

SAMPLING VIA GRADIENT FLOWS IN THE SPACE OF PROBABILITY MEASURES

YIFAN CHEN, DANIEL ZHENG YU HUANG, JIAOYANG HUANG, SEBASTIAN REICH,
AND ANDREW M. STUART

ABSTRACT. Sampling a target probability distribution with an unknown normalization constant is a fundamental challenge in computational science and engineering. Recent work shows that algorithms derived by considering gradient flows in the space of probability measures open up new avenues for algorithm developments. This paper makes three contributions to this approach to sampling, by scrutinizing the design components of such gradient flows. Any instantiation of a gradient flow for sampling needs an energy functional and a metric to determine the flow, as well as numerical approximations of the flow to derive algorithms. Our first contribution is to show that the Kullback-Leibler (KL) divergence, as an energy functional, has the *unique* property (among all f -divergences) that gradient flows resulting from it do not depend on the normalization constant of the target distribution; this justifies the widespread use of the KL divergence in sampling. Our second contribution is to study the choice of metric from the perspective of invariance. The Fisher-Rao metric is known as the *unique* choice (up to scaling) that is diffeomorphism invariant. As a computationally tractable alternative, we introduce a relaxed, affine invariance property for the metrics and gradient flows. In particular, we construct various affine invariant Wasserstein and Stein gradient flows. Affine invariant gradient flows are shown to behave more favorably than their non-affine-invariant counterparts when sampling highly anisotropic distributions, in theory and by using particle methods. Our third contribution is to study, and develop efficient algorithms based on Gaussian approximations of the gradient flows; this leads to an alternative to particle methods. We establish connections between various Gaussian approximate gradient flows, discuss their relation to gradient methods arising from parametric variational inference, and study their convergence properties. Our theory and numerical experiments demonstrate the strengths and potential limitations of the

Received by the editor March 8, 2024, and, in revised form, April 14, 2025, and October 19, 2025.

2020 *Mathematics Subject Classification.* Primary 68Q25, 65D18, 65D15.

Key words and phrases. Bayesian inference, sampling, gradient flow, mean-field dynamics, Gaussian approximation, variational inference, affine invariance.

The first author was supported by the Air Force Office of Scientific Research under MURI award number FA9550-20-1-0358 (Machine Learning and Physics-Based Modeling and Simulation) under the aegis of Houman Owhadi and AMS. The second and fifth authors were supported by NSF award AGS1835860 and by the generosity of Eric and Wendy Schmidt by recommendation of the Schmidt Futures program; the second author was also supported by the National Natural Science Foundation of China through grant 12471403, the Fundamental Research Funds for the Central Universities, and the high-performance computing platform of Peking University; the fifth author was also supported by the Office of Naval Research (ONR) through grant N00014-17-1-2079 and by a Department of Defense Vannevar Bush Faculty Fellowship. The fourth author was supported by Deutsche Forschungsgemeinschaft (DFG) - Project-ID 318763901 - SFB1294. The third author was supported by NSF grant DMS-2054835.

The second author is the corresponding author.

Gaussian approximate Fisher-Rao gradient flow, which is affine invariant, by considering a wide range of target distributions.

1. INTRODUCTION

1.1. The sampling problem. In this paper, we are concerned with the problem of sampling a probability distribution that is known up to normalization. This problem is fundamental in many applications arising in computational science and engineering and is widely studied in applied mathematics, machine learning and statistics communities. A particular application is Bayesian inference for large-scale inverse problems; such problems are ubiquitous, for example in climate science [66, 68, 87, 117], engineering [17, 38, 136], and machine learning [25, 30, 101, 111]. These applications have fueled the need for efficient and scalable sampling algorithms.

Mathematically, the objective is to sample a target probability distribution with density $\rho_{\text{post}}(\cdot)$, for the parameter $\theta \in \mathbb{R}^{N_\theta}$, given by

$$(1.1) \quad \rho_{\text{post}}(\theta) \propto \exp(-\Phi_R(\theta)),$$

where $\Phi_R : \mathbb{R}^{N_\theta} \rightarrow \mathbb{R}$ is a known function. We use the notation ρ_{post} because of the potential application to Bayesian inference, where Φ_R represents the regularized negative log likelihood function; however, we do not explicitly use the Bayesian structure in this paper, and our discussion applies to general target distributions known only up to normalization.

1.2. Gradient flow methodology. Numerous approaches to the sampling problem have been proposed in the literature. Most are based on construction of a dynamical system for densities that converges to the target distribution, or its approximation, after a specified finite time or at infinite time. The most common examples that are widely used in Bayesian inference are sequential Monte Carlo (SMC, specified finite time) [47] and Markov chain Monte Carlo (MCMC, infinite time) [15].

Among them, the Langevin diffusion [106] and its discretization constitute an important class of MCMC algorithms for sampling. It has been shown in the seminal work [70] that the Fokker-Planck equation describing the evolution of densities of the Langevin diffusion is the Wasserstein gradient flow of the Kullback–Leibler (KL) divergence. That is, the dynamical system of densities corresponding to the Langevin diffusion has a gradient flow structure. Such structure has also been identified in other popular sampling algorithms, for example, the Stein variational gradient descent [84], with the Stein variational gradient flow [83] as the continuous limit.

Indeed, extending beyond the above-mentioned examples, gradient flows have profoundly influenced our understanding and development of sampling algorithms. Various gradient flows have been adopted by researchers to address the sampling problems. For instance, the Wasserstein-Fisher-Rao gradient flow was used to sample multi-modal distributions in [90, 91]. In [57, 58], the Kalman-Wasserstein metric (introduced, but not named, in [112]) was used to induce the gradient flow leading to advantageous algorithms for sampling anisotropic distributions. Interpolation between the Wasserstein and Stein metrics was studied in [63]. Accelerated gradient flows in the probability space have been studied in [132]. A recent overview of the use of gradient flows in optimization and sampling can be found in [59, 60]. In

addition to the continuous-time picture, the optimization perspective on gradient flows also leads to new developments in discrete-time algorithms for sampling [134].

Given the many choices of gradient flows in the literature, the aim of this paper is to study their design ingredients and identify key properties that make the gradient flows favorable in the sampling context. Our main focus is the continuous-time formulation of the flow.

1.3. Design ingredients of gradient flows. Given an energy functional $\mathcal{E}$ on the probability space, a class of gradient flows of $\mathcal{E}$ can be formally written as

$$(1.2) \quad \frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}.$$

Here, $\frac{\delta \mathcal{E}}{\delta \rho}$ represents the first variation of $\mathcal{E}$, which is then evaluated at $\rho = \rho_t$. Positive definite operator $M(\rho)$ leads to non-linear preconditioning by its inverse; such operators $M(\rho)$ arise naturally from a Riemannian metric defined in the probability density space.

A key property of the gradient flow is

$$(1.3) \quad \frac{d}{dt} \mathcal{E}(\rho_t) = \left\langle \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}, \frac{\partial \rho_t}{\partial t} \right\rangle = - \left\langle M(\rho_t) \frac{\partial \rho_t}{\partial t}, \frac{\partial \rho_t}{\partial t} \right\rangle \leq 0.$$

This demonstrates that the gradient flow eq. (1.2) will keep decreasing the energy functional. Indeed, if $\mathcal{E}$ is chosen so that ρ_{post} is its unique stationary point, then property (1.3) can be used as the basis of proofs to establish convergence of ρ_t to ρ_{post} as $t \rightarrow \infty$. For more detail concerning the notation we refer to Section 2.

By choosing different $\mathcal{E}(\cdot)$ and $M(\cdot)$, one obtains various gradient flows with varying convergence properties and levels of numerical implementation difficulties. In the sampling context, it is natural to ask which choices lead to the most favorable algorithms, for example, in terms of numerical implementation and convergence rates on specific problem classes. Such questions are the focus of this paper.

1.4. Our contributions and paper organization. The primary contributions of this work are as follows.

- We show in Theorem 3.1 that the KL divergence stands out as a unique energy functional $\mathcal{E}(\cdot)$ among all f -divergences with respect to the target. Specifically, it emerges as the sole choice (up to scaling) that yields gradient flows independent of the normalization constant of the target distribution. Since handling the unknown normalization constant poses a significant challenge in the sampling problem, our finding establishes KL divergence as the desired energy functional for designing gradient flows, ensuring that concerns about normalization constants are eliminated during numerical implementation. We then focus on this choice of energy functional $\mathcal{E}(\cdot)$ in the remainder of the paper.
- We highlight invariance properties stemming from the choice of metric (and hence preconditioner in (1.2)) and elucidate implications for the convergence of the gradient flow. We prove in Theorem 4.1 that the Fisher-Rao metric, which is the unique (up to scaling) metric exhibiting diffeomorphism invariance, achieves a uniform exponential rate of convergence to the target distribution, under quite general conditions. Moreover, we introduce a relaxed (weaker), affine invariance property for gradient flows. As a concrete manifestation, we construct various affine invariant Wasserstein

and Stein gradient flows. These affine invariant gradient flows are more convenient to approximate numerically via particle methods than the diffeomorphism invariant Fisher-Rao gradient flow; pre-existing theory, and numerical experiments in this paper, demonstrate that these affine invariant methods behave favorably compared to their non-affine-invariant versions when sampling highly anisotropic distributions.

- We study efficient implementable sampling algorithms found by restricting the gradient flows to Gaussian measures, as alternatives to particle methods. In Theorem 5.1 we demonstrate that, under mild assumptions, the Gaussian approximation achieved through metric based projection is equivalent to the approximation derived through moment closure; the latter is more convenient for calculations. Furthermore, we establish connections between various Gaussian approximate gradient flows, discuss their relation to gradient methods to solve the Gaussian variational inference problem, and study their convergence properties. Our theory and numerical experiments demonstrate the strengths and limitations of the Fisher-Rao gradient flow under Gaussian approximations, for Gaussian, logconcave, and general target distributions.

The paper is structured as follows. Section 2 introduces gradient flows in the space of probability densities. In Section 3, we discuss the selection of energy functionals, defined over the probability density space and parameterized by the target distribution, showing the primacy of the KL divergence within all f -divergences. In Section 4, we discuss the choice of metrics in the density space and examine the induced gradient flow, with a particular focus on invariance properties such as diffeomorphism invariance and affine invariance. In Section 5, we study Gaussian approximations to efficiently approximate the gradient flows and showcase the advantage and limitation of affine invariance in this context. In Section 6, we demonstrate the effectiveness of the aforementioned approaches on a PDE-constrained Bayesian inverse problem, complementing the simple illustrative numerical examples presented in Section 4 and Section 5. Concluding remarks are presented in Section 7. The appendix contains details of all the proofs.

1.5. Literature survey. We survey the relevant literature in gradient flows, affine invariance, and Gaussian approximations below.

1.5.1. Gradient flows. There is a vast literature concerning the use of gradient flows in probability density space, employing a variety of different metric tensors, to minimize an energy defined as the Kullback–Leibler (KL) divergence between the current density and the target distribution. The most relevant to this paper are the Wasserstein, Fisher-Rao, and Stein gradient flows.

The Wasserstein gradient flow was identified in the seminal work [70]. The authors showed that the Fokker-Planck equation is the Wasserstein gradient flow of the KL divergence of the current density estimate from the target. Since then, Wasserstein gradient flow has played a significant role in optimal transport [115], sampling [28, 78], machine learning [34, 114], partial differential equations [20, 105] and many other areas. The Fisher-Rao metric was introduced by C. R. Rao [110] via the Fisher information matrix. The original definition is in parametric density spaces, and the corresponding Fisher-Rao gradient flow in the parameter space leads to natural gradient descent [1]. The Fisher-Rao metric in infinite dimensional

probability spaces was discussed in [55, 123]. The concept underpins information geometry [2, 6]. The gradient flow of the KL divergence under the Fisher-Rao metric is induced by a mean-field model of birth-death type. The birth-death process has been used in sequential Monte Carlo samplers to reduce the variance of particle weights [43] and to accelerate Langevin sampling [90, 91]. The discovery of the Stein metric [83] follows the introduction of the Stein variational gradient descent algorithm [84]. The study of the Stein gradient flow [48, 83, 89] sheds light on the analysis and improvements of the algorithm [44, 85, 131, 132].

