline{\pi} = \sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot (\overline{\pi} \mathcal{C}_N \{\nabla_{u^{(i)}} \Phi + \nabla_{u^{(i)}} \log \overline{\pi}\}) = -\nabla_{\overline{u}} \cdot \left(\overline{\pi} \mathcal{A} \nabla_{\overline{u}} \log \frac{\overline{\pi}^*}{\overline{\pi}}\right) \tag{72}$$

from which the desired result follows since

$$\frac{\delta \mathrm{KL}(\overline{\pi}|\overline{\pi}^*)}{\delta \overline{\pi}} = \log \frac{\overline{\pi}^*}{\overline{\pi}}. \tag{73}$$

$\square$

For the proof of Proposition 11 we start by defining the set

$$E = \left\{\overline{u} \in \mathbb{R}^L : \quad \mathcal{C}_N(\overline{u}) \text{ is invertible}\right\}. \tag{74}$$

We will use the potential $\mathcal{V}$ defined in (30) as a Lyapunov function. The key calculation is summarised in the following lemma:

Lemma 19. *There exists a constant $\gamma > 0$ such that*

$$(\mathcal{L}\mathcal{V})(\overline{u}) \leq \gamma \mathcal{V}(\overline{u}), \quad \overline{u} \in E, \tag{75}$$

where $\mathcal{L}$ is the generator of (33), that is, the $L^2(\mathbb{R}^L)$-adjoint of $\mathcal{L}^\dagger$ as defined in (72).

Proof. It follows from (72) that the generator of (33) takes the form

$$(\mathcal{L}\mathcal{V})(\overline{u}) = -\sum_{i=1}^{N} \nabla_{u^{(i)}} \Phi(u^{(i)}) \cdot \mathcal{C}_N(\overline{u}) \nabla_{u^{(i)}} \mathcal{V}(\overline{u}) + \sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot (\mathcal{C}_N(\overline{u}) \nabla_{u^{(i)}} \mathcal{V}(\overline{u})). \tag{76}$$

For convenience, let us introduce the notation

$$\mathcal{V}_{\mathcal{C}}(\overline{u}) = -\frac{D+1}{2} \log |\mathcal{C}_N(\overline{u})|, \qquad \mathcal{V}_\Phi(\overline{u}) = \sum_{i=1}^{N} \Phi(u^{(i)}). \tag{77}$$

22

Since $\mathcal{L}$ vanishes on constants, (75) is equivalent to $\mathcal{L}\mathcal{V} \leq \gamma\mathcal{V} + \tilde{C}$ for some constant $\tilde{C}$. Here and in the following, $\tilde{C}$ denotes a generic constant that can change from line to line. Furthermore, by the growth condition on Φ there exists a constant $\tilde{C}$ such that $-2\mathcal{V}_{\mathcal{C}} \leq \mathcal{V}_{\Phi} + \tilde{C}$. Therefore, it is sufficient to show the bound $\mathcal{L}\mathcal{V} \leq \tilde{C}(1 + \mathcal{V}_{\Phi})$. In the remainder of the proof, we achieve the latter bound term-wise for the contributions in (76).

First note that

$$\nabla_{u^{(i)}} \log |\mathcal{C}_N(\overline{u})| = \frac{2}{N}\mathcal{C}_N^{-1}(\overline{u})(u^{(i)} - \mu_N(\overline{u})). \tag{78}$$

Indeed, again following the notation introduced in Remark 16, we have that

$$\frac{\partial}{\partial u^{(i,\gamma)}} \log |\mathcal{C}_N(\overline{u})| = \left(\frac{1}{|\mathcal{C}_N|} \sum_{\alpha,\beta=1}^{D} \frac{\partial |\mathcal{C}_N|}{\partial \mathcal{C}_N^{\alpha\beta}} \frac{\partial \mathcal{C}_N^{\alpha\beta}}{\partial u^{(i,\gamma)}} \right)(\overline{u}) \tag{79a}$$

$$= \sum_{\alpha,\beta=1}^{D} (\mathcal{C}_N^{-1})_{\alpha\beta} \left(\frac{1}{N}\delta_{\alpha\gamma}(u^{(i)} - \mu_N(\overline{u}))_{\beta} + \frac{1}{N}\delta_{\beta\gamma}(u^{(i)} - \mu_N(\overline{u}))_{\alpha} \right) \tag{79b}$$

