Affine Invariant Interacting Langevin Dynamics for Bayesian Inference[*]

Alfredo Garbuno-Inigo[†], Nikolas Nüsken[‡], and Sebastian Reich[§]

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 nondegeneracy and ergodicity. Furthermore, we study its connections to diffusion 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 approximation 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.

Key words. Langevin dynamics, interacting particle systems, Bayesian inference, gradient flow, multiplicative noise, affine invariance, gradient-free

AMS subject classifications. 65N21, 62F15, 65N75, 65C30, 90C56

DOI. 10.1137/19M1304891

1. Introduction. In this paper, we propose an efficient sampling method for Bayesian inference which is based on first-order (overdamped) Langevin dynamics [33] and which satisfies the property 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

[†]Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125 (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).

statistical physics [39] and computational statistics [38]. 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's 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 subsection 3.3 for more details). This methodology was put into the wider context of dynamics-based sampling methods in [22] with a focus on second-order Langevin dynamics.

An interesting link between ensembles of Langevin samplers and the ensemble Kalman filter [10, 20, 35], 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. The same gradient flow structure was previously identified for the time-continuous ensemble Kalman–Bucy filter mean-field equations [36, 35]. Furthermore, if applied to a Bayesian inverse problem with additive Gaussian measurement errors and nonlinear forward map, a gradient-free approximate Langevin dynamics formulation was proposed [11] which is again based on ideas previously exploited in the ensemble Kalman filter literature [10, 4].

The present paper builds upon the unpublished note [30], which identifies a statistically consistent finite ensemble size implementation of the mean-field equations put forward in [11]. More precisely, 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. Furthermore, the invariance of our finite ensemble size evolution equations (with acronym ALDI[1]) under affine coordinate transformations is established through a particular choice of the multiplicative noise term, among all choices consistent with the desired underlying Fokker–Planck equation. We emphasize that ALDI is straightforward to implement, does not require inversion or other matrix factorizations of the empirical covariance matrices (which is important for high-dimensional problems), and is applicable to a wide range of sampling problems.

We have already emphasized that related computational methods have been considered in the literature before. However, none of these contributions has investigated the nondegeneracy and ergodicity properties of such methods. Hence, proof of nondegeneracy and ergodicity of ALDI provides a key theoretical contribution of our paper which holds provided the ensemble size, N, and the dimension, D, of the underlying random variable satisfy $N > D + 1$ and the empirical covariance matrix is nondegenerate at initial time.

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 when making the gradient-free approximation, except for Gaussian likelihood functions, affine invariance is maintained. Numerical experiments are conducted for

---

[1]The acronym stands for a permutation of the capital letters in Affine Invariant Langevin Dynamics.

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. Thus the gradient-free implementation indeed provides an accurate and computationally inexpensive alternative.

The remainder of this paper is structured as follows. Section 2 establishes the mathematical setting of the sampling problems considered in this paper and provides a unifying mathematical framework for ensemble-based first-order Langevin dynamics. Given this framework, we formulate the key algorithmic requirements on the ensemble formulation proposed in this paper. We introduce the concept of affine invariance and prove affine invariance for the nonlinear Fokker–Planck equations put forward in [11]. The algorithmic contributions of this paper can be found in section 3. More specifically, the novel ALDI method is put forward in subsection 3.1 and its gradient-free variant in subsection 3.2. Both methods are put into the context of previous algorithmic work in subsection 3.3. Our theoretical investigations are summarized in section 4, where the affine invariance, nondegeneracy, and ergodicity of ALDI are proven. We also put our approach into the perspective of diffusion processes on Riemannian manifolds [12, 24] and Wasserstein gradient flows [1, 42]. The importance of the correction term is demonstrated for a PDE constrained inverse problem [11] in the numerical example section 5. We also compare the performance of the gradient-based and gradient-free formulations of ALDI and find that 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)

$$(2.1) \qquad \pi_*(u) = \frac{1}{Z} \exp(-\Phi(u)),$$

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

$$(2.2) \qquad Z := \int_{\mathbb{R}^D} \exp(-\Phi(u)) \, \mathrm{d}u < \infty$$

a normalization constant.

Example 2.1 (Bayesian inverse problems). The computational Bayesian inverse problem (BIP) of sampling a random variable u conditioned on an observation $y_{\mathrm{obs}} \in \mathbb{R}^K$ with forward model

$$(2.3) \qquad y = \mathcal{G}(u) + \xi$$

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 that ξ and $u \sim \pi_0$ are independent. Then, by Bayes' theorem, the distribution of the conditional random variable $u|y_{\mathrm{obs}}$ is determined by

$$(2.4) \qquad \pi(\mathrm{d}u|y_{\mathrm{obs}}) = \frac{1}{Z} \exp(-l(u; y_{\mathrm{obs}})) \, \pi_0(\mathrm{d}u),$$

with the least-squares misfit function[2]

$$l(u; y_{\text{obs}}) = \frac{1}{2}\|R^{-\frac{1}{2}}(y_{\text{obs}} - \mathcal{G}(u))\|^2 =: \frac{1}{2}\|y_{\text{obs}} - \mathcal{G}(u)\|_R^2 \tag{2.5}$$

and the normalization constant

$$Z = \int_{\mathbb{R}^D} \exp(-l(u; y_{\text{obs}}))\, \pi_0(\mathrm{d}u) < \infty. \tag{2.6}$$

If the prior PDF π_0 is Gaussian with mean $\mu_0 \in \mathbb{R}^D$ 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

$$\pi_*(u) = \frac{1}{Z} \exp(-\Phi(u; y_{\text{obs}})), \tag{2.7}$$

where

$$\Phi(u; y_{\text{obs}}) := l(u; y_{\text{obs}}) + \frac{1}{2}\|u - \mu_0\|_{P_0}^2. \tag{2.8}$$

We write $\Phi(u)$ for simplicity and ignore the dependence on the data y_{obs} from now on.

The sampling methods considered in this paper are based on stochastic processes of N interacting particles moving in $\mathbb{R}^D$ with the property that the marginal distributions in each of the particles approximate π_* as $t \to \infty$. The position of the ith particle is denoted by $u^{(i)} \in \mathbb{R}^D$ and its value at time $t \geq 0$ by $u_t^{(i)}$, $i = 1, \dots, N$. For ease of reference, we collect all particle positions into the $D \times N$-dimensional matrix

$$U = \left(u^{(1)}, u^{(2)}, \dots, u^{(N)}\right) \in \mathbb{R}^{D \times N}. \tag{2.9}$$

The interacting particle systems considered in this paper obey gradient-based stochastic evolution equations of the form

$$\mathrm{d}u_t^{(i)} = -\mathcal{A}(U_t)\, \nabla_{u^{(i)}} \mathcal{V}(U_t)\, \mathrm{d}t + \Gamma(U_t)\, \mathrm{d}W_t^{(i)}, \qquad i = 1, \dots, N. \tag{2.10}$$

Specific choices for the potential $\mathcal{V} : \mathbb{R}^{D \times N} \to \mathbb{R}$, the positive semidefinite matrix-valued $\mathcal{A}(U) \in \mathbb{R}^{D \times D}$, and $\Gamma(U) \in \mathbb{R}^{D \times L}$ will be discussed below. L is a natural number with typically either $L = D$ or $L = N$. The $W_t^{(i)}$ denote independent L-dimensional standard Brownian motions, and the Itô interpretation [33] of the multiplicative noise term in (2.10) is to be used.

The main algorithmic contribution of this paper consists in developing a particular instance of (2.10) with the following three properties:

(i) The product measure

$$\pi_*^{(N)}(U) := \prod_{i=1}^{N} \pi_*\left(u^{(i)}\right) \tag{2.11}$$

---

[2]Here we have introduced the weighted l_2-norm $\|a\|_B = (a^{\mathrm{T}} B^{-1} a)^{1/2}$ for any symmetric positive definite matrix B.

is invariant under (2.10). Furthermore, $\pi_*^{(N)}$ is ergodic in the sense that the joint law of the process converges towards $\pi_*^{(N)}$ as $t \to \infty$, in an appropriate sense and under suitable conditions on the initialization. See [33] for an introduction to ergodicity in the context of stochastic evolution equations.

(ii) Equations (2.10) are invariant under affine transformations of the state variables, that is, for transformations of the form

$$(2.12) \qquad u = Mv + b$$

for any invertible $M \in \mathbb{R}^{D \times D}$ and any shift vector $b \in \mathbb{R}^D$. A precise definition of affine invariance is provided in Definition 2.2 below. See also [13, 14, 22].

(iii) Equations (2.10) are straightforward and computationally efficient to implement; that is, they do not require the inversion or factorization of D-dimensional matrices and/or higher-order derivatives of the potential $\mathcal{V}$.

Definition 2.2 (affine invariance). *Following* [13, 14, 22], *a formulation* (2.10) *is called* affine invariant *under transformations of the form* (2.12), *that is,*

$$(2.13) \qquad u^{(i)} = Mv^{(i)} + b,$$

if the resulting equations in the transformed particle positions are given by

$$(2.14) \qquad \mathrm{d}v_t^{(i)} = -\mathcal{A}(V_t) \nabla_{v^{(i)}} \widetilde{\mathcal{V}}(V_t) \, \mathrm{d}t + \Gamma(V_t) \, \mathrm{d}W_t^{(i)}, \qquad i = 1, \ldots, N,$$

for any invertible matrix $M \in \mathbb{R}^{D \times D}$ and any shift vector $b \in \mathbb{R}^D$. Here

$$(2.15) \qquad V = \left(v^{(1)}, v^{(2)}, \ldots, v^{(N)} \right) \in \mathbb{R}^{D \times N},$$

and the potential $\widetilde{\mathcal{V}}$ is defined by

$$(2.16) \qquad \widetilde{\mathcal{V}}(V) = \mathcal{V}(U) = \mathcal{V}(MV + b \, 1_N^{\mathrm{T}}),$$

where $1_N \in \mathbb{R}^N$ denotes a column vector of ones.