1.5.2. Affine invariance. The idea of affine invariance was introduced for MCMC methods in [53, 61], motivated by the empirical success of the Nelder-Mead simplex algorithm [102] in optimization. Sampling methods with the affine invariance property can be effective for highly anisotropic distributions; this is because they behave identically in all coordinate systems related through an affine transformation; in particular, they can be understood by studying the best possible coordinate system, which reduces anisotropy to the maximum extent possible within the class of affine transformations. The numerical studies presented in [61] demonstrate that affine-invariant MCMC methods offer significant performance improvements over standard MCMC methods. This idea has been further developed to enhance sampling algorithms in more general contexts. Preconditioning strategies for the Langevin dynamics to achieve affine-invariance were discussed in [79]. In [112], the Kalman-Wasserstein metric was introduced, and gradient flows with respect to this metric were advocated and shown to achieve affine invariance [57, 58]. Moreover, the authors in [57, 58, 107] used the empirical covariance as preconditioners in ensemble Kalman based particle approximations of the Kalman-Wasserstein gradient flow, leading to a family of derivative-free affine invariant sampling approaches. Similarly, the work [86] employed the empirical covariance to precondition second order Langevin dynamics. Affine invariant samplers can also be combined with the pCN (preconditioned Crank–Nicolson) MCMC method [36], to boost the performance of MCMC in function space [37, 49]. Another family of affine-invariant sampling algorithms is based on Newton or Gauss-Newton methods, since the use of the Hessian matrix as the preconditioner naturally induces the affine invariance property. Such methods include the stochastic Newton MCMC [97], the Newton flow with different metrics [44, 131], and mirror Langevin diffusion [33, 64, 140].

1.5.3. Gaussian approximation. Parametric approximations such as those made by Gaussians have been widely used in variational inference [13, 69, 127]. These methods, in the Gaussian setting, aim to solve the problem

$$(1.4) \quad (m^*, C^*) = \arg \min_{m, C} \text{KL}[\mathcal{N}(m, C) \parallel \rho_{\text{post}}].$$

Gradient flows in the space of (m, C) can be used to identify (m^*, C^*) . These gradient flows can also be understood as a Gaussian approximation of the gradient flow in the density space. The Riemannian structure in the density space can be projected to the space of Gaussians, leading to the concept of natural gradient [1, 96, 139] that has the advantage of reparametrization invariance. Other work on the use of Gaussian variational inference methods includes the papers [56, 74, 81, 109, 137].

In addition to their role in parametric variational inference, Gaussian approximations have been widely deployed in various generalizations of Kalman filtering [52, 72, 73, 121, 128]. For Bayesian inverse problems, iterative ensemble Kalman

samplers have been proposed [29, 51, 128]. The paper [65] introduced an ensemble Kalman methodology based on a novel mean-field dynamical system that depends on its own filtering distribution. For all these algorithms based on a Gaussian ansatz, the accuracy depends on some measure of being close to Gaussian. Regarding the use of Gaussian approximations in Kalman inversion we highlight, in addition to the approximate Bayesian methods already cited, the use of ensemble Kalman methods for optimization: see [22, 66, 67, 75, 133]. Kalman filtering has also been used in combination with variational inference [76]. The relation between iterative Kalman filtering and Gauss-Newton or Levenberg Marquardt algorithms was studied in [10, 11, 23, 66], and leads to ensemble Kalman based optimization methods which are affine invariant.

1.6. Notation. We use $\#$ to denote the pushforward operation for general measures. It is defined via duality. More precisely consider a diffeomorphism $\varphi : \mathbb{R}^{N_\theta} \rightarrow \mathbb{R}^{N_\theta}$ and probability measures μ, ν in $\mathbb{R}^{N_\theta}$. Then $\nu = \varphi\#\mu$ if and only if

$$\int f(\theta) d\nu = \int f(\varphi(\theta)) d\mu,$$

for any integrable f under measure ν . If μ, ν admit smooth densities ρ and $\tilde{\rho}$ respectively, we can use the change-of-variable formula to derive the identity $\tilde{\rho}(\tilde{\theta}) = \rho(\varphi^{-1}(\tilde{\theta})) |\nabla_{\tilde{\theta}} \varphi^{-1}(\tilde{\theta})|$.

2. GRADIENT FLOWS

In this section, we introduce the general formulation of gradient flows in probability density space. Methodologically, to define gradient flows, one needs a differential structure in the density space, which then leads to the definition of tangent spaces and metric tensors that determine a gradient flow. Whilst they may be rigorously defined in specific contexts,¹ the technical hurdles must be addressed on a case-by-case basis; we seek to keep such technicalities to a minimum and focus on a formal methodology for deriving the flow equation. Once the flow is formally identified, we can rigorously study its large-time convergence properties and numerical approximations of the flow, which are the key factors in the understanding and design of efficient sampling algorithms.

For the formal methodology, we consider the probability space with smooth positive densities:

$$(2.1) \quad \mathcal{P} = \left\{ \rho \in C^\infty(\mathbb{R}^{N_\theta}) : \int \rho d\theta = 1, \rho > 0 \right\},$$

with the tangent space

$$(2.2) \quad T_\rho \mathcal{P} \subseteq \left\{ \sigma \in C^\infty(\mathbb{R}^{N_\theta}) : \int \sigma d\theta = 0 \right\}.$$

We assume that both the target ρ_{post} , and density ρ given by the gradient flow, are in $\mathcal{P}$: $\rho, \rho_{\text{post}} \in \mathcal{P}$. This allows us to use differential structures under the smooth topology to calculate² gradient flow equations.

¹The rigorous theory of gradient flows in suitable infinite-dimensional functional spaces and its link with evolutionary PDEs is a long-standing subject; see [3, 4] for discussions and a rigorous treatment of gradient flows in metric space.

²The formal Riemannian geometric calculations in the density space were first proposed by Otto in [105]. The calculations in the smooth setting are rigorous if $\mathbb{R}^{N_\theta}$ is replaced by a compact

Denote an energy functional by $\mathcal{E}(\cdot; \rho_{\text{post}}) : \mathcal{P} \rightarrow \mathbb{R}$, where we explicitly highlight its dependence on ρ_{post} and assume that ρ_{post} is the minimizer of $\mathcal{E}(\cdot; \rho_{\text{post}})$ over $\mathcal{P}$. Then, a gradient flow of $\mathcal{E}$ in the probability density space has the form given in (1.2). In the following we use the notation that $\langle \cdot, \cdot \rangle$ is the duality pairing between $T_\rho^* \mathcal{P}$ and $T_\rho \mathcal{P}$; it can be identified as the L^2 inner product when both arguments are classical functions. Note that functions in $T_\rho^* \mathcal{P}$ are not uniquely defined under the L^2 inner product, since $\langle \psi, \sigma \rangle = \langle \psi + c, \sigma \rangle$ for all $\sigma \in T_\rho \mathcal{P}$ and any constant c . A unique representation can be identified by requiring, for example, that $\psi \in T_\rho^* \mathcal{P}$ implies $\mathbb{E}_\rho[\psi] = 0$.³ In the context of this paper, the most important elements in $T_\rho^* \mathcal{P}$ are the first variation $\frac{\delta \mathcal{E}}{\delta \rho}$ of energy functionals $\mathcal{E}(\cdot, \rho_{\text{post}})$, defined by

$$(2.3) \quad \left\langle \frac{\delta \mathcal{E}}{\delta \rho}, \sigma \right\rangle = \lim_{\epsilon \rightarrow 0} \frac{\mathcal{E}(\rho + \epsilon \sigma; \rho_{\text{post}}) - \mathcal{E}(\rho, \rho_{\text{post}})}{\epsilon},$$

for any $\sigma \in T_\rho \mathcal{P}$, assuming the limit exists.

The metric tensor $M(\rho) : T_\rho \mathcal{P} \rightarrow T_\rho^* \mathcal{P}$ can be derived from a Riemannian metric g that one imposes in the probability density space; we have $g_\rho(\sigma_1, \sigma_2) = \langle M(\rho)\sigma_1, \sigma_2 \rangle$ for $\sigma_1, \sigma_2 \in T_\rho \mathcal{P}$, and $M(\rho)^{-1} : T_\rho^* \mathcal{P} \rightarrow T_\rho \mathcal{P}$ is sometimes referred to as the Onsager operator [99, 103, 104]. For simplicity of understanding, we view $M(\rho_t)^{-1}$ in the gradient flow eq. (1.2) as a preconditioning operator for the density dynamics. For more details of the Riemannian perspective we refer to [26, 88]. We note that the gradient flow eq. (1.2) will converge to ρ_{post} if it is the unique stationary point of $\mathcal{E}$. Thus numerical simulations of eq. (1.2) lead to sampling algorithms.

A standard use of particle methods to simulate dynamics of densities is to identify a mean-field stochastic dynamical system, with state space $\mathbb{R}^{N_\theta}$, defined so that its law is given by eq. (1.2). For example, we may introduce the Itô SDE

$$(2.4) \quad d\theta_t = f(\theta_t; \rho_t, \rho_{\text{post}})dt + h(\theta_t; \rho_t, \rho_{\text{post}})dW_t,$$

where $W_t \in \mathbb{R}^{N_\theta}$ is a standard Brownian motion. Because the drift $f : \mathbb{R}^{N_\theta} \times \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}^{N_\theta}$ and diffusion coefficient $h : \mathbb{R}^{N_\theta} \times \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}^{N_\theta \times N_\theta}$ are evaluated at ρ_t , the density of θ_t itself, by definition this is a mean-field model. In fact, this particular mean field model is a McKean SDE. Other types of mean-field dynamics, such as birth-death dynamics, and ODEs, also exist.

When utilizing eq. (1.2) for sampling purposes, two primary properties of the flow become our focal points. Firstly, rapid convergence of the flow to ρ_{post} is desired. Secondly, the numerical implementation of the flow must be tractable, and preferably, as convenient as possible, for example through a mean-field model such as (2.4). Below we describe an example.

Example 2.1. Choose $\mathcal{E}$ to be the KL divergence:

$$(2.5) \quad \mathcal{E}(\rho; \rho_{\text{post}}) = \text{KL}[\rho \| \rho_{\text{post}}] = \int \rho \log\left(\frac{\rho}{\rho_{\text{post}}}\right) d\theta,$$

manifold, as noted in [88]; in such case in the definition of tangent spaces, (2.2) also becomes identity. For rigorous results in general probability space, we refer to [3].

³This choice is also naturally motivated by the definition of Fisher-Rao gradient flows; see (4.1). That is, for $\psi \in T_\rho^* \mathcal{P}$ implying $\mathbb{E}_\rho[\psi] = 0$, the Fisher-Rao metric tensor reduces to a multiplication by the density ρ .

and use the Wasserstein metric tensor [105] which satisfies $M(\rho)^{-1}\psi = -\nabla_\theta \cdot (\rho \nabla_\theta \psi)$, then the gradient flow eq. (1.2) takes the form:

$$(2.6) \quad \frac{\partial \rho_t}{\partial t} = -\nabla_\theta \cdot (\rho_t \nabla_\theta \log \rho_{\text{post}}) + \nabla_\theta \cdot (\nabla_\theta \rho_t).$$

This is the Fokker-Planck equation of the Langevin dynamics [70]:

$$(2.7) \quad d\theta_t = \nabla_\theta \log \rho_{\text{post}}(\theta_t) dt + \sqrt{2} dW_t.$$

This is a trivial (in the sense that it is not actually a mean-field model) example of (2.4). Discretizations of this Langevin dynamics can be used to numerically implement the flow. Regarding the convergence property of eq. (2.6), we have the following classical result [95]:

Proposition 2.2. *Assume that ρ_{post} is α -strongly logconcave: $\log \rho_{\text{post}} \in C^2(\mathbb{R}^{N_\theta})$ and*

$$(2.8) \quad -\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta) \succeq \alpha I.$$

Then, for all $t \geq 0$, it holds that

$$(2.9) \quad \text{KL}[\rho_t \|\rho_{\text{post}}] \leq \text{KL}[\rho_0 \|\rho_{\text{post}}] e^{-2\alpha t}.$$

The convergence rate depends crucially on ρ_{post} . If $\rho_{\text{post}} = \mathcal{N}(0, I)$, the unit Gaussian distribution, the convergence rate is e^{-2t} . However, if the Gaussian distribution is highly anisotropic, then α depends on the lower bound of the eigenvalues of the precision matrix, which can be very small, leading to a slow convergence rate. Related convergence results with relaxed assumptions on the target distribution can be found in [7, 8, 62, 108, 126], based on functional inequalities, and in [98] using coupling methods. There are also many results [31, 39, 50, 118, 125, 135] regarding the non-asymptotic convergence guarantee for discrete algorithms based on the Langevin dynamics.