$$= \frac{2}{N}\left(\mathcal{C}_N^{-1}(u^{(i)} - \mu_N(\overline{u})) \right)_{\gamma}, \tag{79c}$$

using Jacobi's formula for determinants in the second line. For the first term in (76) we thus obtain

$$-\sum_{i=1}^{N} \nabla\Phi(u^{(i)}) \cdot \mathcal{C}_N(\overline{u})\nabla_{u^{(i)}}\mathcal{V}(\overline{u}) \tag{80a}$$

$$= -\frac{D+1}{N}\sum_{i=1}^{N} \nabla\Phi(u^{(i)}) \cdot (u^{(i)} - \mu_N(\overline{u})) - \underbrace{\sum_{i=1}^{N} \nabla\Phi(u^{(i)}) \cdot \mathcal{C}_N(\overline{u})\nabla\Phi(u^{(i)})}_{\leq 0} \tag{80b}$$

$$\leq \tilde{C}\left(1 + \mathcal{V}_{\Phi}(\overline{u}) \right). \tag{80c}$$

To bound the second term in (76), we first notice the estimate

$$\sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot \left(\mathcal{C}_N(\overline{u})\nabla_{u^{(i)}}\Phi(u^{(i)}) \right) \leq \tilde{C}\left(1 + \mathcal{V}(\overline{u}) \right), \tag{81}$$

again easily obtained from Assumption 10. The other contribution is

$$\sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot (\mathcal{C}_N(\overline{u})\nabla_{u^{(i)}}\mathcal{V}_{\mathcal{C}}(\overline{u})) = -\frac{D+1}{N}\sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot \left(u^{(i)} - \mu_N(\overline{u}) \right) \tag{82a}$$

$$= -\frac{(D+1)D}{2N}\left(N - 1 \right). \tag{82b}$$

23

Since the result is a constant, we clearly have the required estimate of the form $\mathcal{L}\mathcal{V} \leq \tilde{C}(1 + \mathcal{V}_\Phi)$. In conjunction with (80) and (81) the claim follows. $\qquad\square$

Proposition 11 now essentially follows from adapting [27, Theorem 2.1]. Textbook accounts of similar arguments can be found in [17, Chapter 5] and [9, Chapter 2]. For the convenience of the reader we provide a self-contained proof:

Proof of Proposition 11. The potential $\mathcal{V}$ is bounded from below by the growth condition on Φ (see Assumption 10). We can therefore choose a constant $c_{\mathcal{V}}$ such that $\mathcal{V}_+ := \mathcal{V} + c_{\mathcal{V}}$ is nonnegative. Since $\mathcal{C}_N(\overline{U}_0)$ is assumed to be nondegenerate, there exists $k_0 \in \mathbb{N}$ such that $\mathcal{V}_+(\overline{U}_0) < k_0$. For $k \geq k_0$, let us define the sets

$$E_k = \{\overline{u} \in E : \quad \mathcal{V}_+(\overline{u}) < k\} \tag{83}$$

and the stopping times

$$\tau_k = \inf\{t \geq 0 \; : \overline{U}_t \notin E_k\} = \inf\left\{t \geq 0 : \mathcal{V}_+(\overline{U}_t) = k\right\}. \tag{84}$$

The stopping times τ_k are increasing in k, and so the limit

$$\lim_{k\to\infty} \tau_k =: \xi$$

exists in $[0, +\infty]$. To prove the claim, it is sufficient to show that $\mathbb{P}[\xi = +\infty] = 1$. We now define

$$g(\overline{u}, t) := e^{-\gamma t}\mathcal{V}_+(\overline{u}), \quad (\overline{u}, t) \in E \times [0, \infty), \tag{85}$$

where γ is the constant obtained in Lemma 19. By using Itô's formula, optional stopping, and the bound (75) we see that

$$\mathbb{E}\left[g(\overline{U}_{t\wedge\tau_k}, t \wedge \tau_k)\right] = g(\overline{U}_0, 0) + \mathbb{E}\left[\int_0^{t\wedge\tau_k} e^{-\gamma s}\left(-\gamma\mathcal{V}_+(\overline{U}_s) + \mathcal{L}\mathcal{V}_+(\overline{U}_s)\right)\mathrm{d}s\right] \tag{86a}$$