Example 2.3 (Langevin dynamics). The classical example of (2.10) is provided by the scaled first-order (overdamped) Langevin dynamics

$$(2.17) \qquad \mathrm{d}u_t^{(i)} = -C \, \nabla_{u^{(i)}} \Phi \left(u_t^{(i)} \right) \, \mathrm{d}t + \sqrt{2} C^{1/2} \mathrm{d}W_t^{(i)},$$

where $W_t^{(i)}$, $i = 1, \ldots, N$, denotes independent D-dimensional Brownian motion, $C \in \mathbb{R}^{D \times D}$ is a constant symmetric positive definite matrix, and $C^{1/2}$ denotes its symmetric positive definite square root. In this case, the particles do not interact and $\mathcal{A} = C$. Furthermore, $\Gamma = \sqrt{2} C^{1/2}$ and the potential $\mathcal{V}$ is given by

$$(2.18) \qquad \mathcal{V}(U) = \sum_{i=1}^{N} \Phi \left(u^{(i)} \right).$$

We note that (2.17) satisfies items (i) and (iii) above for any $N \geq 1$ but not (ii), in general. As pointed out in [22], the failure of (2.17) to be affine invariant potentially leads to inefficient sampling when Φ is poorly scaled with respect to C. More specifically, in the case of Bayesian inverse problems with Gaussian posterior, this scenario occurs when 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 (2.17). Then these PDFs satisfy the Fokker–Planck equation

$$\partial_t \pi_t = \nabla_u \cdot \left(\pi_t \, C \, \nabla_u \frac{\delta \mathrm{KL}(\pi_t | \pi_*)}{\delta \pi_t} \right), \tag{2.19}$$

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

$$\mathrm{KL}(\pi | \pi_*) = \int_{\mathbb{R}^D} \log \left(\frac{\pi(u)}{\pi_*(u)} \right) \pi(\mathrm{d}u). \tag{2.20}$$

It has been shown in [16] that the Fokker–Planck equation (2.19) corresponds to a gradient flow structure in the space of probability measures. Furthermore, since the variational derivative of the Kullback–Leibler divergence is given by

$$\frac{\delta \mathrm{KL}(\pi_t | \pi_*)}{\delta \pi_t} = \log \pi_t - \log \pi_*, \tag{2.21}$$

the invariance of the product measure (2.11) under the stochastic evolution equations (2.17) follows immediately.

An important generalization of the linear Fokker–Planck equation (2.19) was proposed in [11]. It relies on making the matrix C dependent on the PDF π_t itself, thus leading to a nonlinear generalization of (2.19). More specifically, the nonlinear Fokker–Planck equation is given by

$$\partial_t \pi_t = \nabla_u \cdot \left(\pi_t \, C(\pi_t) \, \nabla_u \frac{\delta \mathrm{KL}(\pi_t | \pi_*)}{\delta \pi_t} \right), \tag{2.22}$$

with

$$C(\pi_t) = \mathbb{E}_{\pi_t} \left[(u - \mu_t)(u - \mu_t)^{\mathrm{T}} \right], \qquad \mu_t = \mathbb{E}_{\pi_t} [u]. \tag{2.23}$$

This choice of C is motivated by the ensemble Kalman–Bucy filter [34, 36, 11]. The associated generalized gradient flow structure in the space of probability measures was first stated in [36] in the context of the ensemble Kalman–Bucy filter mean-field equations and has been discussed in detail under the notion of the so-called Kalman–Wasserstein gradient flow structure in [11]. See subsection 3.3 and Remark 4.11 below for more details.

A key observation for the present paper is that, contrary to the classical Fokker–Planck equation (2.19) with constant C, the nonlinear Fokker–Planck equation (2.22) is affine invariant.

Lemma 2.4 (affine invariance of Kalman–Wasserstein dynamics). *The nonlinear Fokker–Planck equation* (2.22) *is affine invariant.*

Proof. We define the pushforward PDFs

$$\widetilde{\pi}_t(v) = |M|\, \pi_t(Mv + b), \qquad \widetilde{\pi}_*(v) = |M|\, \pi_*(Mv + b). \tag{2.24}$$

Then

$$\partial_t \widetilde{\pi}_t = |M|\, \partial_t \pi_t \tag{2.25a}$$

$$= |M|\, \nabla_u \cdot \left(\pi_t\, C(\pi_t)\, \nabla_u \frac{\delta \mathrm{KL}(\pi_t | \pi_*)}{\delta \pi_t} \right) \tag{2.25b}$$

$$= \nabla_v \cdot \left(\widetilde{\pi}_t\, C(\widetilde{\pi}_t)\, \nabla_v \frac{\delta \mathrm{KL}(\widetilde{\pi}_t | \widetilde{\pi}_*)}{\delta \widetilde{\pi}_t} \right). \tag{2.25c}$$

Here we have used that

$$C(\widetilde{\pi}_t) = M\, C(\pi_t)\, M^{\mathrm{T}}, \tag{2.26}$$

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

$$\frac{\delta \mathrm{KL}(\pi_t | \pi_*)}{\delta \pi_t} = \log\left(\frac{\pi}{\pi_*} \right) = \log\left(\frac{\widetilde{\pi}}{\widetilde{\pi}_*} \right) = \frac{\delta \mathrm{KL}(\widetilde{\pi}_t | \widetilde{\pi}_*)}{\delta \widetilde{\pi}_t}. \tag{2.27} \qquad \blacksquare$$

Building upon the affine invariance property of the nonlinear Fokker–Planck equation (2.22), we demonstrate in the following section how to obtain stochastic evolution equations of the form (2.10) which satisfy all three properties (i)–(iii) from above. Their theoretical properties are studied in the subsequent section 4. In particular, we establish nondegeneracy and ergodicity, which provides the key theoretical contribution of this paper.

3. Affine invariant interacting Langevin dynamics. As noted in the previous section, the nonlinear Fokker–Planck evolution (2.22)–(2.23) satisfies invariance of the target measure $\pi_*^{(N)}$ (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 (2.10) which has (2.22) as its mean-field limit while still maintaining properties (i) and (ii) for any finite number of particles. We also introduce a gradient-free approximation which is applicable to BIPs of the form (2.8). This section concludes with a summary of related previous algorithmic work.

3.1. ALDI: An exact gradient-based sampling method. In order to define our interacting particle system, let us first define the empirical covariance matrix

$$\mathcal{C}(U) := \frac{1}{N} \sum_{i=1}^{N} \left(u^{(i)} - m(U) \right) \left(u^{(i)} - m(U) \right)^{\mathrm{T}} \tag{3.1}$$

with empirical mean

$$m(U) := \frac{1}{N} \sum_{i=1}^{N} u^{(i)} = \frac{1}{N} U\, 1_N, \tag{3.2}$$

that is, the particle-based estimators of the quantities defined in (2.23). We also introduce the $D \times N$ matrix of the deviations of the particle positions from their mean value, that is,

$$(3.3) \qquad U' := \left(u^{(1)} - m(U), u^{(2)} - m(U), \ldots, u^{(N)} - m(U) \right) = U - m(U)\, 1_N^{\mathrm{T}},$$

which allows us to write

$$(3.4) \qquad \mathcal{C}(U) = \frac{1}{N} U'(U')^{\mathrm{T}}.$$

Furthermore, we define a generalized (nonsymmetric) square root of $\mathcal{C}(U)$ via

$$(3.5) \qquad \mathcal{C}^{1/2}(U) := \frac{1}{\sqrt{N}} U',$$

that is, $\mathcal{C} = \mathcal{C}^{1/2} \left(\mathcal{C}^{1/2} \right)^{\mathrm{T}}$. For a moment, let us assume that $U \in \mathbb{R}^{D \times N}$ is such that $\mathcal{C}(U)$ is invertible (we will comment on this assumption following Definition 3.1; see also Proposition 4.4) and choose the preconditioning matrix

$$(3.6) \qquad \mathcal{A}(U) = \mathcal{C}(U) = \frac{1}{N} U'(U')^{\mathrm{T}},$$

the potential

$$(3.7) \qquad \mathcal{V}(U) = \sum_{i=1}^{N} \Phi\left(u^{(i)} \right) - \frac{D+1}{2} \log |\mathcal{C}(U)|,$$

and the diffusion matrix

$$(3.8) \qquad \Gamma(U) = \sqrt{2}\, \mathcal{C}^{1/2}(U) = \frac{\sqrt{2}}{\sqrt{N}} U'$$

in (2.10), that is, $L = N$. Note that the potential (3.7) contains the additional $-(D+1)/2 \log |\mathcal{C}(U)|$ term in comparison to (2.18), which is required to keep the target distribution (2.11) invariant under the state-dependent diffusion matrix $\mathcal{C}(U)$. See Proposition 4.1 below and [30] for details.

Using the identity

$$(3.9) \qquad \mathcal{C}(U) \nabla_{u^{(i)}} \log |\mathcal{C}(U)| = \frac{2}{N} \left(u^{(i)} - m(U) \right),$$

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 3.1 (ALDI method). *The affine invariant Langevin dynamics (ALDI) is given by the interacting particle system*

$$(3.10) \qquad \mathrm{d}u_t^{(i)} = -\mathcal{C}(U_t) \nabla_{u^{(i)}} \Phi\left(u_t^{(i)} \right) \mathrm{d}t + \frac{D+1}{N} \left(u_t^{(i)} - m(U_t) \right) \mathrm{d}t + \sqrt{2}\, \mathcal{C}^{1/2}(U_t)\, \mathrm{d}W_t^{(i)}$$

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

We emphasize that the generalized square root $\mathcal{C}^{1/2}(U)$, as defined in (3.5), does not require a computationally expensive Cholesky factorization of $\mathcal{C}(U)$, and hence the formulation (3.10) satisfies the requirement (iii). Note that although defining $\mathcal{V}$ as in (3.7) necessitates $N > D$ in order for the empirical covariance matrix $\mathcal{C}(U)$ to be nonsingular, the terms in (3.10) are well-defined also for $N \leq D$. While a nonsingular $\mathcal{C}(U)$ is required generically for the ALDI method to sample from the desired target measure $\pi_*^{(N)}$ (see the discussion in subsection 4.1), a smaller number of particles, N, is sometimes desirable in order to reduce the computational cost for high-dimensional BIPs.