Note that for the choice of energy functionals in Example 2.1, it is not necessary to know the normalization constant for ρ_{post} ; this is highly desirable. On the other hand, for the Wasserstein metric Example 2.1 shows that the convergence rate depends sensitively on the target; this is undesirable. In Sections 3 and 4 we discuss, within a broader context, choices of the energy functionals and metric tensors. We identify certain choices that can facilitate ease of numerical implementations and fast convergence rates across large classes of targets.

3. CHOICE OF ENERGY FUNCTIONALS

In the literature, the KL divergence is the most commonly used energy functional for deriving gradient flows in the probability space. Using the definition eq. (2.5) of the KL divergence and definition eq. (2.3) of the first variation, we can formally deduce that

$$(3.1) \quad \frac{\delta \mathcal{E}}{\delta \rho} = \log \rho - \log \rho_{\text{post}} - \mathbb{E}_\rho[\log \rho - \log \rho_{\text{post}}] \in T_\rho^* \mathcal{P}.$$

From eq. (3.1) we observe that, for the KL divergence, $\frac{\delta \mathcal{E}}{\delta \rho}$ remains unchanged if we scale ρ_{post} by any positive constant $c > 0$, i.e. if we change ρ_{post} to $c\rho_{\text{post}}$. This property eliminates the need to know the normalization constant of ρ_{post} in order to calculate the first variation. As a consequence, the gradient flow eq. (1.2) derived from this energy functional is also independent of the normalization constant. This

has significant advantages in terms of numerical implementations, as handling the unknown normalization constant is a well-known challenge in sampling.

We can also formulate the property in terms of the energy functional itself, without delving into the technicality of dealing with the mathematical well-definedness of first variations.⁴ An equivalent formulation is that $\mathcal{E}(\rho; c\rho_{\text{post}}) - \mathcal{E}(\rho; \rho_{\text{post}})$ is independent of ρ , for any $c \in (0, \infty)$. In fact, this implies that

$$\left\langle \frac{\delta \mathcal{E}}{\delta \rho}(\rho; c\rho_{\text{post}}) - \frac{\delta \mathcal{E}}{\delta \rho}(\rho; \rho_{\text{post}}), \sigma \right\rangle = 0$$

for $\sigma \in T_\rho \mathcal{P}$ when the first variations exist; thus in such case, first variations do not depend on the scaling of ρ_{post} .

Can this property be satisfied for other energy functionals? In Theorem 3.1, we show that the answer is no among all f -divergences with continuously differentiable f . Here the f -divergence between two continuous density functions ρ and ρ_{post} , positive everywhere, is defined as

$$D_f[\rho \parallel \rho_{\text{post}}] = \int \rho_{\text{post}} f\left(\frac{\rho}{\rho_{\text{post}}}\right) d\theta.$$

For convex f with $f(1) = 0$, Jensen's inequality implies that $D_f[\rho \parallel \rho_{\text{post}}] \geq 0$. The KL divergence used in (2.5) corresponds to the choice $f(x) = x \log x$.

Theorem 3.1. *Assume that $f : (0, \infty) \rightarrow \mathbb{R}$ is continuously differentiable and $f(1) = 0$. Then the KL divergence is the only f -divergence (up to scalar factors) such that $D_f[\rho \parallel c\rho_{\text{post}}] - D_f[\rho \parallel \rho_{\text{post}}]$ is independent of $\rho \in \mathcal{P}$ for any $c \in (0, \infty)$ and for any $\rho_{\text{post}} \in \mathcal{P}$.*

The proof can be found in Appendix A.1. As a consequence of Theorem 3.1, gradient flows defined by the energy (2.5) do not depend on the normalization constant of ρ_{post} . Hence numerical approximations of gradient flows stemming from this energy $\mathcal{E}$ are more straightforward to implement in comparison with the use of other divergences or distances as energy. This justifies the choice of the KL divergence as an energy functional for sampling, and our developments in most of this paper are hence specific to the energy eq. (2.5).

Remark 3.2. Other energy functionals can be, and are, used for constructing gradient flows; for example, the chi-squared divergence [32, 82]:

$$(3.2) \quad \chi^2(\rho \parallel \rho_{\text{post}}) = \int \rho_{\text{post}} \left(\frac{\rho}{\rho_{\text{post}}} - 1 \right)^2 d\theta = \int \frac{\rho^2}{\rho_{\text{post}}} d\theta - 1.$$

The normalization constant can appear explicitly in the gradient flow equation when general energy functionals are used. Additional techniques need to be explored to simulate such flows. For example, when the energy functional is the chi-squared divergence, in [32], kernelization is used to avoid the normalization constant in the Wasserstein gradient flow. Moreover, in [82] where a modification of the Fisher-Rao metric is used, ensemble methods with birth-death type dynamics are adopted to derive numerical methods; the normalization constant can be absorbed into the birth-death rate. These techniques may change the flow equation completely or its time scales, thus affecting its convergence behavior.

⁴For the f -divergences considered in this paper, we can show the limit in eq. (2.3) exists so the first variation is well-defined. For details see the proof of Theorem 3.1 in Appendix A.1 using a localization argument.

4. CHOICE OF METRICS

We now fix the energy functional to be the KL divergence and discuss the choices of $M(\cdot)$. The metric tensor can significantly influence the rate at which the gradient flow converges. Within this section, we delve into a range of metrics, focusing on *invariance* properties; as we will show in this and the next section, these play an important role in determining convergence rates of the flow and algorithms based on it. In Section 4.1, we discuss the Fisher-Rao gradient flow, highlight its diffeomorphism invariance property, and showcase the remarkable uniform convergence rate to the target distribution which stems from this invariance; it is also known that the Fisher-Rao metric is the unique diffeomorphism invariant metric (up to scaling) in the probability density space. However, implementing the Fisher-Rao gradient flow via particle methods is resource-demanding and does not currently lead to efficient algorithms (see discussion below). In Section 4.2, we introduce a relaxed, affine invariance property for gradient flows. It is worth noting that certain widely used gradient flows, namely the Wasserstein and Stein gradient flows, are not affine invariant. Some modifications of the Wasserstein and Stein gradient flows to achieve affine invariance are introduced in Section 4.3 and Section 4.4. These affine invariant gradient flows are more convenient to approximate numerically via particle methods compared to the diffeomorphism invariant Fisher-Rao gradient flow. In Section 4.5, we conduct illustrative numerical experiments showing that these affine invariant gradient flows behave more favorably compared to their non-affine-invariant versions when sampling highly anisotropic distributions.

4.1. Fisher-Rao gradient flow and diffeomorphism invariance. The Fisher-Rao Riemannian metric⁵ [2, 6] is

$$g_{\rho}^{\text{FR}}(\sigma_1, \sigma_2) = \int \frac{\sigma_1 \sigma_2}{\rho} d\theta, \text{ for } \sigma_1, \sigma_2 \in T_{\rho}\mathcal{P}.$$

The metric tensor then admits the following form:

$$(4.1) \quad M^{\text{FR}}(\rho)^{-1}\psi = \rho\psi \in T_{\rho}\mathcal{P}$$

for any $\psi \in T_{\rho}^*\mathcal{P}$; recall that $\psi \in T_{\rho}^*\mathcal{P}$ implies $\mathbb{E}_{\rho}[\psi] = 0$. Combining eq. (4.1) and eq. (3.1), we obtain the Fisher-Rao gradient flow of the KL divergence as

$$(4.2) \quad \begin{aligned} \frac{\partial \rho_t}{\partial t} &= -M^{\text{FR}}(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t}, \\ &= -\rho_t \left((\log \rho_t - \log \rho_{\text{post}}) - \mathbb{E}_{\rho_t}[\log \rho_t - \log \rho_{\text{post}}] \right). \end{aligned}$$

An important property of the Fisher-Rao gradient flow is its *diffeomorphism invariance*. More precisely, consider a diffeomorphism in the parameter space $\varphi : \mathbb{R}^{N_{\theta}} \rightarrow \mathbb{R}^{N_{\theta}}$ and correspondingly $\tilde{\rho}_t = \varphi \# \rho_t$, $\tilde{\rho}_{\text{post}} = \varphi \# \rho_{\text{post}}$. Then, it holds that

$$(4.3) \quad \frac{\partial \tilde{\rho}_t}{\partial t} = -\tilde{\rho}_t \left((\log \tilde{\rho}_t - \log \tilde{\rho}_{\text{post}}) - \mathbb{E}_{\tilde{\rho}_t}[\log \tilde{\rho}_t - \log \tilde{\rho}_{\text{post}}] \right).$$

That is, the form of the equation remains invariant under diffeomorphic transformations. As a consequence, suppose there is a diffeomorphism satisfying $\tilde{\rho}_{\text{post}} = \varphi \# \rho_{\text{post}} = \mathcal{N}(0, I)$, then we can study the convergence rate of eq. (4.3) for a

⁵For some intuitions for the definition of the Fisher-Rao Riemannian metric, see Appendix C.1.

Gaussian target distribution and the rate will apply directly to eq. (4.2), due to the following property of the KL divergence:

$$(4.4) \quad \text{KL}[\rho_t \|\rho_{\text{post}}] = \text{KL}[\varphi \# \rho_t \|\varphi \# \rho_{\text{post}}] = \text{KL}[\tilde{\rho}_t \|\tilde{\rho}_{\text{post}}].$$

This property, which also holds for other f -divergences beyond the KL divergence, implies that the Fisher-Rao gradient flow may converge at the same rate for general distributions and Gaussian distributions, indicating its favorable convergence behavior. Indeed, in Theorem 4.1 we prove its uniform convergence rate across a wide class of target distributions. The proof can be found in Appendix B.1; we note that the diffeomorphism invariance property is not explicitly used in the current proof.

Theorem 4.1. *Assume that there exist constants $K, B > 0$ such that the initial density ρ_0 satisfies*

$$(4.5) \quad e^{-K(1+|\theta|^2)} \leq \frac{\rho_0(\theta)}{\rho_{\text{post}}(\theta)} \leq e^{K(1+|\theta|^2)},$$

and both $\rho_0, \rho_{\text{post}}$ have bounded second moments:

$$(4.6) \quad \int |\theta|^2 \rho_0(\theta) d\theta \leq B, \quad \int |\theta|^2 \rho_{\text{post}}(\theta) d\theta \leq B.$$

Let ρ_t solve the Fisher-Rao gradient flow (4.2). Then, for any $t \geq \log((1+B)K)$,

$$(4.7) \quad \text{KL}[\rho_t \|\rho_{\text{post}}] \leq (2+B+eB)Ke^{-t}.$$

Similar propositions are presented in concurrent works [90, Theorem 3.3] and [91, Theorem 2.3]. Our results do not require $\rho_0(\theta)/\rho_{\text{post}}(\theta)$ to be bounded, thus applying to a broader class of distributions including Gaussians. An explicit expansion of the KL divergence along the Fisher-Rao gradient flow is also given in [46], which indicates the optimal asymptotic convergence rate is $\mathcal{O}(e^{-2t})$. Additionally, detailed discussions on uniform convergence rates with respect to other f -divergences can be found in the concurrent works [19, 100].

As mentioned earlier, the diffeomorphism invariance of the Fisher-Rao gradient flow provides an intuitive explanation as to why it achieves such an exceptional uniform convergence. This is in sharp contrast to the Wasserstein gradient flow (Proposition 2.2). Nevertheless, effectively approximating the Fisher-Rao gradient flow is not straightforward, in particular because naïve particle methods lead to algorithms which do not evolve the support of the measure, requiring enormous resources therefore. In [90, 91], birth-death dynamics are used to simulate the flow, using kernel density estimators to approximate the current density and adding transport steps to move the support of particles; the latter in fact changes the flow equation. The diffeomorphism invariance property is also not preserved at the particle level.

On the other hand, the invariance of the gradient flow is due to the invariance of the Riemannian metric itself. The Fisher-Rao metric is known to be the only metric, up to constants, that satisfies the diffeomorphism invariance property [5, 9, 21]. Therefore there are no alternatives if we ask for such strong diffeomorphism invariance. In the following subsection, we explore a relaxed, affine invariance property for the gradient flows. In contrast to the diffeomorphism invariance property, we are able to construct various affine invariant gradient flows, in particular, the affine invariant Wasserstein and Stein gradient flows. Moreover, when using particle

methods to simulate the flow, we can show that the affine invariance property is also preserved at the particle level.