$$\leq g(\overline{U}_0, 0) = \mathcal{V}_+(\overline{U}_0), \tag{86b}$$

for any $t \geq 0$ and $k \geq k_0$. On the other hand,

$$\mathbb{E}\left[g(\overline{U}_{t\wedge\tau_k}, t \wedge \tau_k)\right] \geq e^{-\gamma t}\mathbb{E}\left[\mathcal{V}_+(\overline{U}_{t\wedge\tau_k})\right] \tag{87a}$$

$$\geq e^{-\gamma t}\left(\mathbb{E}\left[\mathbf{1}_{t<\tau_k}\mathcal{V}_+(\overline{U}_t)\right] + \mathbb{P}\left[\tau_k \leq t\right] \cdot k\right) \geq e^{-\gamma t}\left(\mathbb{P}\left[\tau_k \leq t\right] \cdot k\right), \tag{87b}$$

where the last estimate uses the fact that $\mathcal{V}_+ \geq 0$. Combining (86) and (87), we see that

$$e^{\gamma t}\mathcal{V}_+(\overline{U}_0) \geq \mathbb{P}\left[\tau_k \leq t\right] \cdot k, \tag{88}$$

for every $t \geq 0$ and $k \geq k_0$. It follows immediately that

$$\lim_{k\to\infty} \mathbb{P}\left[\tau_k \leq t\right] = 0, \tag{89}$$

24

and further

$$\mathbb{P}\left[\xi \le t\right] = 0 \tag{90}$$

by monotone convergence. Since (90) holds for all $t \ge 0$, we conclude that $\mathbb{P}[\xi = \infty] = 1$, as required. $\qquad\square$

For the proof of Proposition 12 we will need the following lemma:

Lemma 20. *Let $N \ge D + 2$. Then E is path-connected.*

Proof. We begin by fixing some additional notation. For $j \in \{1, \ldots, N\}$, define the 'leave-one-out' versions of the empirical mean and covariance,

$$\mu_N^{-j}(\overline{u}) = \frac{1}{N-1} \sum_{\substack{i=1 \\ i \ne j}}^{N} u^{(i)}, \qquad \mathcal{C}_N^{-j}(\overline{u}) = \frac{1}{N-1} \sum_{\substack{i=1 \\ i \ne j}}^{N} (u^{(i)} - \mu_N^{-j}(\overline{u}))(u^{(i)} - \mu_N^{-j}(\overline{u}))^{\mathrm{T}}. \tag{91}$$

Notice the update formula

$$\mathcal{C}_N(\overline{u}) = \frac{N-1}{N} \mathcal{C}_N^{-j}(\overline{u}) + \frac{N-1}{N^2} (u^{(j)} - \mu_N^{-j}(\overline{u}))(u^{(j)} - \mu_N^{-j}(\overline{u}))^{\mathrm{T}}, \tag{92}$$

holding for any $j \in \{1, \ldots N\}$. Consider now the set

$$\widetilde{E} := \left\{ \overline{u} \in \mathbb{R}^L : \quad \mathcal{C}_N^{-j} \text{ is invertible for all } j \in \{1, \ldots, N\} \right\}. \tag{93}$$

We see that $\widetilde{E} \subset E$ since the second term on the right-hand side of (92) is positive semidefinite. Importantly, the condition $N \ge D + 2$ ensures that $\widetilde{E}$ is nonempty.

Observe that $\widetilde{E}$ has the representation

$$\widetilde{E} = \left\{ \overline{u} \in \mathbb{R}^L : \quad \prod_{j=1}^{N} |\mathcal{C}_N^{-j}(\overline{u})| > 0 \right\}, \tag{94}$$

immediately implying that $\widetilde{E}$ is open. We now show that $\widetilde{E}$ is dense in $\mathbb{R}^L$. To this end, let $\overline{x} \in \mathbb{R}^L$ and $\overline{y} \in \widetilde{E}$. It is sufficient to prove that for every $\varepsilon > 0$ there exists $t \in (0, \varepsilon)$ such that $(1-t)\overline{x} + t\overline{y} \in \widetilde{E}$. For this, define $P : \mathbb{R} \to \mathbb{R}$ by

$$P(t) = \prod_{j=1}^{N} |\mathcal{C}_N^{-j}((1-t)\overline{x} + ty)|, \tag{95}$$

which is clearly a polynomial. Since $\overline{y} \in \widetilde{E}$ we have that $P(1) > 0$, and so P has only finitely many zeroes. This proves that indeed for all $\varepsilon > 0$ there exists $t \in (0, \varepsilon)$ such that $P(t) > 0$, and hence $(1-t)\overline{x} + t\overline{y} \in \widetilde{E}$.