If indeed $N \leq D$, then $\mathcal{C}(U)$ is singular and the dynamics of the interacting particle system (3.10) is restricted to the linear subspace spanned by the N initial particle positions $u_0^{(i)}$, that is,

$$(3.11) \qquad u_t^{(i)} = \sum_{j=1}^{N} m_t^{ij} u_0^{(j)}.$$

Stochastic differential equations (SDEs) in the N^2 scalar coefficients m_t^{ij} can easily be derived from (3.10) using the *ansatz* (3.11). In other words, provided that the initial samples $u_0^{(i)}$ are appropriately chosen, an implementation of (3.10) with $N \leq D$ can lead to a computationally efficient reduction of the BIP onto a lower-dimensional linear subspace. The affine invariance of (3.10) holds regardless of the ensemble size and is discussed in subsection 4.2 in more detail.

3.2. Approximate gradient-free sampling. A central idea put forward in [11] (see also [32]) in the context of BIPs described in Example 2.1 is to combine the preconditioned Langevin dynamics with gradient-free formulations of the ensemble Kalman filter. Recalling the forward map $\mathcal{G}$ from (2.3), the empirical cross-correlation matrix $\mathcal{D}(U) \in \mathbb{R}^{D \times K}$ is defined via

$$(3.12) \qquad \mathcal{D}(U) = \frac{1}{N} \sum_{i=1}^{N} \left(u^{(i)} - m(U) \right) \left(\mathcal{G}(u^{(i)}) - m(\mathcal{G}(U)) \right)^{\mathrm{T}}$$

with empirical mean

$$(3.13) \qquad m(\mathcal{G}(U)) = \frac{1}{N} \sum_{i=1}^{N} \mathcal{G}(u^{(i)}) = \frac{1}{N} \mathcal{G}(U) \, 1_N.$$

We now make the approximation $\mathcal{C}(U)\nabla_u \mathcal{G}(u) \approx \mathcal{D}(U)$, motivated by the fact that this approximation becomes exact for affine forward maps, $\mathcal{G}(u) = Gu + c$. We refer the reader to [10, Appendix A.1] for more details. In terms of the ALDI formulation (3.10) and the potential $\Phi(u)$, given by (2.8), we obtain the following definition.

Definition 3.2 (gradient-free ALDI). *Given a potential $\Phi(u)$ of the form (2.8), the gradient-free ALDI formulation is given by the interacting particle system*

$$(3.14\mathrm{a}) \qquad \mathrm{d}u_t^{(i)} = - \left\{ \mathcal{D}(U_t)R^{-1} \left(\mathcal{G}\left(u_t^{(i)}\right) - y_{\mathrm{obs}} \right) + \mathcal{C}(U_t)P_0^{-1}(u_t^{(i)} - \mu_0) \right\} \mathrm{d}t$$

$$(3.14\mathrm{b}) \qquad + \frac{D+1}{N} \left(u_t^{(i)} - m(U_t) \right) \mathrm{d}t + \sqrt{2}\mathcal{C}^{1/2}(U_t) \, \mathrm{d}W_t^{(i)}$$

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

While the invariance of $\pi_*^{(N)}$ is lost under the gradient-free formulation (3.14) (except, of course, for affine forward operators), affine invariance of the equations of motions is maintained; see subsection 4.2.

3.3. Related previous algorithmic 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's thesis [14]. More specifically, the author proposes an affine invariant modification of the popular MALA algorithm [38, 12], where each particle $u_k^{(i)}$, $i = 1, \ldots, N$, is sequentially updated at time step k using the proposal

$$(3.15) \qquad u_{k+1}^{(i)} = u_k^{(i)} - h M_k^{(i)} \nabla_{u_k^{(i)}} \Phi(u_k^{(i)}) + \sqrt{2h} L_k^{(i)} \xi_k^{(i)},$$

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 the Cholesky factor 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 was developed in [22], studying in detail affine invariant implementations of second-order Langevin dynamics using empirical covariance matrices.

More recently, ensemble preconditioned first-order Langevin dynamics was revisited in [11] with an emphasis on its mean-field limit and its connection to the ensemble Kalman filter [10, 20, 35]. In fact, (3.10) appeared first in [11] with the potential (3.7) replaced by (2.18), that is, without the correction term

$$(3.16) \qquad \frac{D+1}{N} \left(u_t^{(i)} - m(U_t) \right),$$

and with $\mathcal{C}^{1/2}(U)$ being replaced by the symmetric matrix square root of the covariance matrix $\mathcal{C}(U)$. The resulting method is called the ensemble Kalman sampler (EKS) in [11]. The correction term (3.16) is, however, needed in (3.10) in order for $\pi_*^{(N)}$ to be an invariant distribution under the resulting interacting particle system (2.10) and first appeared in the unpublished note [30]. The invariance of $\pi_*^{(N)}$ under (3.10) is proven in subsection 4.1.

The correction term (3.16) vanishes as $N \to \infty$ for D fixed, which justifies the nonlinear Fokker–Planck equation (2.22) 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 [38, 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 were investigated in [8, section 2.2] and [25, 22].

The gradient-free approximation of the form $\mathcal{C}(U) \nabla_u \mathcal{G}(u) \approx \mathcal{D}(U)$ originated in the ensemble Kalman filter literature [10]. More precisely, the time-continuous formulation of the ensemble Kalman filter, the so-called ensemble Kalman–Bucy filter given by

$$(3.17) \qquad \mathrm{d}u_t^{(i)} = -\mathcal{C}(U_t) \nabla_{u^{(i)}} \mathcal{G}\left(u_t^{(i)} \right) R^{-1} \left(\frac{1}{2} \left\{ \mathcal{G}\left(u_t^{(i)} \right) + m(\mathcal{G}(U_t)) \right\} - y_{\mathrm{obs}} \right),$$

fits into the interacting particle dynamics framework (2.10) with $\mathcal{A}(U) = \mathcal{C}(U)$, $\Gamma(U) = 0$, and

$$(3.18) \qquad \mathcal{V}(U) = \frac{1}{4} \sum_{i=1}^{N} \left\| y_{\text{obs}} - \mathcal{G}\left(u^{(i)}\right) \right\|_R^2 + \frac{1}{4} \| y_{\text{obs}} - m(\mathcal{G}(U)) \|_R^2.$$

See [34, 36, 35] for more details. Its gradient-free formulation becomes

$$(3.19) \qquad \mathrm{d}u_t^{(i)} = -\mathcal{D}(U_t) R^{-1} \left(\frac{1}{2} \left\{ \mathcal{G}\left(u_t^{(i)}\right) + m(\mathcal{G}(U_t)) \right\} - y_{\text{obs}} \right)$$

[4, 35]. The derivative-free ensemble Kalman inversion (EKI) method [40, 19] is a slight modification of (3.19) with the mean contribution $m(\mathcal{G}(U))$ replaced by $\mathcal{G}(u_t^{(i)})$. This modification leads to a faster decay in the ensemble deviations U_t' and, hence, in the covariance matrix $\mathcal{D}(U_t)$ while retaining the evolution equation in the ensemble mean $m(U_t)$.

The extension of such gradient-free formulations to Langevin dynamics was first proposed in [11]. Gradient-free formulations have been found to work well for unimodal posterior distributions in [11], but fail for multimodal distributions, as demonstrated in [37]. A localized covariance formulation of ALDI was proposed in [37] to overcome this limitation. Localized covariance matrices were already considered in [22], but not in the context of gradient-free formulations.

4. Theoretical analysis of ALDI. The aim of this section is to analyze some of the properties of the dynamics (3.10), in particular verifying conditions (i) and (ii) outlined in section 2. The key observation (crucially depending on the correction term $\frac{D+1}{2} \log |\mathcal{C}(U)|$ to the potential $\mathcal{V}$ in (3.7)) is that the corresponding Fokker–Planck equation has the same mathematical structure as its counterpart (2.22) for the mean-field regime.

Proposition 4.1 (linear Fokker–Planck equation). *Let U_t, as defined by (2.9), satisfy the stochastic evolution equations (3.10) and assume that the time-marginal PDF $\pi_t^{(N)}$ of U_t is smooth. Then $\pi_t^{(N)}$ satisfies the linear Fokker–Planck equation*

$$(4.1) \qquad \partial_t \pi_t^{(N)} = \sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot \left(\pi_t^{(N)} \, \mathcal{C} \, \nabla_{u^{(i)}} \frac{\delta \mathrm{KL}\left(\pi_t^{(N)} | \pi_*^{(N)}\right)}{\delta \pi_t^{(N)}} \right).$$

Proof. The proof can be found in the appendix. See also the technical report [30]. ∎

Note that the PDF $\pi_t^{(N)}$ in (4.1) is defined on the extended space $\mathbb{R}^{D \times N}$, whereas π_t in (2.22) is defined on $\mathbb{R}^D$. In contrast to (4.1), the mean-field equation (2.22) is nonlinear since $C(\pi_t)$ depends on the solution π_t itself.

4.1. Nondegeneracy and ergodicity. As a first result, we have that property (i) is satisfied for the extended target measure (2.11) on the joint state space $\mathbb{R}^{D \times N}$. This follows directly from Proposition 4.1.