4.2. Affine invariance. Roughly speaking, affine invariant gradient flows are invariant under any invertible *affine transformations* of the density variables; as a consequence, the convergence rate is independent of the affine transformation. It is thus natural to expect that algorithms with this property have an advantage for sampling highly anisotropic target distributions. We note that the affine invariance property has been studied in MCMC algorithms [53, 61] and certain gradient flows [57, 58]; see the literature survey in Section 1.5.2. Our goal here is to bring this notion to general gradient flows in the density space.

4.2.1. Affine invariant gradient flow. In the following, we define affine invariant gradient flows.

Definition 4.2 (Affine invariant gradient flow). Consider the gradient flow

$$(4.8) \quad \frac{\partial \rho_t}{\partial t} = -M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho}(\rho; \rho_{\text{post}}) \Big|_{\rho=\rho_t},$$

and the affine transformation $\tilde{\theta} = \varphi(\theta) = A\theta + b$, with $A \in \mathbb{R}^{N_\theta \times N_\theta}$ invertible and $b \in \mathbb{R}^{N_\theta}$. Let $\tilde{\rho}_t := \varphi \# \rho_t$ and $\tilde{\rho}_{\text{post}} = \varphi \# \rho_{\text{post}}$. The *gradient flow is affine invariant* if

$$(4.9) \quad \frac{\partial \tilde{\rho}_t}{\partial t} = -M(\tilde{\rho}_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho}(\tilde{\rho}; \tilde{\rho}_{\text{post}}) \Big|_{\tilde{\rho}=\tilde{\rho}_t},$$

for any invertible affine transformation φ .

The key idea in Definition 4.2 is that, after the affine transformation, the dynamics of $\tilde{\rho}_t$ is itself a gradient flow, in the same metric as the gradient flow in the original variables. Thus affine invariance is similar to diffeomorphism invariance, but with the diffeomorphism constrained to invertible affine mappings. In Sections 4.3 and 4.4, we will construct concrete examples such as affine invariant Wasserstein and Stein gradient flows that satisfy the affine invariance.

4.2.2. Affine invariant mean-field dynamics. As mentioned in the Fisher-Rao gradient flow, another important factor when using the flow for sampling is its numerical implementation, for example via particle methods. In the following we introduce the affine invariance property at the particle level. More specifically we consider the use of mean-field Itô SDEs of the form (2.4) and study its affine invariance property. The density of solution to (2.4) is governed by a non-linear Fokker-Planck equation

$$(4.10) \quad \frac{\partial \rho_t}{\partial t} = -\nabla_\theta \cdot (\rho_t f) + \frac{1}{2} \nabla_\theta \cdot (\nabla_\theta \cdot (h h^T \rho_t)).$$

By choice of f, h it may be possible to ensure that eq. (4.10) coincides with eq. (1.2). Then an interacting particle system can be used to approximate eq. (2.4), generating an empirical measure which approximates ρ_t .

Definition 4.3 (Affine invariant mean-field dynamics). Consider the mean-field dynamics eq. (2.4) and the affine transformation $\tilde{\theta} = \varphi(\theta) = A\theta + b$, with $A \in \mathbb{R}^{N_\theta \times N_\theta}$ invertible and $b \in \mathbb{R}^{N_\theta}$. The mean-field dynamics is called affine invariant, if

$$(4.11a) \quad Af(\theta; \rho, \rho_{\text{post}}) = f(\varphi(\theta); \varphi\#\rho, \varphi\#\rho_{\text{post}}),$$

$$(4.11b) \quad Ah(\theta; \rho, \rho_{\text{post}}) = h(\varphi(\theta); \varphi\#\rho, \varphi\#\rho_{\text{post}}),$$

for any such invertible affine transformation φ . Equivalently, this implies that $\tilde{\theta}_t = \varphi(\theta_t)$ satisfies a SDE of the same form as eq. (2.4):

$$(4.12) \quad d\tilde{\theta}_t = f(\tilde{\theta}_t; \tilde{\rho}_t, \tilde{\rho}_{\text{post}})dt + h(\tilde{\theta}_t; \tilde{\rho}_t, \tilde{\rho}_{\text{post}})dW_t,$$

where $\tilde{\rho}_t = \varphi\#\rho_t$ and $\tilde{\rho}_{\text{post}} = \varphi\#\rho_{\text{post}}$.

Again, the key idea in Definition 4.3 is that, after the affine transformation, the mean field dynamics eq. (2.4) remains the same. We will present affine invariant mean-field dynamics for the affine invariant gradient flows we constructed.

4.3. Affine invariant Wasserstein gradient flow. The Wasserstein Riemannian metric tensor⁶ is defined as (see Example 2.1)

$$(4.13) \quad M^W(\rho)^{-1}\psi = -\nabla_\theta \cdot (\rho \nabla_\theta \psi) \in T_\rho \mathcal{P}.$$

Recall the form of the Wasserstein gradient flow of the KL divergence given in Example 2.1, along with the (trivial) mean-field Langevin dynamics. The Wasserstein gradient flow is not affine invariant. Indeed, from Proposition 2.2, we observe that its convergence rate is not invariant upon affine transformations of the target distribution. To make it affine invariant we consider the following generalized Wasserstein metric tensor

$$(4.14) \quad M^{\text{AIW}}(\rho)^{-1}\psi = -\nabla_\theta \cdot (\rho P(\theta, \rho) \nabla_\theta \psi) \in T_\rho \mathcal{P},$$

where $P : \mathbb{R}^{N_\theta} \times \mathcal{P} \rightarrow \mathbb{R}_{>0}^{N_\theta \times N_\theta}$ is a preconditioner satisfying Condition 4.4:

Condition 4.4. Consider any invertible affine transformation $\tilde{\theta} = \varphi(\theta) = A\theta + b$ with $A \in \mathbb{R}^{N_\theta \times N_\theta}$ invertible and $b \in \mathbb{R}^{N_\theta}$ and correspondingly $\tilde{\rho} = \varphi\#\rho$. The preconditioning matrix satisfies

$$(4.15) \quad P(\tilde{\theta}, \tilde{\rho}) = AP(\theta, \rho)A^T.$$

Remark 4.5. Examples of preconditioning matrices that satisfy eq. (4.15) include the covariance matrix $P(\theta, \rho) = C(\rho)$ and some local preconditioners, such as

$$P(\theta, \rho) = (\theta - m(\rho))(\theta - m(\rho))^T,$$

or more generally,

$$P(\theta, \rho) = \int (\theta' - m(\rho))(\theta' - m(\rho))^T \kappa(\theta, \theta', \rho) \rho(\theta') d\theta'.$$

Here $\kappa : \mathbb{R}^{N_\theta} \times \mathbb{R}^{N_\theta} \times \mathcal{P} \rightarrow \mathbb{R}$ is assumed to be a positive definite kernel for any fixed $\rho \in \mathcal{P}$ which is, moreover, affine invariant, namely $\kappa(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}) = \kappa(\theta, \theta', \rho)$ under any

⁶For completeness, we include some intuitions regarding the definition of the Wasserstein Riemannian metric in Appendix C.2.

invertible affine transformation $\tilde{\theta} = \varphi(\theta) = A\theta + b$ and correspondingly $\tilde{\rho} = \varphi\#\rho$. A potential choice is

$$\kappa(\theta, \theta', \rho) = \exp\left(-\frac{1}{2}(\theta - \theta')^T C(\rho)^{-1}(\theta - \theta')\right).$$

When such condition is satisfied, the gradient flow is affine invariant:

Theorem 4.6. *Under the assumption on P given in Condition 4.4, the associated gradient flow of the KL divergence with respect to the metric tensor M^{AIW} , namely*

$$(4.16) \quad \frac{\partial \rho_t(\theta)}{\partial t} = \nabla_{\theta} \cdot \left(\rho_t P(\theta, \rho_t) (\nabla_{\theta} \log \rho_t - \nabla_{\theta} \log \rho_{\text{post}}) \right),$$

is affine invariant.

The proof of Theorem 4.6 is provided in Appendix D.2. Henceforth we refer to M^{AIW} satisfying the condition of Theorem 4.6 as an affine invariant Wasserstein metric tensor. For the specific choice of the preconditioner $P(\theta, \rho) = C(\rho)$, the covariance, the flow is called the Kalman-Wasserstein gradient flow [57, 58]; the underlying metric structure was first identified in [112]. When ρ_{post} belongs to the Gaussian family, the Kalman-Wasserstein gradient flow provably achieves convergence rate $\mathcal{O}(e^{-t})$ in KL divergence [18, 57, 58]. It is worth mentioning that there will be an additional computational cost associated with evaluating $P(\theta, \rho)$. However, for many practical problems, the dominant computational expense lies in evaluating Φ_R and its derivatives, and the additional computational cost is negligible in comparison.

Define $h : \mathbb{R}^{N_{\theta}} \times \mathcal{P} \rightarrow \mathbb{R}^{N_{\theta} \times N_{\theta}}$. Once defined we may then also define $D(\theta, \rho) = \frac{1}{2} h(\theta, \rho) h(\theta, \rho)^T$ and $d(\theta, \rho) = \nabla_{\theta} \cdot D(\theta, \rho)$. Now consider the following mean field dynamics:

$$(4.17) \quad \begin{aligned} d\theta_t &= P(\theta_t, \rho_t) \nabla_{\theta} \log \rho_{\text{post}}(\theta_t) dt \\ &+ \left((D(\theta_t, \rho_t) - P(\theta_t, \rho_t)) \nabla_{\theta} \log \rho_t(\theta_t) + d(\theta_t, \rho_t) \right) dt + h(\theta_t, \rho_t) dW_t. \end{aligned}$$

For any choice of h the Fokker-Planck equation associated with this mean-field dynamics is the affine invariant Wasserstein gradient flow (4.16); see a proof in Appendix D.3. When $D(\theta, \rho_t) \neq P(\theta_t, \rho_t)$, the equation requires knowledge of the score function $\nabla_{\theta} \log \rho_t(\theta_t)$ of the current density, which is computationally undesirable. However, for this purpose, various approaches have been adopted in the literature [14, 94, 119, 130]; see also [120] and references therein for discussions of the score estimation.

At the particle level, we can prove the above dynamics is affine invariant; see Theorem 4.7 and its proof in Appendix D.4.

Theorem 4.7. *The mean-field dynamics eq. (4.17) is affine invariant under the assumption on the preconditioner P given in Condition 4.4 and the assumptions on h given in eq. (4.11b).*

In particular, let $C(\rho)$ denote the covariance matrix of ρ . If we take $P(\theta, \rho) = C(\rho)$ and $h(\theta, \rho) = \sqrt{2C(\rho)}$, then we get the following affine invariant overdamped Langevin equation, introduced in [57, 58]:

$$(4.18) \quad d\theta_t = C(\rho_t) \nabla_{\theta} \log \rho_{\text{post}}(\theta_t) dt + \sqrt{2C(\rho_t)} dW_t.$$

Comparison with eq. (2.7) demonstrates that it is a preconditioned version of the standard overdamped Langevin dynamics (2.6).

4.4. Affine invariant Stein gradient flow. The Stein Riemannian metric tensor⁷ is defined as

$$(4.19) \quad M^S(\rho)^{-1}\psi = -\nabla_\theta \cdot \left(\rho(\theta) \int \kappa(\theta, \theta', \rho) \rho(\theta') \nabla_{\theta'} \psi(\theta') d\theta' \right) \in T_\rho \mathcal{P}.$$

Here $\kappa : \mathbb{R}^{N_\theta} \times \mathbb{R}^{N_\theta} \times \mathcal{P} \rightarrow \mathbb{R}$ is a positive definite kernel for any fixed ρ . The Stein gradient flow has the form

$$(4.20) \quad \begin{aligned} \frac{\partial \rho_t(\theta)}{\partial t} &= - \left(M^S(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t} \right) (\theta) \\ &= \nabla_\theta \cdot \left(\rho_t(\theta) \int \kappa(\theta, \theta', \rho_t) \rho_t(\theta') \nabla_{\theta'} (\log \rho_t(\theta') - \log \rho_{\text{post}}(\theta')) d\theta' \right), \end{aligned}$$

and it has the mean-field counterpart [83, 84] in θ_t with the law ρ_t :

$$(4.21) \quad \begin{aligned} \frac{d\theta_t}{dt} &= \int \kappa(\theta_t, \theta', \rho_t) \rho_t(\theta') \nabla_{\theta'} (\log \rho_{\text{post}}(\theta') - \log \rho_t(\theta')) d\theta' \\ &= \int \kappa(\theta_t, \theta', \rho_t) \rho_t(\theta') \nabla_{\theta'} \log \rho_{\text{post}}(\theta') + \rho_t(\theta') \nabla_{\theta'} \kappa(\theta_t, \theta', \rho_t) d\theta'. \end{aligned}$$

Here, the second equality is obtained using integration by parts; it facilitates an expression that avoids the score (gradient of the log density function of ρ_t). This is useful because, when implementing particle methods, the resulting integral can then be approximated directly by Monte Carlo methods.