We now show how to construct a continuous path between arbitrary $\overline{x}, \overline{y} \in E$. By density of $\widetilde{E}$, it is enough to find a path between $\widetilde{x} \in \widetilde{E}$ and $\widetilde{y} \in \widetilde{E}$ lying in connected neighbourhoods of $\overline{x}$ and $\overline{y}$ respectively. Since $\widetilde{E}$ is open, there exist open neighbourhoods $U_{\widetilde{x}}, U_{\widetilde{y}} \subset \widetilde{E}$. It is then sufficient to find points in these neighbourhoods that can be connected by a continuous path. Denoting $\widetilde{x} = (x^{(1)}, \ldots, x^{(N)})$ and $\widetilde{y} = (y^{(1)}, \ldots, y^{(N)})$, we can choose a continuous path $\gamma^{(1)} : [0, 1] \to \mathbb{R}^D$ with $\gamma^{(1)}(0) = x^{(1)}$ and $\gamma^{(1)}(1) = y^{(1)}$, and set $\overline{\gamma}^{(1)}(t) = (\gamma^1(t), x^{(2)}, \ldots, x^{(N)})$. By (92), it is clear that $\overline{\gamma}^{(1)}(t) \in E$ for all $t \in [0, 1]$. By density of $\widetilde{E}$ we can perturb $\overline{\gamma}^{(1)}(1)$ in order to ensure that $\overline{\gamma}^{(1)}(1) \in \widetilde{E}$. We can now proceed iteratively to move the remaining particles using paths $\overline{\gamma}^{(2)}, \ldots \overline{\gamma}^{(N)}$ and concatenate them, yielding the required total path. Note that the perturbation of the endpoints $\overline{\gamma}^{(i)}$ can be chosen arbitrarily small in order to ensure that the final point $\overline{\gamma}^{(N)}(1)$ is in $U_{\widetilde{y}}$. $\qquad \square$

Proof of Proposition 12. Since the diffusion matrix $\Gamma(\overline{u})\Gamma(\overline{u})^{\mathrm{T}}$ is strictly positive definite on E, E is path-connected by Lemma 20, and the process $(\overline{U}_t)_{t \geq 0}$ admits an invariant measure with strictly positive Lebesgue-density by Corollary 9, the process is positively recurrent and irreducible by the result in [18]. We also refer to [37, Section 2.2.2.1]. The convergence in total variation distance then follows from [26, Theorem 6.1]. $\qquad \square$

References

[1] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows: In metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.

[2] J. Amezcua, E. Kalnay, K. Ide, and S. Reich. Ensemble transform Kalman-Bucy filters. *Q.J.R. Meteor. Soc.*, 140:995–1004, 2014.

[3] J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge–Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

[4] K. Bergemann and S. Reich. A localization technique for ensemble Kalman filters. *Q. J. R. Meteorological Soc.*, 136:701–707, 2010.

[5] JA Carrillo, MG Delgadino, and GA Pavliotis. A proof of the mean-field limit for λ-convex potentials by Γ-convergence. *arXiv preprint arXiv:1906.04601*, 2019.

[6] J.A. Carrillo and U. Vaes. Wasserstein stability estimates for covariance-preconditioned Fokker–Planck equations. *arXiv preprint arXiv:1910.07555*, 2019.

[7] Z. Ding and Q. Li. Ensemble Kalman sampling: Mean-field limit and convergence analysis. *arXiv preprint arXiv:1910.12923*, 2019.

[8] A.B. Duncan, N. Nüsken, and G.A. Pavliotis. Using perturbed underdamped Langevin dynamics to efficiently sample from probability distributions. *Journal of Statistical Physics*, 169(6):1098–1131, 2017.

[9] A. Eberle. Markov processes. *Lecture Notes at University of Bonn*, 2009.

[10] G. Evensen. *Data Assimilation. The Ensemble Kalman Filter*. Springer-Verlag, New York, 2006.