Corollary 4.2 (invariance of the posterior measure). *The extended target measure (2.11) is invariant for (3.10); that is, if $U_0 \sim \pi_*^{(N)}$, then $U_t \sim \pi_*^{(N)}$ for all $t \geq 0$.*

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

$$\left.\frac{\delta \mathrm{KL}\left(\pi^{(N)}|\pi_*^{(N)}\right)}{\delta \pi^{(N)}}\right|_{\pi^{(N)} = \pi_*^{(N)}} = 0.$$

Using (4.1), we immediately see that $\partial_t \pi_*^{(N)} = 0$, implying the claimed result. ∎

Note that $\pi_*^{(N)}$ is not the unique invariant measure for the dynamics (3.10). For instance, if $U = \left(u^{(1)}, \ldots, u^{(N)}\right)$ with $u^{(1)} = u^{(2)} = \cdots = u^{(N)}$, then $\mathcal{C}(U) = 0$ and $u^{(i)} = m(U)$, and hence δ_U (the Dirac measure centered at U) is invariant. To ensure favorable ergodic properties, we need to prove that $\pi_*^{(N)}$ is 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:

Assumption 4.3 (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 exist a compact set $K \subset \mathbb{R}^D$ and constants $c_2 > c_1 > 0$ such that

(4.2a) $$c_1|u|^2 \leq \Phi(u) \leq c_2|u|^2,$$

(4.2b) $$c_1|u| \leq |\nabla\Phi(u)| \leq c_2|u|,$$

(4.2c) $$c_1 I_{D \times D} \leq \mathrm{Hess}\,\Phi(u) \leq c_2 I_{D \times D}$$

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

The bound (4.2c) is to be understood in the sense of quadratic forms. Assumption 4.3 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 $S \in \mathbb{R}^{D \times D}$ strictly positive definite), and $\Phi_1 \in C_c^\infty(\mathbb{R}^D)$ is a smooth perturbation with compact support. We would like to emphasize that Assumption 4.3 can be relaxed with minimal effort, but we refrain from doing so for ease of exposition.

Due to the fact that $\mathcal{C}(U)$ is not uniformly bounded from below on $\mathbb{R}^{D \times N}$, the associated Fokker–Planck operator is not uniformly elliptic and standard ergodicity results are not applicable. However, we have the following nondegeneracy result.

Proposition 4.4 (nondegeneracy of the empirical covariance matrix). *Let Assumption* 4.3 *be satisfied and assume that $\mathcal{C}(U_0)$ is strictly positive definite. Then (3.10) admits a unique global strong solution, and $\mathcal{C}(U_t)$ stays strictly positive definite for all $t \geq 0$, almost surely.*

Proof. The proof rests on the identity (3.9) so that (3.10) can be written in the form

(4.3) $$\mathrm{d}u_t^{(i)} = -\mathcal{C}(U_t)\nabla_{u^{(i)}}\mathcal{V}(U_t)\,\mathrm{d}t + \sqrt{2}\mathcal{C}^{1/2}(U_t)\,\mathrm{d}W_t^{(i)}, \quad i = 1, \ldots, N,$$

with the potential $\mathcal{V}$ given by (3.7), making use of the repulsive effect of the term

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

Details can be found in the appendix. ∎

With Proposition 4.4 in place, the proof of the following ergodicity result is relatively straightforward.

Proposition 4.5 (ergodicity). *Assume the conditions from Proposition 4.4, and furthermore that $N > D + 1$. Then the dynamics is ergodic, that is, $\pi_t^{(N)} \to \pi_*^{(N)}$ as $t \to \infty$ in total variation distance.*

Proof. The proof can be found in the appendix. ∎

Remark 4.6. 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 3.1. In the case when $N = D + 1$ one can show that the set

$$(4.4) \qquad E = \left\{ U \in \mathbb{R}^{D \times N} : \quad \mathcal{C}(U) \text{ is invertible} \right\}$$

has two connected components. The dynamics will then be ergodic with respect to $\pi_*^{(N)}$ 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.

4.2. Affine invariance. We show that (3.10) and its gradient-free variant (3.14) are affine-invariant in the terminology introduced in [13, 14] and summarized in Definition 2.2.

Lemma 4.7 (affine invariance of ALDI). *The Fokker–Planck equation (4.1), its associated interacting particle system (3.10), as well as its gradient-free formulation (3.14) are all affine invariant.*

Proof. We follow the proof of Lemma 2.4. Since $\mathcal{C}(U) = M\mathcal{C}(V)M^\mathrm{T}$ we also have $\mathcal{A}(U) = M\mathcal{A}(V)M^\mathrm{T}$. Furthermore,

$$(4.5) \qquad \nabla_{v^{(i)}} \widetilde{f}(V) = \nabla_{v^{(i)}} f\left(MV + b\,\mathbb{1}_N^\mathrm{T} \right) = M^\mathrm{T} \nabla_{u^{(i)}} f(U)$$

for functions $\widetilde{f}(V) = f(U) = f\left(MV + b\,\mathbb{1}_N^\mathrm{T} \right)$, and an analogous statement holds for the divergence operator. Finally, equality (2.27) 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 (3.10). In particular, it holds that $\mathcal{C}^{1/2}(U) = M\mathcal{C}^{1/2}(V)$. Furthermore,

$$(4.6) \qquad \mathcal{D}(U) = M\mathcal{D}(V)$$

with $\widetilde{\mathcal{G}}(v) = \mathcal{G}(Mu + b)$ and $\mathcal{D}(V)$ the empirical covariance matrix between v and $\widetilde{\mathcal{G}}(v)$. This implies the affine invariance of the gradient-free formulation (3.14). ∎

Remark 4.8 (pathwise versus distributional affine invariance). Definition 2.2 is based on pathwise affine invariance at the level of the SDEs (2.10). Pathwise invariance implies affine invariance of the associated time-marginal distributions $\pi_t^{(N)}$, that is, affine invariance of the implied Fokker–Planck equation. The converse is not true, in general.

4.3. Geometric properties and gradient flow structure. In this section, we place the dynamics (3.10) in a geometric context, viewing (a suitable subset of) $\mathbb{R}^{D \times N}$ as a Riemannian manifold when equipped with an appropriate metric tensor. This approach has been pioneered

in [12]; we also recommend the review paper [24]. Leveraging this perspective, we show that the evolution induced by (3.10) 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 (3.10) is ergodic on the set E, as defined in (4.4), according to Proposition 4.5. Extending the framework to the case when $N \leq D + 1$ is the subject of ongoing work. We now turn E into a $D \times N$-dimensional Riemannian manifold. Denoting the γth coordinate of the ith particle by $U^{(\gamma,i)}$, we introduce the metric tensor

$$(4.7) \qquad g = \sum_{i=1}^{N} \sum_{\gamma,\sigma=1}^{D} \mathcal{C}_{\gamma\sigma}^{-1} \mathrm{d}U^{(\gamma,i)} \, \mathrm{d}U^{(\sigma,i)}.$$

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 the reader to [15, 21] and, in the context of computational statistics, to [24]. Using these objects induced by g, both the SDE (3.10) and the corresponding Fokker–Planck equation (4.1) admit a compact formulation.

Proposition 4.9 (Riemannian interpretation of ALDI). *Let $\pi^{(N),g}$ denote the density of $\pi^{(N)}$ with respect to the Riemannian volume, that is, $\pi^{(N),g} \mathrm{dvol}_g = \pi^{(N)} \mathrm{d}U$ and $\pi_*^{(N),g} \mathrm{dvol}_g = \pi_*^{(N)} \mathrm{d}U$. Then the dynamics (3.10) can be written in the form*

$$(4.8) \qquad \mathrm{d}U_t = \nabla_U^g \log \pi_*^{(N),g}(U_t) \, \mathrm{d}t + \sqrt{2} \, \mathrm{d}W_t^g,$$

and the Fokker–Planck equation (4.1) can be written in the form

$$(4.9) \qquad \partial_t \pi_t^{(N),g} = \nabla_U^g \cdot \left(\pi_t^{(N),g} \, \nabla_U^g \frac{\delta \mathrm{KL}(\pi_t^{(N)} | \pi_*^{(N)})}{\delta \pi_t^{(N)}} \right).$$

Remark 4.10. 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 [26, 29, 31] $\phi \mapsto -\nabla_U^g \cdot \left(\pi_t^{(N),g} \, \nabla_U^g \phi \right)$.

Proof. Using the results from [24], in particular (46)–(47), the proof of the first statement reduces to verifying that

$$(4.10) \qquad \partial_J g^{IJ} = \frac{D+1}{N} \left(u^{(i)} - m(U) \right)_\gamma,$$

where g^{IJ} stands for the components of the inverse of g, and we have used the notation $I = (\gamma, i)$ and $J = (\sigma, j)$. Furthermore, we apply Einstein's summation convention here and in the remainder of this proof. The statement (4.10) follows directly from the definition of g and the identity [30]

$$(4.11) \qquad \nabla_{u^{(i)}} \cdot \mathcal{C}(U) = \frac{D+1}{N} (u^{(i)} - m(U)),$$

giving rise to the drift correction (3.16). Indeed, together with the coordinate expressions

$$(4.12) \qquad \nabla_U^g \cdot f = \frac{1}{\sqrt{|g|}} \partial_I \left(\sqrt{|g|} f^I \right), \quad (\nabla_U^g V)^I = g^{IJ} \partial_J V$$

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

For the second statement, note that $\mathrm{dvol}_g = \sqrt{|g|}\, \mathrm{d}U$, and hence $\pi^{(N),g} = |g|^{-1/2} \pi^{(N)}$. ∎