The Stein gradient flow is in general not affine invariant. To make it affine invariant, we introduce the following metric tensor

$$(4.22) \quad M^{\text{AIS}}(\rho)^{-1}\psi = -\nabla_\theta \cdot \left(\rho(\theta) \int \kappa(\theta, \theta', \rho) \rho(\theta') P(\theta, \theta', \rho) \nabla_{\theta'} \psi(\theta') d\theta' \right) \in T_\rho \mathcal{P}.$$

Here, the preconditioner $P : \mathbb{R}^{N_\theta} \times \mathbb{R}^{N_\theta} \times \mathcal{P} \rightarrow \mathbb{R}^{N_\theta \times N_\theta}$ is assumed to be positive definite, which may take the form of $P(\theta, \theta', \rho) = L(\theta, \rho) L(\theta', \rho)^T$. The kernel function and the preconditioner are required to satisfy Condition 4.8:

Condition 4.8. Consider any invertible affine transformation $\tilde{\theta} = \varphi(\theta) = A\theta + b$ with $A \in \mathbb{R}^{N_\theta \times N_\theta}$ invertible and $b \in \mathbb{R}^{N_\theta}$ and correspondingly $\tilde{\rho} = \varphi \# \rho$; moreover $\tilde{\theta}' = \varphi(\theta')$. The preconditioning matrix and the kernel satisfy

$$\kappa(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}) P(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}) = \kappa(\theta, \theta', \rho) A P(\theta, \theta', \rho) A^T.$$

Remark 4.9. Examples satisfying Condition 4.8 can be obtained using an affine invariant kernel that fulfills the condition $\kappa(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}) = \kappa(\theta, \theta', \rho)$, along with preconditioning matrices that satisfy eq. (4.15).

⁷For completeness, we include some intuitions in the definition of the Stein Riemannian metric in Appendix C.3.

Then, the corresponding gradient flow becomes affine invariant:

Theorem 4.10. *Under the assumption on P and κ given in Condition 4.8. The associated gradient flow of the KL divergence with respect to the metric tensor M^{AIS} , namely*

$$(4.23) \quad \begin{aligned} \frac{\partial \rho_t(\theta)}{\partial t} &= \nabla_{\theta} \cdot (\mathbf{f}), \\ \mathbf{f} &= \left(\rho_t(\theta) \int \kappa(\theta, \theta', \rho_t) \rho_t(\theta') P(\theta, \theta', \rho_t) \nabla_{\theta'} (\log \rho_t(\theta') - \log \rho_{\text{post}}(\theta')) d\theta' \right), \end{aligned}$$

is affine invariant.

The proof of Theorem 4.10 can be found in Appendix D.5. Furthermore, we can construct the following mean-field dynamics to simulate the flow:

$$(4.24) \quad \begin{aligned} \frac{d\theta_t}{dt} &= \int \left(\kappa(\theta_t, \theta', \rho_t) \rho_t(\theta') P(\theta_t, \theta', \rho_t) \nabla_{\theta'} \log \rho_{\text{post}}(\theta') \right. \\ &\quad \left. + \nabla_{\theta'} \cdot (\kappa(\theta_t, \theta', \rho_t) P(\theta_t, \theta', \rho_t)) \rho_t(\theta') \right) d\theta'. \end{aligned}$$

Similarly, we can prove the affine invariance of the above mean field dynamics; see Theorem 4.11 and its proof in Appendix D.6.

Theorem 4.11. *The mean-field dynamics eq. (4.24) is affine invariant under Condition 4.8.*

For a survey of convergence analysis of these gradient flows we refer to [26, Section 3.5].

4.5. Illustrative numerical experiments. In this subsection, we conduct numerical experiments comparing the affine invariant Wasserstein and Stein gradient flows with their non-affine invariant versions. We focus our experiments on three two-dimensional posterior distributions. In defining them we use the notation $\theta = [\theta^{(1)}, \theta^{(2)}]^T \in \mathbb{R}^2$.

(1) *Gaussian Posterior.*

$$\Phi_R(\theta) = \frac{1}{2} \theta^T \begin{bmatrix} 1 & 0 \\ 0 & \lambda \end{bmatrix} \theta \quad \text{with } \lambda = 0.01, 0.1, 1.$$

We initialize the gradient flows as follows:

$$(4.25) \quad \theta_0 \sim \mathcal{N}\left(\begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix}\right).$$

(2) *Logconcave Posterior.*

$$\Phi_R(\theta) = \frac{(\sqrt{\lambda} \theta^{(1)} - \theta^{(2)})^2}{20} + \frac{(\theta^{(2)})^4}{20} \quad \text{with } \lambda = 0.01, 0.1, 1.$$

We initialize the gradient flows according to (4.25).

(3) *General Posterior.*

$$\Phi_R(\theta) = \frac{\lambda(\theta^{(2)} - (\theta^{(1)})^2)^2}{20} + \frac{(1 - \theta^{(1)})^2}{20} \quad \text{with } \lambda = 0.01, 0.1, 1.$$

This example is known as the Rosenbrock function [61]. We initialize the gradient flows according to (4.25).

In all three examples, the parameter λ controls the anisotropy of the distribution. We will use several summary statistics to compare the resulting samples with the ground truth. They are $\mathbb{E}[\theta]$, the covariance $\text{Cov}[\theta]$, and $\mathbb{E}[\cos(\omega^T \theta + b)]$; for the last summary statistics we randomly draw $\omega \sim \mathcal{N}(0, I)$ and $b \sim \text{Uniform}(0, 2\pi)$ and report the average over 20 random draws of ω and b . The ground truths of these summary statistics are evaluated by integrating ρ_{post} numerically (See Appendix F for detailed formula of numerical integration). We will compare the performance of the following gradient flows (GFs) in our experiments.

- Wasserstein GF: The Wasserstein gradient flow, which is implemented by the Langevin dynamics (2.7).
- Affine invariant Wasserstein GF: The affine invariant Wasserstein gradient flow with $P(\theta, \rho) = C(\rho)$ and $h(\theta, \rho) = \sqrt{2C(\rho)}$; this is implemented by the stochastic interacting particle approximation of the preconditioned Langevin dynamics (4.18).
- Stein GF: The Stein gradient flow with

$$(4.26) \quad \kappa(\theta, \theta', \rho) = (1 + 4 \log(J + 1)/N_\theta)^{N_\theta/2} \exp(-\frac{1}{h} \|\theta - \theta'\|^2),$$

which is implemented as a deterministic interacting particle dynamics (4.21). Here $h = \text{med}^2 / \log(J + 1)$; med^2 is the squared median of the pairwise Euclidean distance between the current particles, following [84].

- Affine invariant Stein GF: The affine invariant Stein gradient flow with

$$(4.27) \quad P = C(\rho), \quad \kappa(\theta, \theta', \rho) = (1 + 2/N_\theta)^{N_\theta/2} \exp(-\frac{1}{2N_\theta} (\theta - \theta')^T C(\rho)^{-1} (\theta - \theta')),$$

which is implemented as a deterministic interacting particle dynamics (4.24).

Here the scaling constants⁸ in the above definition of kernel functions are chosen such that

$$(4.28) \quad \int \int \kappa(\theta, \theta', \rho) \mathcal{N}(\theta, m, C(\rho)) \mathcal{N}(\theta', m, C(\rho)) d\theta d\theta' = 1.$$

This choice makes the Stein gradient flows comparable with the Wasserstein gradient flows in terms of time scales, for a fair comparison. In practical implementation, it is recommended to drop these scaling constants. All the gradient flows are implemented by interacting particle systems with $J = 400$ particles (except for the Wasserstein GF, using (2.7), where particles do not interact). To reduce the influence of temporal discretization errors, all simulations use a fixed time step size of $\Delta t = 5 \times 10^{-3}$ with either the forward Euler or Euler–Maruyama integrator.

For the Gaussian posterior, the convergence of different gradient flows, according to the three summary statistics, are presented in Figure 1. The imposition of the affine invariance property makes the convergence rate independent of the anisotropy λ and accelerates the sampling for badly scaled Gaussian ($\lambda = 0.01$). We note that all these gradient flows do not converge within machine precision because of the limited number of particles.

For the logconcave posterior, the results are shown in Figure 2. Again, the imposition of the affine invariance property makes the convergence rate independent

⁸In the case of the Stein GF, there is no analytical formula for the integral eq. (4.28). To deal with this, we estimate the scaling constant by replacing $\text{med}^2 I$ with $N_\theta C(\rho)$, for which we can analytically compute eq. (4.28).




FIGURE 1. Gaussian posterior case: Convergence of different gradient flows in terms of the L^2 error of $\mathbb{E}[\theta_t]$, the relative Frobenius norm error of the covariance $\frac{\|\text{Cov}[\theta_t] - \text{Cov}[\theta_{\text{true}}]\|_F}{\|\text{Cov}[\theta_{\text{true}}]\|_F}$, and the error of $\mathbb{E}[\cos(\omega^T \theta_t + b)]$. We report the mean error across 10 random realizations, with shaded regions representing one standard deviation.

of the anisotropy λ and accelerates the sampling in the highly anisotropic case ($\lambda = 0.01$).

For the general posterior, we note that the Rosenbrock function is a non-convex function. Its minimum is at $[1, 1]$. The expectation and covariance of the posterior density function is (See Appendix F)

$$\mathbb{E}[\theta] = \begin{bmatrix} 1 \\ 11 \end{bmatrix} \quad \text{Cov}[\theta] = \begin{bmatrix} 10 & 20 \\ 20 & \frac{10}{\lambda} + 240 \end{bmatrix}.$$

The particles obtained by different gradient flows at $t = 15$ are depicted in Figure 3, and their convergence behaviors according to the three summary statistics are depicted in Figure 4. For small λ (e.g., $\lambda = 0.01$), $\theta^{(2)}$ is the stretch direction,




FIGURE 2. Logconcave posterior case: Convergence of different gradient flows in terms of the L^2 error of $\mathbb{E}[\theta_t]$, the relative Frobenius norm error of the covariance $\frac{\|\text{Cov}[\theta_t] - \text{Cov}[\theta_{\text{true}}]\|_F}{\|\text{Cov}[\theta_{\text{true}}]\|_F}$, and the error of $\mathbb{E}[\cos(\omega^T \theta_t + b)]$. We report the mean error across 10 random realizations, with shaded regions representing one standard deviation.

and therefore the imposition of the affine invariance property makes the convergence faster. However, when λ increases, the posterior density concentrates on a manifold with significant curvature (See Figure 3). Although the particle positions match well with the density contours, the convergence of summary statistics significantly deteriorates; the imposition of affine invariance does not significantly relieve the issue. This indicates that additional structures need to be explored for sampling densities such as this, which concentrate on a curved manifold.

Moreover, the convergence curves for the affine invariant Wasserstein gradient flow, implemented with the Langevin dynamics, oscillate slightly due to the added noise; those obtained from the affine invariant Stein gradient flow, implemented by




FIGURE 3. General posterior case: Particles obtained by different gradient flows at $t = 15$. Grey lines represent the contour of the true posterior.

Stein variational gradient descent, are smooth (See Figures 1, 2 and 4). The added noise helps for sampling non-Gaussian, highly anisotropic posteriors in comparison with the affine invariant Stein gradient flow (See Figures 3 and 4).

As a summary, affine invariant gradient flows, when implemented via particle methods, typically outperform their non-affine invariant versions, especially when the target distribution is highly anisotropic.