[11] A. Garbuno-Inigo, F. Hoffmann, W. Li, and A.M. Stuart. Interacting Langevin diffusions: Gradient structure and ensemble Kalman sampler. Technical Report arXiv:1903:08866v2, Caltech, 2019.

[12] M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Statist. Soc. B*, 73:123–214, 2011.

[13] J. Goodman and J. Weare. Ensemble samplers with affine invariance. *Comm. Appl. Math. and Comput. Science*, 5:65–80, 2010.

[14] P. Greengard. An ensemblized Metropolized Langevin sampler. Master's thesis, Courant Institute, New York University, 2015.

[15] E. P. Hsu. *Stochastic analysis on manifolds*, volume 38. American Mathematical Soc., 2002.

[16] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker–Planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.

[17] R. Khasminskii. *Stochastic stability of differential equations*, volume 66. Springer Science & Business Media, 2011.

[18] W. Kliemann. Recurrence and invariant measures for degenerate diffusions. *The annals of probability*, 15(2):690–707, 1987.

[19] K. Law, A. Stuart, and K. Zygalakis. *Data assimilation: A mathematical introduction*. Springer-Verlag, New York, 2015.

[20] J.M. Lee. *Riemannian manifolds: An introduction to curvature*, volume 176. Springer Science & Business Media, 2006.

[21] B. Leimkuhler, Ch. Matthews, and J. Weare. Ensemble preconditioning for Markov chain Monte Carlo simulations. *Stat. Comput.*, 28:277–290, 2018.

[22] S. Lisini. Nonlinear diffusion equations with variable coefficients as gradient flows in Wasserstein spaces. *ESAIM: Control, Optimisation and Calculus of Variations*, 15(3):712–740, 2009.

[23] S. Livingstone and M. Girolami. Information-geometric Markov chain Monte Carlo methods using diffusions. *Entropy*, 16(6):3074–3102, 2014.

[24] Y.-A. Ma, T. Chen, and E. Fox. A complete recipe for stochastic gradient MCMC. In *Advances in Neural Information Processing Systems*, pages 2917–2925, 2015.

[25] S. Machlup and L. Onsager. Fluctuations and irreversible process. ii. systems with kinetic energy. *Physical Review*, 91(6):1512, 1953.

[26] S.P. Meyn and R.L. Tweedie. Stability of Markovian processes ii: Continuous-time processes and sampled chains. *Advances in Applied Probability*, 25(3):487–517, 1993.

[27] S.P. Meyn and R.L. Tweedie. Stability of Markovian processes iii: Foster–Lyapunov criteria for continuous-time processes. *Advances in Applied Probability*, 25(3):518–548, 1993.

[28] A. Mielke, D.R.M Renger, and M.A. Peletier. A generalization of Onsagers reciprocity relations to gradient flows with nonlinear mobility. *Journal of Non-Equilibrium Thermodynamics*, 41(2):141–149, 2016.

[29] N. Nüsken and S. Reich. Note on interacting Langevin diffusion: Gradient structure and ensemble Kalman sampler. Technical Report arXiv:1908.10890v1, University of Potsdam, 2019.

[30] H.C. Öttinger. *Beyond equilibrium thermodynamics.* John Wiley & Sons, 2005.

[31] S. Pathiraja and S. Reich. Discrete gradients for computational Bayesian inference. *Journal of Computational Dynamics*, 6:236–251, 2019.

[32] G.A. Pavliotis. *Stochastic processes and applications.* Springer–Verlag, New York, 2014.

[33] S. Reich and C.J. Cotter. *Probabilistic forecasting and Bayesian data assimilation.* Cambridge University Press, Cambridge, 2015.

[34] S. Reich and S. Weissmann. Fokker–Planck particle systems for Bayesian inference: Computational approaches. Technical Report arXiv:1911.10832, University of Potsdam, 2019.

[35] G.O. Roberts and O. Stramer. Langevin diffusions and Metropolis–Hastings algorithms. *Methodol. Comput. App. Probab.*, 4:337–358, 2003.

[36] P.J. Rossky, J.D. Doll, and H.L. Friedmann. Brownian dynamics as smart Monte Carlo simulation. *J. Chem. Phys.*, 69:4628–4633, 1978.

[37] G. Stoltz and M. Rousset. *Free energy computations: A mathematical perspective.* World Scientific, 2010.

[38] C. Villani. *Optimal transport: Old and new*, volume 338. Springer Science & Business Media, 2008.