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

$$(4.13) \qquad \mathcal{W}_g^2 \left(\mu^{(N)}, \nu^{(N)} \right) = \inf_{\gamma \in \Pi\left(\mu^{(N)}, \nu^{(N)} \right)} \int_{E \times E} d_g^2(U, V) \, \mathrm{d}\gamma(U, V),$$

where $\Pi\left(\mu^{(N)}, \nu^{(N)} \right)$ denotes the set of probability measures on $E \times E$ with marginals $\mu^{(N)}$ and $\nu^{(N)}$. It is well known that the evolution (4.9) can be interpreted as gradient flow dynamics of the Kullback–Leibler divergence on the set of probability measures equipped with the distance (4.13); see, for instance, [42, Chapter 15] or [23]. By the Benamou–Brenier formula [3], we have the representation

$$(4.14a) \quad \mathcal{W}_g^2 \left(\mu^{(N)}, \nu^{(N)} \right) = \inf_{\{\pi_t, \Phi_t\}} \left\{ \int_0^1 \int_E g(\nabla_U^g \Phi_t, \nabla_U^g \Phi_t) \, \mathrm{d}\pi_t \, \mathrm{d}t : \right.$$

$$(4.14b) \qquad \left. \partial_t \pi_t^g + \nabla_U^g \cdot (\pi_t^g \nabla_U^g \Phi_t) = 0, \quad \pi_0 = \mu^{(N)}, \ \pi_1 = \nu^{(N)} \right\},$$

where the constraining continuity equation in (4.14b) is to be interpreted in a weak form and we again denoted by π^g the density of π with respect to dvol_g. In standard coordinates (using the definition (4.7) as well as the formulas (4.12)) we see that

$$(4.15a) \quad \mathcal{W}_g^2 \left(\mu^{(N)}, \nu^{(N)} \right) = \inf_{\{\pi_t, \Phi_t\}} \left\{ \int_0^1 \int_E \nabla_U \Phi_t \cdot \mathcal{C} \, \nabla_U \Phi_t \, \mathrm{d}\pi_t \mathrm{d}t : \right.$$

$$(4.15b) \qquad \left. \partial_t \pi_t + \nabla_U \cdot (\pi_t \mathcal{C} \, \nabla_U \Phi_t) = 0, \quad \pi_0 = \mu^{(N)}, \ \pi_1 = \nu^{(N)} \right\},$$

revealing a close similarity with the Kalman–Wasserstein distance (here denoted by $\mathcal{W}_{\mathrm{Kalman}}$) introduced in [11]. Indeed, let us choose $\mu^{(N)} := \otimes_{i=1}^N \mu^{(i)}$ and $\nu^{(N)} := \otimes_{i=1}^N \nu^{(i)}$, the product measures on $\mathbb{R}^{D \times N}$ associated to $\mu, \nu \in \mathcal{P}(\mathbb{R}^D)$, where $\mu^{(i)}$ and $\nu^{(i)}$, $i = 1, \ldots, N$, are understood to be identical copies of μ and ν. We formally expect that

$$(4.16) \qquad \frac{1}{N} \mathcal{W}_g^2 \left(\mu^{(N)}, \nu^{(N)} \right) \xrightarrow{N \to \infty} \mathcal{W}_{\mathrm{Kalman}}(\mu, \nu),$$

using that $\mathcal{C}(U) \approx C(\pi)$ for sufficiently large N, where $C(\pi)$ was defined in (2.23). 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) was carried out recently in [5].

Remark 4.11 (gradient flow structure of the ensemble Kalman–Bucy filter). Taking the formal mean-field limit of the ensemble Kalman–Bucy filter (3.17) leads to the following evolution equation in the marginal densities π_t:

$$\partial_t \pi_t = \nabla_u \cdot \left(\pi_t \, C(\pi_t) \, \nabla_u \frac{\delta \mathcal{F}_{\mathrm{EnKBF}}(\pi_t)}{\delta \pi_t} \right) \tag{4.17}$$

with potential

$$\mathcal{F}_{\mathrm{EnKBF}}(\pi) = \frac{1}{4} \int_{\mathbb{R}^D} \|y_{\mathrm{obs}} - \mathcal{G}(u)\|_R^2 \, \pi(u) \, \mathrm{d}u + \frac{1}{4} \|y_{\mathrm{obs}} - \mathbb{E}_\pi[\mathcal{G}(u)]\|_R^2, \tag{4.18}$$

which arises naturally from (3.18) in the limit $N \to \infty$ [36, 35]. Note that (4.17) is exactly of the form (2.22) with the Kullback–Leibler divergence replaced by the potential (4.18). Its gradient flow structure in the space of probability measures was first discussed in [36, 35] and is equivalent to the Kalman–Wasserstein gradient flow structure introduced in [11]. The mean-field limit of the EKI [40, 19] also fits within this framework with the potential $\mathcal{F}_{\mathrm{EnKBF}}(\pi)$ replaced by

$$\mathcal{F}_{\mathrm{EKI}}(\pi) = \frac{1}{2} \int_{\mathbb{R}^D} \|y_{\mathrm{obs}} - \mathcal{G}(u)\|_R^2 \, \pi(u) \, \mathrm{d}u. \tag{4.19}$$

The affine invariance of both the EnKBF and EKI follows along the lines of Lemma 2.4. As for the finite ensemble size formulations, one expects a slower decay of $\mathcal{F}_{\mathrm{EnKBF}}(\pi_t)$ compared to $\mathcal{F}_{\mathrm{EKI}}(\pi)$.

5. Numerical experiment: A PDE constrained inverse problem.
We consider the inverse problem of determining the permeability field $a(x) > 0$ in the elliptic PDE

$$-\partial_x(a(x)\partial_x p(x)) = f(x), \qquad x \in \Omega = [0, 2\pi), \tag{5.1}$$

from $K = 10$ observed grid values

$$y_j = p(\mathrm{x}_j) + \eta_j, \qquad \mathrm{x}_j = \frac{2\pi(j-1)}{K}, \tag{5.2}$$

$j = 1, \ldots, K$, of the pressure field p for a given forcing f. Both p and f are assumed to integrate to zero over the domain Ω. The measurement errors η_j in (5.2) are i.i.d. Gaussian with mean zero and variance $\sigma_R = 10^{-4}$. A related 2-dimensional Darcy flow problem was studied in [11]. In this paper, we restrict the simulations to the 1-dimensional formulation (5.1) for computational simplicity.

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

$$x_i = \frac{2\pi i}{D}, \qquad i = 0, \ldots, D-1, \tag{5.3}$$

with $D = 50$ grid points. Hence (5.1) gets replaced by the finite-difference formulation

$$\frac{a_{i+1/2}(p_{i+1} - p_i) - a_{i-1/2}(p_i - p_{i-1})}{h^2} = -f_i, \tag{5.4}$$

$i = 1, \ldots, D$. Here $h = 2\pi/D$ denotes the mesh size and $p_i \approx p(x_i)$, 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 nonnegative, we set

$$(5.5) \qquad a_{i-1/2} = \exp(u_i)$$

for $i = 1, \ldots, D$. The computational forward problem is now given by the solution $\{p_i\}_{i=0}^{D-1}$ to (5.4) for given $\{f_i\}_{i=0}^{D-1}$ and $\{u_i\}_{i=1}^{D}$ and its restriction to the observation grid $\{x_j\}_{j=1}^{K}$. We denote this map by $\mathcal{G}(u)$, suppressing the dependence on the forcing given by

$$(5.6) \qquad f_i = \exp\left(-\frac{(2x_i - L)^2}{40}\right) - c_f,$$

where $c_f > 0$ is chosen such that the forcing has mean zero. The measurement error covariance matrix is given by $R = \sigma_R I_{K \times K}$. This completes the description of our forward model (2.3).

The prior distribution on $u \in \mathbb{R}^D$ is assumed to be Gaussian with mean zero and covariance matrix P_0 defined by

$$(5.7) \qquad P_0^{-1} = 4h \left(\frac{\mu}{D} 1_D 1_D^{\mathrm{T}} - \Delta_h\right)^2,$$

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 (5.4) with $a_{i\pm1/2} = 1$. The parameter $\mu > 0$ is set to $\mu = 10^2$, leading to a penalty on the (spatial) mean of $u = \{u_i\}_{i=1}^{D}$ to be close to zero.

The observations (5.2) are generated numerically by solving (5.4) with the reference permeability field given by

$$(5.8) \qquad a_{i-1/2}^{\dagger} = \exp(u_i^{\dagger}),$$

where

$$(5.9) \qquad u_i^{\dagger} = \frac{1}{2}\sin(x_i - h/2)$$

for $i = 1, \ldots, D$, and setting

$$(5.10) \qquad y_j = p_l + \eta_j, \qquad l = \frac{D}{K} j = 5j, \qquad \eta_j \sim \mathcal{N}(0, \sigma_R),$$

$j = 1, \ldots, K$.

We implemented the gradient-based ALDI formulation (3.10) as well as the gradient-free ALDI formulation (3.14) using the Euler–Maruyama method with step size $\Delta t = 0.01$ over a time interval $t \in [0, 20]$. In line with [11] we refer to the ALDI implemented without the correction term (3.16) as the ensemble Kalman sampler (EKS). The ensemble sizes were taken as $N = 25, 52, 100, 200$. Except for the smallest ensemble size, all other choices resulted in nonsingular empirical covariance matrices $\mathcal{C}(U_t)$.

We compare the simulation results based on the estimation bias

$$(5.11) \qquad \mathrm{BIAS} = \frac{h}{T} \int_{\tau}^{\tau+T} \| m(\exp(U_t)) - a^{\dagger} \|_1 \, \mathrm{d}t,$$

where $\|a\|_1$ denotes the l_1-norm of a vector $a \in \mathbb{R}^D$, and the ensemble spread

$$(5.12) \qquad \mathrm{SPREAD} = \frac{h}{T} \int_{\tau}^{\tau+T} \mathrm{trace}\,(\mathcal{C}(U_t))$$

computed along numerical solutions for $\tau = 12$ and $T = 8$. 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 implementations yield results which are essentially indistinguishable from those based on the exact gradient while being computationally much more efficient. Finally, the results for ALDI indicate that it is entirely sufficient to implement it with $N = D + 2 = 52$ particles, the minimum size required for ergodicity to hold.

Table 1
Computed estimation bias (5.11) for ensemble sizes $N \in \{25, 51, 100, 200\}$ and implementations of ALDI and EKS as well as with exact gradient (g) and gradient-free (gf).