5. GAUSSIAN APPROXIMATE GRADIENT FLOW

An alternative to using particle methods to approximate a density is through parametric approximations. In this section, we focus on Gaussian approximations

of the gradient flows in the density space. The resulting Gaussian approximate gradient flows can be understood as gradient flows in the parameter space of Gaussian distributions, within the framework of variational inference [13, 69, 127].

In Section 5.1 and Section 5.2 we introduce Gaussian approximations based on metric projection and moment closure respectively. We show their equivalence under certain conditions in Section 5.3. Based on this equivalence, we use the moment closure approach to derive different Gaussian approximate gradient flows. In Section 5.4, we introduce several explicit Gaussian approximate gradient flows, establish their connections and discuss their relation to natural gradient methods in variational inference. In Section 5.5 we provide convergence analysis for these Gaussian approximate gradient flows, showing that the affine invariance property is beneficial for the convergence rates of the flows, for Gaussian and logconcave target distributions. We conduct illustrative numerical experiments in Section 5.6 to compare these gradient flows further; the results validate our theoretical analysis.

5.1. Gaussian approximation via metric projection. Consider the Gaussian density space

(5.1a)

$$\mathcal{P}^G := \left\{ \rho_a : \rho_a(\theta) = \frac{\exp(-\frac{1}{2}(\theta - m)^T C^{-1}(\theta - m))}{\sqrt{|2\pi C|}} \text{ with } a = (m, C) \in \mathcal{A} \right\},$$

$$(5.1b) \quad \mathcal{A} = \left\{ (m, C) : m \in \mathbb{R}^{N_\theta}, C \succ 0 \in \mathbb{R}^{N_\theta \times N_\theta} \right\}.$$

Instead of minimizing ρ over $\mathcal{P}$ as in the gradient flow in the whole probability density space, one can minimize over Gaussian density space $\mathcal{P}^G \subset \mathcal{P}$. Given an energy functional $\mathcal{E}$, this consideration leads to the following constrained minimization problem:

$$(5.2) \quad \min_{a \in \mathcal{A}} \mathcal{E}(\rho_a; \rho_{\text{post}}).$$

This forms the basis of Gaussian variational inference. For the reasons explained in Section 3 we choose $\mathcal{E}$ to be the KL divergence. The tangent space at any point of $\mathcal{A}$ is $\{(\delta m, \delta C) : \delta m \in \mathbb{R}^{N_\theta}, \delta C \in \mathbb{R}^{N_\theta \times N_\theta} \text{ is symmetric}\}$, which is isomorphic to $\mathbb{R}^{N_a}$, where $N_a = N_\theta + \frac{1}{2}(N_\theta + 1)N_\theta$ is the dimension of the parameter a . Given a metric tensor $M(\rho)$ in the tangent space of $\mathcal{P}$, we can define the projected metric tensor in the tangent space of $\mathcal{A}$ as

$$(5.3) \quad \mathfrak{M}(a) := \int \nabla_a \rho_a(\theta) (M(\rho_a) \nabla_a \rho_a^T(\theta)) d\theta,$$

defined by

$$(5.4) \quad \langle \mathfrak{M}(a) \sigma_1, \sigma_2 \rangle_{\mathbb{R}^{N_a}} = \langle M(\rho_a) \nabla_a \rho_a \cdot \sigma_1, \nabla_a \rho_a \cdot \sigma_2 \rangle, \quad \forall \sigma_1, \sigma_2 \in \mathbb{R}^{N_a}.$$

Given the metric tensor, we can write down the resulting gradient flow equation in $\mathcal{A}$:

$$(5.5) \quad \frac{\partial a_t}{\partial t} = -\mathfrak{M}(a_t)^{-1} \left. \frac{\partial \mathcal{E}(\rho_a; \rho_{\text{post}})}{\partial a} \right|_{a=a_t}.$$

This leads to a flow of the density ρ_{a_t} that stays in the Gaussian space $\mathcal{P}^G$. We remark that the above metric projection approach is standard in Riemannian geometry to obtain constrained metrics and gradient flows [45].




FIGURE 4. General posterior case: Convergence of different gradient flows in terms of the L^2 error of $\mathbb{E}[\theta_t]$, the relative Frobenius norm error of the covariance $\frac{\|\text{Cov}[\theta_t] - \text{Cov}[\theta_{\text{true}}]\|_F}{\|\text{Cov}[\theta_{\text{true}}]\|_F}$, and the error of $\mathbb{E}[\cos(\omega^T \theta_t + b)]$. We report the mean error across 10 random realizations, with shaded regions representing one standard deviation.

5.2. Gaussian approximation via moment closures. For any gradient flow in the density space eq. (1.2), we can also consider the following moment closure approach to get a Gaussian approximation. First, we write down evolution equations for the mean and covariance under eq. (1.2) noting that they satisfy the following identities:

$$\begin{aligned}
 \frac{dm_t}{dt} &= \frac{d}{dt} \int \rho_t(\theta) \theta d\theta = - \int M(\rho_t)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t} \theta d\theta, \\
 \frac{dC_t}{dt} &= \frac{d}{dt} \int \rho_t(\theta) (\theta - m_t)(\theta - m_t)^T d\theta \\
 &= - \int M(\rho_t b)^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_t} (\theta - m_t)(\theta - m_t)^T d\theta.
 \end{aligned}
 \tag{5.6}$$

This is not, in general, a closed system for the mean and covariance; this is because ρ_t is not, in general, determined by only the first and second moments. To close the system, we replace ρ_t by $\rho_{a_t} = \mathcal{N}(m_t, C_t)$ in eq. (5.6). Then we obtain the following closed system for the evolution of (m_t, C_t) :

$$(5.7) \quad \begin{aligned} \frac{dm_t}{dt} &= - \int M(\rho_{a_t})^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_{a_t}} \theta d\theta, \\ \frac{dC_t}{dt} &= - \int M(\rho_{a_t})^{-1} \frac{\delta \mathcal{E}}{\delta \rho} \Big|_{\rho=\rho_{a_t}} (\theta - m_t)(\theta - m_t)^T d\theta. \end{aligned}$$

The moment closure approach has been used in Kalman filtering [116] and Wasserstein gradient flows [78] to obtain a reduced system of equations.

5.3. Equivalence between two approaches. The metric projection approach encapsulated in identity (5.4) is rooted in Riemannian geometry, whilst the moment closure approach leading to (5.7) is based on probability. However these two seemingly distinct approaches lead to equivalent Gaussian approximations of the gradient flow; see Theorem 5.1 and its proof in Appendix E.2. In fact, such equivalence has been pointed out for the Wasserstein gradient flow in [78]. Our result provides a general condition for such equivalence.

Theorem 5.1. *Suppose the following condition holds:*

$$(5.8) \quad M(\rho_a)T_{\rho_a}\mathcal{P}^G = \text{span}\{\theta_i, \theta_i\theta_j, 1 \leq i, j \leq N_\theta\}.$$

Here $M(\rho_a)$ is the metric tensor and $T_{\rho_a}\mathcal{P}^G$ is the tangent space of the Gaussian density space $\mathcal{P}^G$. Moreover, $\theta_i, \theta_i\theta_j$ are understood as functions of θ . Then the mean and covariance evolution eq. (5.7) are equivalent to (5.5).

We have Theorem 5.2 showing that many metric tensors we have considered in Section 4 indeed satisfy the condition eq. (5.8). Its proof can be found in Appendix E.2.

Theorem 5.2. *Assumption (5.8) holds for the Fisher-Rao metric tensor, the affine invariant Wasserstein metric tensor with a preconditioner P independent of θ , and the affine invariant Stein metric tensor with a preconditioner P independent of θ and with a bilinear kernel $\kappa(\theta, \theta', \rho) = (\theta - m)^T A(\rho)(\theta' - m) + b(\rho)$; here we require $b \neq 0$, and A non-singular.*

This equivalence is useful as it allows us to use the moment closure approach to calculate the projected gradient flows. The moment closure approach is more convenient than the metric projection approach, since the former involves $M(\rho_{a_t})^{-1}$ while the latter involves $M(\rho_{a_t})$. We know from the examples of the Wasserstein and Stein metric tensors that $M(\rho_{a_t})^{-1}$ is a differential operator which is simpler in calculations compared to its inverse. On the other hand the projected gradient flow is an intrinsic object and so it is desirable that the easy-to-implement moment closure recovers it.

5.4. Equations of Gaussian approximate gradient flows. In this subsection, we use the moment closure approach to derive Gaussian approximate gradient flows for different metric tensors, discuss their connections, and draw relations to natural gradient methods in variational inference.

5.4.1. *Gaussian approximate Fisher-Rao gradient flow.* Applying the moment closure approach to Fisher-Rao gradient flow (4.2) leads to the following equation:

$$\begin{aligned}\frac{dm_t}{dt} &= \int \theta (\rho_{a_t} (\log \rho_{\text{post}} - \log \rho_{a_t}) - \rho_{a_t} \mathbb{E}_{\rho_{a_t}} [\log \rho_{\text{post}} - \log \rho_{a_t}]) d\theta \\ &= \text{Cov}_{\rho_{a_t}} [\theta, \log \rho_{\text{post}} - \log \rho_{a_t}] \\ &= C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} (\log \rho_{\text{post}} - \log \rho_{a_t})], \\ \frac{dC_t}{dt} &= \mathbb{E}_{\rho_{a_t}} [(\theta - m_t)(\theta - m_t)^T (\log \rho_{\text{post}} - \log \rho_{a_t} - \mathbb{E}_{\rho_{a_t}} [\log \rho_{\text{post}} - \log \rho_{a_t}])] \\ &= C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} (\log \rho_{\text{post}} - \log \rho_{a_t})] C_t,\end{aligned}$$

where $\rho_{a_t} \sim \mathcal{N}(m_t, C_t)$, and we have used the Stein's lemma (see Lemma E.1) in the last steps of the derivation. Furthermore noting that $\mathbb{E}_{\rho_{a_t}} [\nabla \log \rho_{a_t}] = 0$, we obtain

$$(5.9) \quad \begin{aligned}\frac{dm_t}{dt} &= C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{\text{post}}], \\ \frac{dC_t}{dt} &= C_t + C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] C_t.\end{aligned}$$

Remark 5.3. Eq. (5.9) is the same as the continuous limit of *natural gradient descent* [1, 96, 139] for Gaussian variational inference, where the Fisher information matrix is used to precondition the gradient descent dynamics for the optimization problem in eq. (5.2). In fact, the Fisher information matrix is equal to $\mathfrak{M}(\cdot)$ in eq. (5.3) when $M(\cdot)$ is chosen as the Fisher-Rao metric tensor [93].

The Gaussian approximate Fisher-Rao gradient flow (5.9) is affine invariant. Under any invertible affine transformation $\tilde{\theta} = \varphi(\theta) = A\theta + b$, the Gaussian density ρ_{a_t} with $a_t = (m_t, C_t)$ in eq. (5.9) will be transformed to another Gaussian density $\tilde{\rho}_a = \varphi \# \rho_a = \rho_{\tilde{a}_t}$ with $\tilde{a}_t = (\tilde{m}_t, \tilde{C}_t) = (Am, ACA^T)$. Let $\tilde{\rho}_{\text{post}} = \varphi \# \rho_{\text{post}}$. Then, the dynamics of $(\tilde{m}_t, \tilde{C}_t)$ remains the same form as (5.9) satisfying

$$(5.10) \quad \begin{aligned}\frac{d\tilde{m}_t}{dt} &= \tilde{C}_t \mathbb{E}_{\rho_{\tilde{a}_t}} [\nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}}], \\ \frac{d\tilde{C}_t}{dt} &= \tilde{C}_t + \tilde{C}_t \mathbb{E}_{\rho_{\tilde{a}_t}} [\nabla_{\tilde{\theta}} \nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}}] \tilde{C}_t.\end{aligned}$$

The above equation can be derived using Lemma D.1.