N	gf-EKS	gf-ALDI	g-EKS	g-ALDI
25	0.5035	0.4113	0.4940	0.4042
52	0.3748	0.3028	0.3706	0.2957
100	0.3215	0.3070	0.3166	0.3016
200	0.3088	0.3081	0.3030	0.3009

Table 2
As in Table 1, but reporting the results for the ensemble spread (5.12).

N	gf-EKS	gf-ALDI	g-EKS	g-ALDI
25	0.0082	0.0724	0.0083	0.0738
52	0.0135	0.0475	0.0134	0.0476
100	0.0219	0.0457	0.0218	0.0457
200	0.0337	0.0453	0.0336	0.0453

In order to provide better insight into the impact of the correction term (3.16) 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 (3.16) 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.
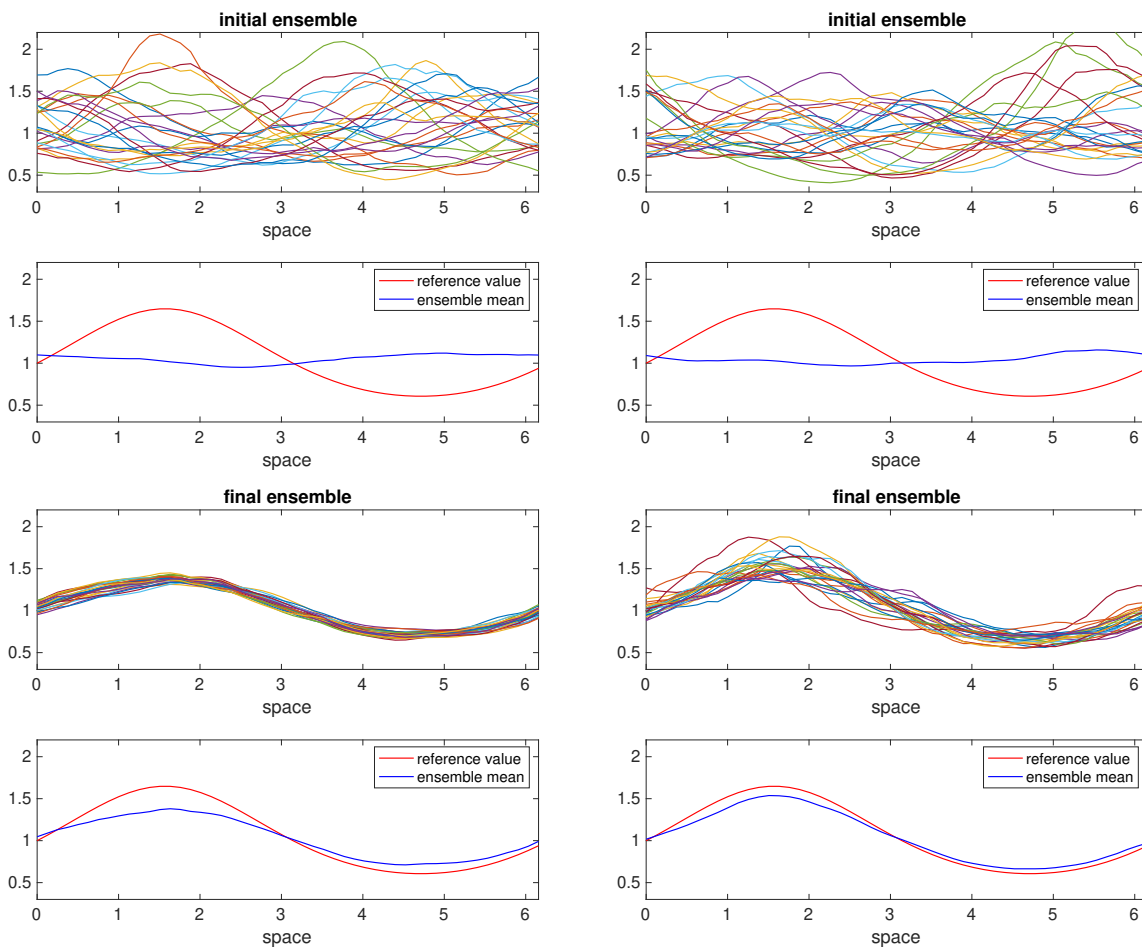
Figure 1. *Displayed are the initial (top row) and final (bottom row) ensembles of the permeability fields $a(x) = \exp(u(x))$ for $N = 25$. The left column is from the EKS, while the right column is from the ALDI method.*

6. 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 (3.16) due to the multiplicative nature of the noise in the Langevin equations (3.10) [30]. 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 nonsingular initial empirical covariance matrix $\mathcal{C}(U_0)$ ensure that $|\mathcal{C}(U_t)| \neq 0$ for all $t \geq 0$ and that the equations of motion (3.10) are ergodic with invariant measure $\pi_*^{(N)}$. Further computational savings can be achieved through the gradient-free implementation (3.14) for BIPs as introduced in Example 2.1. 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
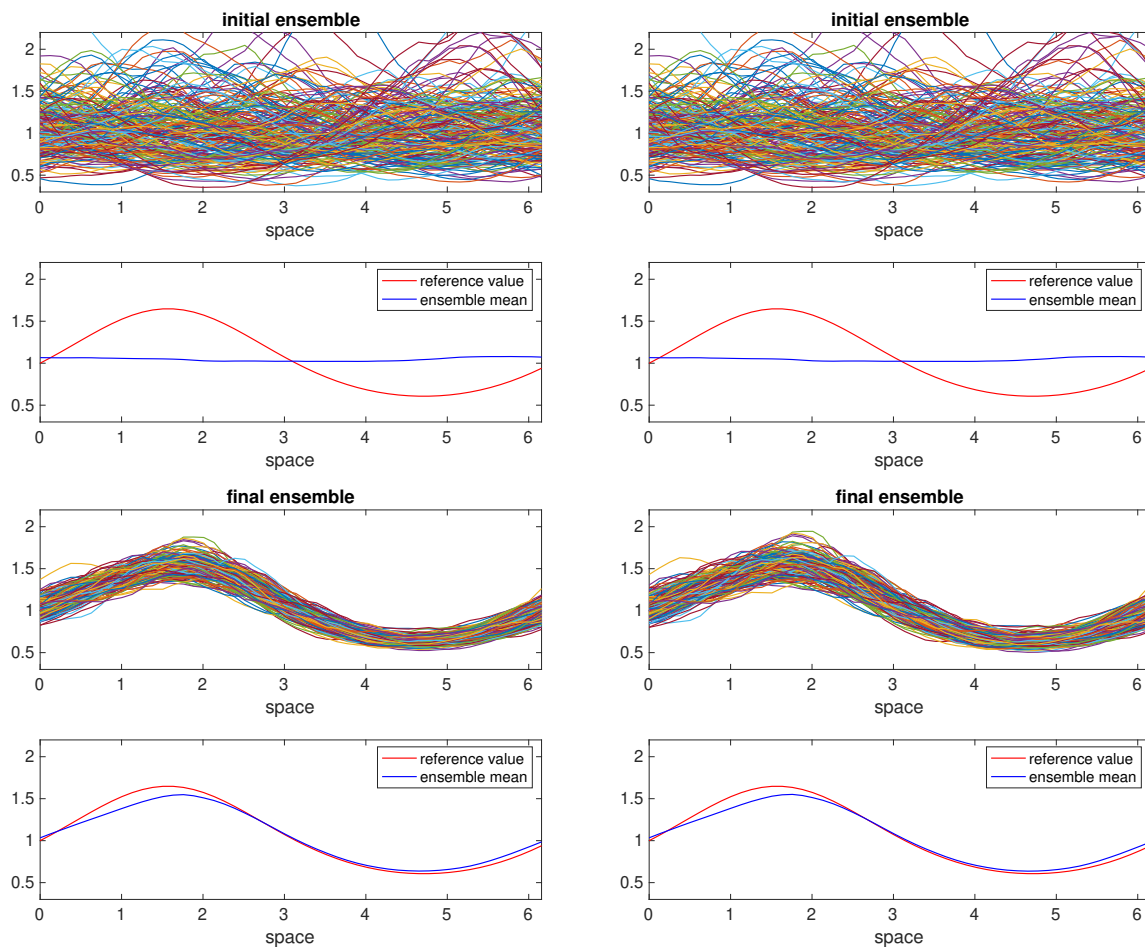
Figure 2. *As in Figure 1, except for ensemble size $N = 200$.*

for reducing estimation errors both for $N < D$ as well as for $N = \mathcal{O}(D)$ implementations of the ALDI method (3.10).

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 was studied for the related continuous-time ensemble Kalman–Bucy filter in [2]. We also reemphasize that multimodal target distributions might require localized empirical covariance matrices in (3.10), as first suggested in [22] and further explored in [37].

While this paper has focused on a theoretical investigation and computational implementation of finite-sample size interacting Langevin dynamics, we wish to point out that the Kalman–Wasserstein gradient flows proposed in [36, 11] have 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 nondegeneracy and ergodicity.

Proof of Proposition 4.1. The Fokker–Planck equation is given by

$$
(A.1) \qquad \partial_t \pi_t^{(N)} = \mathcal{L}^\dagger \pi_t^{(N)},
$$

where $\mathcal{L}$ denotes the infinitesimal generator of (3.10) and $\mathcal{L}^\dagger$ refers to its adjoint in $L^2(\mathbb{R}^{D \times N})$, given by

$$
(A.2a) \qquad \left(\mathcal{L}^\dagger \pi^{(N)}\right)(U) = \sum_{i=1}^N \nabla_{u^{(i)}} \cdot \left(\pi^{(N)}(U)\left\{\mathcal{C}(U)\nabla_{u^{(i)}}\Phi(u^{(i)}) - \frac{D+1}{N}(u^{(i)} - m(U))\right\}\right)
$$

$$
(A.2b) \qquad + \sum_{i=1}^N \nabla_{u^{(i)}} \cdot \left\{\pi^{(N)}(U)\,\nabla_{u^{(i)}} \cdot \mathcal{C}(U) + \mathcal{C}(U)\nabla_{u^{(i)}}\pi^{(N)}(U)\right\};
$$

see [33, Chapter 4] and [22]. Here the divergence of the matrix-valued $\mathcal{C}(U)$ is componentwise given by (in terms of the notation introduced in subsection 4.3)

$$
(A.3) \qquad \{(\nabla_{u^{(i)}} \cdot \mathcal{C}(U))\}_k = \sum_{\gamma=1}^D \frac{\partial}{\partial U^{(\gamma,i)}} \{\mathcal{C}(U)\}_{k\gamma}, \qquad k = 1, \ldots, D.
$$

An explicit calculation [30] leads to (4.11), and the Fokker–Planck operator $\mathcal{L}^\dagger$ reduces to

$$
(A.4a) \qquad \mathcal{L}^\dagger \pi^{(N)} = \sum_{i=1}^N \nabla_{u^{(i)}} \cdot \left(\pi^{(N)}\,\mathcal{C}\left\{\nabla_{u^{(i)}}\Phi + \nabla_{u^{(i)}}\log\pi^{(N)}\right\}\right)
$$

$$
(A.4b) \qquad = \sum_{i=1}^N \nabla_{u^{(i)}} \cdot \left(\pi^{(N)}\,\mathcal{C}\,\nabla_{u^{(i)}}\log\frac{\pi^{(N)}}{\pi_*^{(N)}}\right),
$$

from which the desired result follows since

$$
(A.5) \qquad \frac{\delta \mathrm{KL}(\pi^{(N)}|\pi_*^{(N)})}{\delta \pi^{(N)}} = \log\frac{\pi^{(N)}}{\pi^{(N)*}}. \qquad \blacksquare
$$

For the proof of Proposition 4.4 we recall the definition (4.4) of the set $E \subset \mathbb{R}^{D \times N}$. We will use the potential $\mathcal{V}$ defined in (3.7) as a Lyapunov function. The key calculation is summarized in the following lemma.

Lemma A.1. *There exists a constant $\gamma > 0$ such that*

$$
(A.6) \qquad (\mathcal{L}\mathcal{V})(U) \le \gamma \mathcal{V}(U), \quad U \in E,
$$

where $\mathcal{L}$ is the generator of (3.10), that is, the $L^2(\mathbb{R}^{D \times N})$-adjoint of $\mathcal{L}^\dagger$ as defined in (A.4).

Proof. It follows from (A.4) that the generator of (3.10) takes the form

$$
(A.7) \qquad (\mathcal{L}\mathcal{V})(U) = -\sum_{i=1}^N \nabla_{u^{(i)}}\Phi(u^{(i)}) \cdot \mathcal{C}(U)\nabla_{u^{(i)}}\mathcal{V}(U) + \sum_{i=1}^N \nabla_{u^{(i)}} \cdot (\mathcal{C}(U)\nabla_{u^{(i)}}\mathcal{V}(U)).
$$

For convenience, let us introduce the notation

$$\text{(A.8)} \qquad \mathcal{V}_{\mathcal{C}}(U) = -\frac{D+1}{2} \log |\mathcal{C}(U)|, \qquad \mathcal{V}_{\Phi}(U) = \sum_{i=1}^{N} \Phi(u^{(i)}).$$

Since $\mathcal{L}$ vanishes on constants, (A.6) 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 termwise for the contributions in (A.7).

First note that

$$\text{(A.9)} \qquad \nabla_{u^{(i)}} \log |\mathcal{C}(U)| = \frac{2}{N}\mathcal{C}^{-1}(U)(u^{(i)} - m(U)), \qquad i = 1, \ldots, N.$$

Indeed, again following the notation introduced in subsection 4.3, we have that

$$\text{(A.10a)} \quad \frac{\partial}{\partial U^{(\gamma,i)}} \log |\mathcal{C}(U)| = \left(\frac{1}{|\mathcal{C}|} \sum_{\alpha,\beta=1}^{D} \frac{\partial |\mathcal{C}|}{\partial \mathcal{C}^{\alpha\beta}} \frac{\partial \mathcal{C}^{\alpha\beta}}{\partial U^{(\gamma,i)}} \right)(U)$$

$$\text{(A.10b)} \qquad\qquad = \sum_{\alpha,\beta=1}^{D} (\mathcal{C}^{-1})_{\alpha\beta} \left(\frac{1}{N}\delta_{\alpha\gamma}(u^{(i)} - m(U))_{\beta} + \frac{1}{N}\delta_{\beta\gamma}(u^{(i)} - m(U))_{\alpha} \right)$$

$$\text{(A.10c)} \qquad\qquad = \frac{2}{N} \left(\mathcal{C}^{-1}(u^{(i)} - m(U)) \right)_{\gamma},$$

using Jacobi's formula for determinants in the second line. For the first term in (A.7) we thus obtain

$$\text{(A.11a)} \quad -\sum_{i=1}^{N} \nabla_{u^{(i)}} \Phi(u^{(i)}) \cdot \mathcal{C}(U) \nabla_{u^{(i)}} \mathcal{V}(U)$$

$$\text{(A.11b)} \qquad = -\frac{D+1}{N} \sum_{i=1}^{N} \nabla_{u^{(i)}} \Phi(u^{(i)}) \cdot (u^{(i)} - m(U)) \underbrace{- \sum_{i=1}^{N} \nabla_{u^{(i)}} \Phi(u^{(i)}) \cdot \mathcal{C}(U) \nabla_{u^{(i)}} \Phi(u^{(i)})}_{\leq 0}$$

$$\text{(A.11c)} \qquad \leq \tilde{C} \left(1 + \mathcal{V}_{\Phi}(U) \right).$$

To bound the second term in (A.7), we first notice the estimate

$$\text{(A.12)} \qquad \sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot \left(\mathcal{C}(U) \nabla_{u^{(i)}} \Phi(u^{(i)}) \right) \leq \tilde{C} \left(1 + \mathcal{V}(U) \right),$$

again easily obtained from Assumption 4.3. The other contribution is

$$\text{(A.13a)} \qquad \sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot (\mathcal{C}(U) \nabla_{u^{(i)}} \mathcal{V}_{\mathcal{C}}(U)) = -\frac{D+1}{N} \sum_{i=1}^{N} \nabla_{u^{(i)}} \cdot \left(u^{(i)} - m(U) \right)$$

$$\text{(A.13b)} \qquad\qquad\qquad\qquad\qquad = -\frac{(D+1)D}{2N} \left(N - 1 \right).$$

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 (A.11) and (A.12), the claim follows. ∎

Proposition 4.4 now essentially follows from adapting [28, 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 the following self-contained proof.

Proof of Proposition 4.4. The potential $\mathcal{V}$ is bounded from below by the growth condition on Φ (see Assumption 4.3). We can therefore choose a constant $c_\mathcal{V}$ such that $\mathcal{V}_+ := \mathcal{V} + c_\mathcal{V}$ is nonnegative. Since $\mathcal{C}(U_0)$ is assumed to be nondegenerate, there exists $k_0 \in \mathbb{N}$ such that $\mathcal{V}_+(U_0) < k_0$. For $k \geq k_0$, let us define the sets

$$(\text{A.14}) \qquad E_k = \{U \in E : \quad \mathcal{V}_+(U) < k\}$$

and the stopping times

$$(\text{A.15}) \qquad \tau_k = \inf\{t \geq 0 \ : U_t \notin E_k\} = \inf\{t \geq 0 : \mathcal{V}_+(U_t) = k\}.$$

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

$$(\text{A.16}) \qquad g(U,t) := e^{-\gamma t}\mathcal{V}_+(U), \quad (U,t) \in E \times [0,\infty),$$

where γ is the constant obtained in Lemma A.1. By using Itô's formula, optional stopping, and the bound (A.6) we see that

$$(\text{A.17a}) \qquad \mathbb{E}\left[g(U_{t\wedge\tau_k}, t \wedge \tau_k)\right] = g(U_0,0) + \mathbb{E}\left[\int_0^{t\wedge\tau_k} e^{-\gamma s}\left(-\gamma\mathcal{V}_+(U_s) + \mathcal{L}\mathcal{V}_+(U_s)\right)\mathrm{d}s\right]$$

$$(\text{A.17b}) \qquad \leq g(U_0,0) = \mathcal{V}_+(U_0)$$

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

$$(\text{A.18a}) \qquad \mathbb{E}\left[g(U_{t\wedge\tau_k}, t \wedge \tau_k)\right] \geq e^{-\gamma t}\mathbb{E}\left[\mathcal{V}_+(U_{t\wedge\tau_k})\right]$$

$$(\text{A.18b}) \qquad \geq e^{-\gamma t}\left(\mathbb{E}\left[\mathbf{1}_{t<\tau_k}\mathcal{V}_+(U_t)\right] + \mathbb{P}\left[\tau_k \leq t\right]\cdot k\right) \geq e^{-\gamma t}\left(\mathbb{P}\left[\tau_k \leq t\right]\cdot k\right),$$

where the last estimate uses the fact that $\mathcal{V}_+ \geq 0$. Combining (A.17) and (A.18), we see that

$$(\text{A.19}) \qquad e^{\gamma t}\mathcal{V}_+(U_0) \geq \mathbb{P}\left[\tau_k \leq t\right]\cdot k$$

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

$$(\text{A.20}) \qquad \lim_{k\to\infty}\mathbb{P}\left[\tau_k \leq t\right] = 0,$$

and further

$$(\text{A.21}) \qquad \mathbb{P}\left[\xi \leq t\right] = 0$$

by monotone convergence. Since (A.21) holds for all $t \geq 0$, we conclude that $\mathbb{P}[\xi = \infty] = 1$, as required. ∎

For the proof of Proposition 4.5 we will need the following lemma.