5.4.2. *Gaussian approximate Wasserstein gradient flow.* By applying the moment closure approach to (2.6), we get the mean and covariance evolution equations for the Gaussian approximate Wasserstein gradient flow with θ -independent P as follows:

$$(5.11) \quad \begin{aligned}\frac{dm_t}{dt} &= \int [\rho_{a_t} P(\rho_{a_t}) \nabla_{\theta} (\log \rho_{\text{post}} - \log \rho_{a_t})] d\theta = P(\rho_{a_t}) \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{\text{post}}], \\ \frac{dC_t}{dt} &= \int \nabla_{\theta} \cdot [\rho_{a_t} P(\rho_{a_t}) \nabla_{\theta} (\log \rho_{a_t} - \log \rho_{\text{post}})] (\theta - m_t)(\theta - m_t)^T d\theta \\ &= 2P(\rho_{a_t}) + P(\rho_{a_t}) \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] C_t + C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] P(\rho_{a_t}),\end{aligned}$$

where $\rho_{a_t} \sim \mathcal{N}(m_t, C_t)$, and we have used integration by parts and Stein's lemma (see Lemma E.1), and the fact that $\rho \nabla_{\theta} \log \rho = \nabla_{\theta} \rho$ and $\mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{a_t}] = 0$ in the above derivation.

When we set $P(\rho) \equiv I$, the evolution equation becomes

$$(5.12) \quad \begin{aligned} \frac{dm_t}{dt} &= \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{\text{post}}], \\ \frac{dC_t}{dt} &= 2I + \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] C_t + C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}]. \end{aligned}$$

This corresponds to the gradient flow under the constrained Wasserstein metric in the Gaussian density space [72, 78, 116].

Remark 5.4. When P is the identity operator, the metric tensor $\mathfrak{M}(a)$ defined in eq. (5.3) has an explicit formula [12, 28, 80, 92, 124], and the corresponding Gaussian density space is called the Bures–Wasserstein space [126].

If we set $P(\theta, \rho) = C(\rho)$, a choice that leads to affine invariant Wasserstein gradient flows in the density space, then the resulting evolution equations for the corresponding mean and covariance are

$$(5.13) \quad \begin{aligned} \frac{dm_t}{dt} &= C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{\text{post}}], \\ \frac{dC_t}{dt} &= 2C_t + 2C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] C_t. \end{aligned}$$

Eq. (5.13) is similar to the Gaussian approximate Fisher-Rao gradient flow (5.9), but with scaling factor 2 in the covariance evolution. This equation is also affine invariant, akin to the Gaussian approximate Fisher-Rao gradient flow (5.9).

5.4.3. Gaussian approximate Stein gradient flow. We apply the moment closure approach to (4.20). The mean and covariance evolution equations of the preconditioned Stein gradient flow (4.23) with a θ -independent P are

$$(5.14) \quad \begin{aligned} \frac{dm_t}{dt} &= - \int \left(\rho_{a_t}(\theta) \int \kappa(\theta, \theta', \rho_{a_t}) \rho_{a_t}(\theta') P_t \nabla_{\theta'} (\log \rho_{a_t}(\theta') - \log \rho_{\text{post}}(\theta')) d\theta' \right) d\theta, \\ \frac{dC_t}{dt} &= \\ &- \int \left(\rho_{a_t}(\theta) \int \kappa(\theta, \theta', \rho_{a_t}) \rho_{a_t}(\theta') P_t \nabla_{\theta'} (\log \rho_{a_t}(\theta') - \log \rho_{\text{post}}(\theta')) d\theta' \right) (\theta - m_t)^T \\ &+ (\theta - m_t) \left(\rho_{a_t}(\theta) \int \kappa(\theta, \theta', \rho_{a_t}) \rho_{a_t}(\theta') P_t \nabla_{\theta'} (\log \rho_{a_t}(\theta') - \log \rho_{\text{post}}(\theta')) d\theta' \right)^T d\theta, \end{aligned}$$

where $\rho_{a_t} \sim \mathcal{N}(m_t, C_t)$, $P_t := P(\rho_{a_t})$, and we have used integration by parts in the above derivation.

Imposing the form of the bilinear kernel mentioned in Theorem 5.2 and using the Stein's lemma (See Lemma E.1) and the fact that $\mathbb{E}_{\rho_{a_t}} [\nabla \log \rho_{a_t}] = 0$, we obtain

$$(5.15) \quad \begin{aligned} \frac{dm_t}{dt} &= b_t P_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{\text{post}}], \\ \frac{dC_t}{dt} &= P_t A_t C_t + C_t A_t P_t \\ &\quad + P_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] C_t A_t C_t + C_t A_t C_t \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] P_t. \end{aligned}$$

Here, we used the notations that $A_t = A(\rho_{a_t})$ and $b_t = b(\rho_{a_t})$.

Different choices of the preconditioner P and the bilinear kernel κ in Theorem 5.2 allow us to construct various Gaussian approximate gradient flows. Setting

$P(\theta, \rho) = C(\rho)$ and choosing the bilinear kernel with $A(\rho) = \frac{1}{2}C(\rho)^{-1}$ and $b(\rho) = 1$ lead to the Gaussian approximate Fisher-Rao gradient flow in eq. (5.9). Choosing the preconditioner $P(\rho) = I$ and the bilinear kernel with $A(\rho) = C(\rho)^{-1}$ and $b(\rho) = 1$ leads to the Gaussian approximate Wasserstein gradient flow in eq. (5.12). Moreover, setting the preconditioner $P_t = I$ and the bilinear kernel with $A(\rho) = I$ and $b(\rho) = 1$ recovers the Gaussian sampling approach introduced in [56]:

$$(5.16) \quad \begin{aligned} \frac{dm_t}{dt} &= \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{\text{post}}], \\ \frac{dC_t}{dt} &= 2C_t + \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] C_t^2 + C_t^2 \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}]. \end{aligned}$$

We note that these dynamics are not affine invariant.

The above discussion implies that the Gaussian approximate Stein gradient flow is quite general: with various preconditioners and bilinear kernels, it can recover many Gaussian dynamics used in sampling.

Remark 5.5. In many sampling problems, we may evaluate $\nabla_{\theta} \log \rho_{\text{post}}$ but not the Hessian matrix $\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}$. Nevertheless we note that Stein's identities, as presented in Lemma E.1, allow one to eliminate one derivative in the aforementioned dynamics, through

$$(5.17) \quad \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] = \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{\text{post}} (\theta - m_t)^T] C_t^{-1}.$$

5.5. Convergence analysis. In the last subsection, we obtained various Gaussian approximate gradient flows, all of which can be used to obtain a Gaussian approximation of the target distribution. Some of the Gaussian dynamics are affine invariant and some are not. The goal of the subsection is to provide convergence analysis of these dynamics and understand the effect of affine invariance for different target distributions. We will consider three classes of target distribution: the Gaussian posterior case, logconcave posterior case, and general posterior case.

5.5.1. Gaussian posterior case. We start with the Gaussian posterior case. For such case, we are able to compute the explicit formula of the dynamics and then establish convergence rates. We consider the Gaussian approximate Fisher-Rao and Wasserstein gradient flows here. Additionally, for the purpose of comparison, we also consider the standard gradient flow for Gaussian variational inference, employing the Euclidean inner-product metric in $\mathbb{R}^{N_a}$ (i.e., setting $\mathfrak{M}(a_t) = I$ in (5.5)). This choice leads to the following dynamical system

$$(5.18) \quad \begin{aligned} \frac{dm_t}{dt} &= \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \log \rho_{\text{post}}], \\ \frac{dC_t}{dt} &= \frac{1}{2} C_t^{-1} + \frac{1}{2} \mathbb{E}_{\rho_{a_t}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}]. \end{aligned}$$

In this article, we call it the vanilla Gaussian approximate gradient flow. In Theorem 5.6, we provide convergence rates for the three different dynamics.

Theorem 5.6. *Assume the posterior distribution is Gaussian $\rho_{\text{post}}(\theta) = \mathcal{N}(\theta, m_{\star}, C_{\star})$. Denote the largest eigenvalue of $C_{\star}$ by $\lambda_{\star, \max}$. For gradient flows with initialization $C_0 = \lambda_0 I$, where $\lambda_0 > 0$, the following hold:*

- (1) *for the Gaussian approximate Fisher-Rao gradient flow (5.9):*

$$\|m_t - m_{\star}\|_2 = \mathcal{O}(e^{-t}), \|C_t - C_{\star}\|_2 = \mathcal{O}(e^{-t});$$

(2) for the Gaussian approximate Wasserstein gradient flow (5.12):

$$\|m_t - m_\star\|_2 = \mathcal{O}(e^{-t/\lambda_{\star, \max}}), \|C_t - C_\star\|_2 = \mathcal{O}(e^{-2t/\lambda_{\star, \max}});$$

(3) for the vanilla Gaussian approximate gradient flow (5.18):

$$\|m_t - m_\star\|_2 = \mathcal{O}(e^{-t/\lambda_{\star, \max}}), \|C_t - C_\star\|_2 = \mathcal{O}(e^{-t/(2\lambda_{\star, \max}^2)}),$$

where the implicit constants depend on m_0 , λ_0 , $m_\star$, and $C_\star$.

The proof can be found in Appendix E.5, which also provides details about the implicit constants. In particular, it shows that the convergence is exponential in time, with constants that depend linearly on the discrepancy between the initial mean and the posterior mean. For the Gaussian approximate gradient flow (5.18) and the Gaussian approximate Wasserstein gradient flow (5.12), if $\lambda_{\star, \max}$ is large, their convergence rate is much slower than the Gaussian approximate Fisher-Rao gradient flow. The affine invariance property of the Gaussian approximate Fisher-Rao gradient flow allows it to achieve a uniform exponential rate of convergence $\mathcal{O}(e^{-t})$ for Gaussian posteriors.

The uniform convergence rate $\mathcal{O}(e^{-t})$ has also been studied for the mean and covariance dynamics of the Kalman-Wasserstein gradient flow (5.13) in [57, Lemma 3.2], [16, 18], which is affine invariant. As a consequence, the affine invariance property accelerates the convergence of the flow when the target distribution is Gaussian.

Can affine invariance also accelerate the dynamics beyond Gaussian posteriors? To investigate this question, in the following we study the convergence property of the Gaussian approximate Fisher-Rao gradient flow, which is affine invariant, for logconcave posteriors and general posteriors.

5.5.2. Logconcave posterior case. We consider an α -strongly logconcave posterior distribution: for some $\alpha > 0$

$$-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta) \succeq \alpha I;$$

note that the Gaussian posterior is a special case. The Gaussian variational inference problem eq. (5.2) with the KL divergence admits a unique minimizer, since under the α -strongly logconcave assumption, the KL divergence is strongly convex over the space of Gaussian densities equipped with the Wasserstein metric [78, Lemma 2]. However, unlike the Gaussian posterior case (as shown in Theorem 5.6), the convergence of the Gaussian approximate Fisher-Rao gradient flow is no longer universal, across a class of posteriors ρ_{post} . We can nonetheless establish Theorem 5.7 regarding lower and upper bounds on the local convergence rate of the Gaussian approximate Fisher-Rao gradient flow ODE system (5.9). Here, the local convergence rate refers to the exponential convergence rate when the initial condition is sufficiently close to the unique stationary point of (5.9), and is characterized by the eigenvalues of the linearized Jacobian matrix.

Theorem 5.7. For $0 < \alpha < \beta < \infty$, assume the posterior distribution ρ_{post} is α -strongly logconcave and β -log-smooth such that

$$\log \rho_{\text{post}} \in C^2(\mathbb{R}^{N_\theta}), -\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta) \succeq \alpha I, -\nabla_\theta \nabla_\theta \log \rho_{\text{post}} \preceq \beta I.$$

Denote the unique minimizer of the Gaussian variational inference problem eq. (5.2) with the KL divergence as the energy functional by $\rho_{a_\star} := \mathcal{N}(m_\star, C_\star)$. For $N_\theta = 1$, let $\lambda_{\star, \max} < 0$ denote the largest eigenvalue of the linearized Jacobian matrix of the

Gaussian approximate Fisher-Rao gradient flow (5.9) around $m_\star$ and $C_\star$. Then we have

$$(5.19) \quad -\lambda_{\star, \max} \geq \frac{1}{(7 + \frac{4}{\sqrt{\pi}})(1 + \log(\frac{\beta}{\alpha}))}.$$

Moreover, the bound is sharp: it is possible to construct a sequence of triplets $\rho_{\text{post},n}$, α_n and β_n , where $\lim_{n \rightarrow \infty} \frac{\beta_n}{\alpha_n} = \infty$, such that, if we let $\lambda_{\star, \max, n}$ denote the corresponding largest eigenvalues of the linearized Jacobian matrix for the n -th triple, then, it holds that

$$-\lambda_{\star, \max, n} = \mathcal{O}\left(1 / \log \frac{\beta_n}{\alpha_n}\right).$$

The proof can be found in Appendix E.6. For the counterexample, we construct $\rho_{\text{post},n}$ such that $-\nabla_\theta \nabla_\theta \rho_{\text{post},n}$ is a bump function between α_n and β_n with the width of the bump gradually approaching 0. Our analysis reveals that this construction can drive $-\lambda_{\star, \max}$ as small as possible.