Lemma A.2. *Let $N \geq D + 2$. 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,

$$(A.22) \quad m^{-j}(U) = \frac{1}{N-1} \sum_{\substack{i=1 \\ i \neq j}}^{N} u^{(i)}, \qquad \mathcal{C}^{-j}(U) = \frac{1}{N-1} \sum_{\substack{i=1 \\ i \neq j}}^{N} (u^{(i)} - m^{-j}(U))(u^{(i)} - m^{-j}(U))^{\mathrm{T}}.$$

Notice the update formula

$$(A.23) \quad \mathcal{C}(U) = \frac{N-1}{N} \mathcal{C}^{-j}(U) + \frac{N-1}{N^2}(u^{(j)} - m^{-j}(U))(u^{(j)} - m^{-j}(U))^{\mathrm{T}},$$

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

$$(A.24) \quad \widetilde{E} := \left\{ U \in \mathbb{R}^{D \times N} : \quad \mathcal{C}^{-j} \text{ is invertible for all } j \in \{1, \ldots, N\} \right\}.$$

We see that $\widetilde{E} \subset E$ since the second term on the right-hand side of (A.23) is positive semidefinite. Importantly, the condition $N \geq D + 2$ ensures that $\widetilde{E}$ is nonempty.

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

$$(A.25) \quad \widetilde{E} = \left\{ U \in \mathbb{R}^{D \times N} : \quad \prod_{j=1}^{N} |\mathcal{C}^{-j}(U)| > 0 \right\},$$

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

$$(A.26) \quad P(t) = \prod_{j=1}^{N} |\mathcal{C}^{-j}((1-t)X + tY)|,$$

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

We now show how to construct a continuous path between arbitrary $X, Y \in E$. 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 neighborhoods of X and Y, respectively. Since $\widetilde{E}$ is open, there exist open neighborhoods $U_{\widetilde{X}}, U_{\widetilde{Y}} \subset \widetilde{E}$. It is then sufficient to find points in these neighborhoods 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 (A.23), 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 of $\overline{\gamma}^{(i)}$ can be chosen arbitrarily small in order to ensure that the final point $\overline{\gamma}^{(N)}(1)$ belongs to $U_{\widetilde{Y}}$. ∎

Proof of Proposition 4.5. Since the diffusion matrix $\Gamma(U)\Gamma(U)^{\mathrm{T}}$ is strictly positive definite on E, E is path-connected by Lemma A.2, and the process $(U_t)_{t\geq 0}$ admits an invariant measure with strictly positive Lebesgue density by Corollary 4.2, the process is positively recurrent and irreducible by the result in [18]. We also refer the reader to [41, section 2.2.2.1]. The convergence in total variation distance then follows from [27, Theorem 6.1]. ∎

REFERENCES

[1] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*, Birkhäuser Verlag, Basel, 2008.

[2] J. Amezcua, E. Kalnay, K. Ide, and S. Reich, *Ensemble transform Kalman-Bucy filters*, Quart. J. Roy. Meteor. Soc., 140 (2014), pp. 995–1004.

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

[4] K. Bergemann and S. Reich, *An ensemble Kalman-Bucy filter for continuous data assimilation*, Meteorolog. Z., 21 (2012), pp. 213–219.

[5] J. Carrillo, M. Delgadino, and G. Pavliotis, *A Proof of the Mean-Field Limit for λ-Convex Potentials by Γ-Convergence*, preprint, https://arxiv.org/abs/1906.04601, 2019.

[6] J. Carrillo and U. Vaes, *Wasserstein Stability Estimates for Covariance-Preconditioned Fokker–Planck Equations*, preprint, https://arxiv.org/abs/1910.07555, 2019.

[7] Z. Ding and Q. Li, *Ensemble Kalman Sampling: Mean-Field Limit and Convergence Analysis*, preprint, https://arxiv.org/abs/1910.12923, 2019.

[8] A. Duncan, N. Nüsken, and G. Pavliotis, *Using perturbed underdamped Langevin dynamics to efficiently sample from probability distributions*, J. Stat. Phys., 169 (2017), pp. 1098–1131.

[9] A. Eberle, *Markov Processes*, lecture notes, 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*, SIAM J. Appl. Dyn. Syst., 19 (2020), pp. 412–441, https://doi.org/10.1137/19M1251655.

[12] M. Girolami and B. Calderhead, *Riemann manifold Langevin and Hamiltonian Monte Carlo methods*, J. R. Stat. Soc. Ser. B Stat. Methodol., 73 (2011), pp. 123–214.

[13] J. Goodman and J. Weare, *Ensemble samplers with affine invariance*, Commun. Appl. Math. Comput. Sci., 5 (2010), pp. 65–80.

[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*, Grad. Stud. Math. 38, American Mathematical Society, Providence, RI, 2002.

[16] R. Jordan, D. Kinderlehrer, and F. Otto, *The variational formulation of the Fokker–Planck equation*, SIAM J. Math. Anal., 29 (1998), pp. 1–17, https://doi.org/10.1137/S0036141096303359.

[17] R. Khasminskii, *Stochastic Stability of Differential Equations*, Springer Science & Business Media, Heidelberg, 2012.

[18] W. Kliemann, *Recurrence and invariant measures for degenerate diffusions*, Ann. Probab., 15 (1987), pp. 690–707.

[19] N. B. Kovachki and A. M. Stuart, *Ensemble Kalman inversion: A derivative-free technique for machine learning tasks*, Inverse Problems, 35 (2019), 095005.

[20] K. Law, A. Stuart, and K. Zygalakis, *Data Assimilation: A Mathematical Introduction*, Springer-Verlag, New York, 2015.

[21] J. LEE, *Riemannian Manifolds: An Introduction to Curvature*, Springer Science & Business Media, New York, 2006.

[22] B. LEIMKUHLER, C. MATTHEWS, AND J. WEARE, *Ensemble preconditioning for Markov chain Monte Carlo simulations*, Stat. Comput., 28 (2018), pp. 277–290.

[23] S. LISINI, *Nonlinear diffusion equations with variable coefficients as gradient flows in Wasserstein spaces*, ESAIM Control Optim. Calc. Var., 15 (2009), pp. 712–740.

[24] S. LIVINGSTONE AND M. GIROLAMI, *Information-geometric Markov chain Monte Carlo methods using diffusions*, Entropy, 16 (2014), pp. 3074–3102.

[25] Y.-A. MA, T. CHEN, AND E. FOX, *A complete recipe for stochastic gradient MCMC*, in Proceedings of the 28th International Conference on Neural Information Processing Systems, Volume 2, MIT Press, Cambridge, MA, 2015, pp. 2917–2925.

[26] S. MACHLUP AND L. ONSAGER, *Fluctuations and irreversible process.* II. *Systems with kinetic energy*, Phys. Rev., 91 (1953), pp. 1512–1515.

[27] S. MEYN AND R. TWEEDIE, *Stability of Markovian processes* II: *Continuous-time processes and sampled chains*, Adv. Appl. Probab., 25 (1993), pp. 487–517.

[28] S. MEYN AND R. TWEEDIE, *Stability of Markovian processes* III: *Foster–Lyapunov criteria for continuous-time processes*, Adv. Appl. Probab., 25 (1993), pp. 518–548.

[29] A. MIELKE, D. RENGER, AND M. PELETIER, *A generalization of Onsager's reciprocity relations to gradient flows with nonlinear mobility*, J. Non-Equilibrium Thermodyn., 41 (2016), pp. 141–149.

[30] N. NÜSKEN AND S. REICH, *Note on Interacting Langevin Diffusion: Gradient Structure and Ensemble Kalman Sampler*, preprint, https://arxiv.org/pdf/1908.10890.pdf, 2019.

[31] H. ÖTTINGER, *Beyond Equilibrium Thermodynamics*, John Wiley & Sons, 2005.

[32] S. PATHIRAJA AND S. REICH, *Discrete gradients for computational Bayesian inference*, J. Comput. Dyn., 6 (2019), pp. 236–251.

[33] G. PAVLIOTIS, *Stochastic Processes and Applications*, Springer-Verlag, New York, 2014.

[34] S. REICH, *A dynamical systems framework for intermittent data assimilation*, BIT, 51 (2011), pp. 235–249.

[35] S. REICH AND C. COTTER, *Probabilistic Forecasting and Bayesian Data Assimilation*, Cambridge University Press, Cambridge, UK, 2015.

[36] S. REICH AND C. J. COTTER, *Ensemble filter techniques for intermittent data assimilation*, in Large Scale Inverse Problems: Computational Methods and Applications in the Earth Sciences, M. Cullen, M. A. Freitag, S. Kindermann, and R. Scheichl, eds., Radon Ser. Comput. Appl. Math. 13, Walter de Gruyter, Berlin, 2013, pp. 91–134.

[37] S. REICH AND S. WEISSMANN, *Fokker–Planck Particle Systems for Bayesian Inference: Computational Approaches*, preprint, https://arxiv.org/abs/1911.10832, 2019.

[38] G. ROBERTS AND O. STRAMER, *Langevin diffusions and Metropolis–Hastings algorithms*, Methodol. Comput. Appl. Probab., 4 (2003), pp. 337–358.

[39] P. ROSSKY, J. DOLL, AND H. FRIEDMANN, *Brownian dynamics as smart Monte Carlo simulation*, J. Chem. Phys., 69 (1978), pp. 4628–4633.

[40] C. SCHILLINGS AND A. M. STUART, *Analysis of the ensemble Kalman filter for inverse problems*, SIAM J. Numer. Anal., 55 (2017), pp. 1264–1290, https://doi.org/10.1137/16M105959X.

[41] G. STOLTZ AND M. ROUSSET, *Free Energy Computations: A Mathematical Perspective*, World Scientific, Hackensack, NJ, 2010.

[42] C. VILLANI, *Optimal Transport: Old and New*, Grundlehren Math. Wiss. 338, Springer Science & Business Media, New York, 2008.