Theorem 5.7 implies that the local convergence rate of the Gaussian approximate Fisher-Rao gradient flow is $\log(\frac{\beta}{\alpha})$ for logconcave posteriors. Therefore, such affine invariant flow is still advantageous in sampling highly anisotropic logconcave posteriors, since the dependence of the local convergence rate on the anisotropic ratio or condition number β/α is only logarithmic.

Recently, the work [78] proved the global convergence of the Gaussian approximate Wasserstein gradient flow (which is not affine invariant) when the posterior is logconcave. We leave it as a future work to prove the global convergence rate for the Gaussian approximate Fisher-Rao gradient flows.

5.5.3. General posterior case. In general, we can construct posteriors such that the convergence of the Gaussian approximate Fisher Rao gradient flow to a stationary point can be arbitrarily slow.

Theorem 5.8. *For any $K > 0$ there exists a target ρ_{post} such that, for the three Gaussian approximate gradient flows eq. (5.9), eq. (5.12) and eq. (5.18), the convergence rate to their stationary points can be as slow as $\Theta(t^{-\frac{1}{2K}})$.*

The proof can be found in Appendix E.7. This demonstrates that the use of affine invariant properties cannot achieve a universal acceleration for all posterior distributions.

5.6. Illustrative numerical examples. In this subsection, we perform numerical experiments comparing different Gaussian approximate gradient flows. The test target distributions from Section 4.5 are used. We consider the three dynamical models for mean and covariance given in eq. (5.9), eq. (5.12) and eq. (5.18). Letting $\rho_{a_n} = \mathcal{N}(m_n, C_n)$ denote the approximated Gaussian density at time $n\Delta t$, the Gaussian approximate Fisher-Rao gradient flow (5.9) is discretized as

$$(5.20) \quad \begin{aligned} m_{n+1} &= m_n + \Delta t C_n \mathbb{E}_{\rho_{a_n}} [\nabla_\theta \log \rho_{\text{post}}], \\ C_{n+1}^{-1} &= C_n^{-1} - \Delta t (C_n^{-1} + \mathbb{E}_{\rho_{a_n}} [\nabla_\theta \nabla_\theta \log \rho_{\text{post}}]). \end{aligned}$$

Here, we use the forward Euler scheme to discretize the dynamics of C^{-1} rather than C , for which we observe better numerical stability. Following [78], the Gaussian

approximate Wasserstein gradient flow (5.12) is discretized as

$$\begin{aligned}
 m_{n+1} &= m_n + \Delta t \mathbb{E}_{\rho_{a_n}} [\nabla_{\theta} \log \rho_{\text{post}}], \\
 (5.21) \quad C_{n+1} &= \left(I + \Delta t (\mathbb{E}_{\rho_{a_n}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] + C_n^{-1}) \right) C_n \\
 &\quad \times \left(I + \Delta t (\mathbb{E}_{\rho_{a_n}} [\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}] + C_n^{-1}) \right).
 \end{aligned}$$

The Gaussian approximate gradient flow (5.18) is discretized by the forward Euler scheme directly. The expectations in the evolution equations are calculated by the unscented transform [65, 66, 71] with $J = 2N_{\theta} + 1 = 5$ quadrature points. In this implementation the Gaussian approximation has considerable speedup in comparison with the previously mentioned particle-based sampling approaches, where $J = 400$. We use the same summary statistics as in Section 4.5 to compare the obtained Gaussian with the true target distribution.

For the Gaussian posterior, the convergence of different gradient flows, according to the three summary statistics, are presented in Figure 5. The imposition of the affine invariance property makes the convergence rate independent of the anisotropy λ and accelerates the sampling for badly scaled Gaussian ($\lambda = 0.01$). The convergence rates of Gaussian approximate gradient flows match well with the predictions of Theorem 5.6.

For the logconcave posterior, the convergence behaviors of different gradient flows, according to the three summary statistics, are presented in Figure 6. On the one hand, we observe that for two out of three summary statistics in Figure 6, the relative error reaches a plateau. This stagnation stems from the fact that ρ_{a_t} converges to the unique minimizer ρ_{a_*} , which differs from the true posterior ρ_{post} . On the other hand, the imposition of affine invariance accelerates sampling in highly anisotropic settings (e.g., $\lambda = 0.01$), rendering the early convergence behavior largely insensitive to the anisotropy parameter λ . Notably, the performance of the Gaussian approximate Fisher–Rao gradient flow does not deteriorate with increasing anisotropy. This suggests that the local convergence rate established in Theorem 5.7 for this flow may be extended to arbitrary dimensionalities, suggesting the possibility of achieving a global convergence rate as well.

For the general posterior, the estimated posterior densities (3 standard deviations) obtained by different Gaussian approximate gradient flows are presented in Figure 7, and their convergence behaviors according to the three summary statistics are depicted in Figure 8. Here, the use of Gaussian approximations cannot represent the posterior distribution well because the posterior is far from Gaussians.

In summary, along with the numerical results obtained by particle methods in Section 4.5, we observe that affine-invariant modifications improve convergence for both Gaussian and particle approximations. Generally, Gaussian approximation methods are more efficient for sampling Gaussian and logconcave posterior distributions, as well as general posteriors with small λ , in terms of the number of function evaluations. However, this efficiency comes at the expense of computing the Hessian matrix. For general posterior distributions with $\lambda = 1$ that deviate significantly from Gaussian distributions, Gaussian approximation may not accurately represent the distribution (See Figure 7). In such cases, particle methods, due to their non-parametric nature, might offer improved representations (See Figure 3), but achieving convergence remains challenging and could require a large number of particles, with or without affine-invariant modifications.




FIGURE 5. Gaussian posterior case: Convergence of different dynamics in terms of L^2 error of $\mathbb{E}[\theta_t]$, the relative Frobenius norm error of the covariance $\frac{\|\text{Cov}[\theta_t] - \text{Cov}[\theta_{\text{true}}]\|_F}{\|\text{Cov}[\theta_{\text{true}}]\|_F}$, and the error of $\mathbb{E}[\cos(\omega^T \theta_t + b)]$

6. APPLICATION: DARCY FLOW

In this section, we assess the effectiveness of the aforementioned approaches in addressing Bayesian inverse problems, specifically focusing on a 1D Darcy flow problem. The Darcy equation describes the pressure field $p(x)$ in a porous medium with a parameterized, positive permeability field $a(x, \theta)$, through the following PDE

$$(6.1) \quad \begin{aligned} -\partial_x(a(x, \theta)\partial_x p(x)) &= f(x), & x \in D, \\ p(x) &= 0, & x \in \partial D. \end{aligned}$$

Here the computational domain is $D = [0, 1]$. Homogeneous Dirichlet boundary conditions are applied on ∂D , and f is the source of the fluid (see the left of




FIGURE 6. Logconcave posterior case: Convergence of different dynamics in terms of L^2 error of $\mathbb{E}[\theta_t]$, the relative Frobenius norm error of the covariance $\frac{\|\text{Cov}[\theta_t] - \text{Cov}[\theta_{\text{true}}]\|_F}{\|\text{Cov}[\theta_{\text{true}}]\|_F}$, and the error of $\mathbb{E}[\cos(\omega^T \theta_t + b)]$

Figure 9):

$$(6.2) \quad f(x) = \begin{cases} 2000 & 0 \leq x \leq \frac{1}{3} \\ 1000 & \frac{1}{3} < x \leq \frac{2}{3} \\ 0 & \frac{2}{3} < x \leq 1 \end{cases}.$$

The Bayesian inverse problem considered here is to determine parameter θ of the field $a(\cdot; \theta)$ (see the right of Figure 9) from observations y_{ref} , which consist of point-wise measurements of the pressure $p(x)$ at 7 equidistant points in the domain (see the middle of Figure 9), corrupted with observation noises $\eta \sim \mathcal{N}(0, I)$. The field




FIGURE 7. General posterior case: Density functions (3 standard deviations) obtained by different dynamics at $t = 15$. Grey lines represent the contour of the true posterior

$a(x, \theta)$ is parametrized as

$$(6.3) \quad \log a(x, \theta) = \sum_{l=1} \theta_{(l)} \sqrt{\lambda_l} \psi_l(x),$$

where the basis functions and eigenvalues are defined by

$$(6.4) \quad \psi_l(x) = \sqrt{2} \cos(\pi l x), \quad \lambda_l = (\pi^2 |l|^2 + \tau^2)^{-2},$$

and $\theta_{(l)} \sim \mathcal{N}(0, 10^2)$ i.i.d. We note that these considerations amount to assuming that $\log a(x, \theta)$ is a mean zero Gaussian random field with covariance

$$(6.5) \quad \mathbf{C} = (-\Delta + \tau^2 I)^{-2}.$$

Here $-\Delta$ is the Laplacian operator on D subject to homogeneous Neumann boundary conditions, acting on the space of spatial mean zero functions [40]. In fact, (6.3) is the Karhunen-Loève expansion of the Gaussian random field. Here, for numerical studies, we truncate the sum (6.3) to the first $N_\theta = 16$ terms and we take $\tau = 3$.

The forward problem, denoted by $y = \mathcal{G}(\theta)$, is defined as the map from θ to $p(x)$ at the 7 equidistant points. Numerically, the map is implemented through solving the Darcy eq. (6.1) by a finite difference method on a size-128 grid to obtain $p(x)$ at the 7 observation locations. A typical computational goal in Bayesian inverse problems is to sample the posterior distribution

$$(6.6) \quad \rho_{\text{post}} \propto \exp(-\Phi_R(\theta)), \quad \Phi_R(\theta) = \frac{1}{2} (y_{\text{ref}} - \mathcal{G}(\theta))^T (y_{\text{ref}} - \mathcal{G}(\theta)) + \frac{1}{200} \theta^T \theta,$$

where we specify a Gaussian prior $\mathcal{N}(0, 10^2 I)$ on θ . Note that here $y_{\text{ref}}, \mathcal{G}(\theta) \in \mathbb{R}^7$ and $\theta \in \mathbb{R}^{16}$. The evaluation of $\Phi_R(\theta)$ requires solving the Darcy eq. (6.1), and the corresponding gradients and Hessians are computed automatically through the Julia ForwardDiff library [113].

The reference, ground truth posterior distribution is obtained by using the preconditioned Crank-Nicolson algorithm (pCN) [36] (which is a specific function-space MCMC algorithm). We run pCN to obtain 2×10^8 samples (with 5×10^7 samples for the burn-in period), using step size 0.02, yielding an average acceptance rate of 0.51. The estimated covariance matrix is plotted in Figure 10; this figure demonstrates that there is significant anisotropy. We run and compare all the aforementioned gradient flow approaches. We initialize them at $\mathcal{N}(0, I)$ instead of the prior; this initialization allows for larger stable time step sizes. Details of




FIGURE 8. General posterior case: Convergence of different dynamics in terms of L_2 error of $\mathbb{E}[\theta_t]$, the relative Frobenius norm error of the covariance $\frac{\|\text{Cov}[\theta_t] - \text{Cov}[\theta_{\text{true}}]\|_F}{\|\text{Cov}[\theta_{\text{true}}]\|_F}$, and the error of $\mathbb{E}[\cos(\omega^T \theta_t + b)]$

particle based approaches, including Wasserstein GF, affine invariant Wasserstein GF, Stein GF, and affine invariant Stein GF have been presented in Section 4.5. The scaling constants in the definition of kernels κ in eqs. (4.26) and (4.27) are dropped, since they only affect the time step size. We use interacting particle systems with $J = 100$ particles. Details of Gaussian approximated approaches, including Gaussian approximated Fisher-Rao GF, Gaussian approximated Wasserstein GF, and Gaussian approximated GF have been presented in Section 5.6. The expectations in the evolution equations are calculated by the unscented transform with $J = 2N_\theta + 1 = 33$ quadrature points.




FIGURE 9. From left to right: Source term $f(x)$, reference pressure field $p(x)$ with 7 equidistant pointwise measurements corrupted by observation noises $\eta \sim \mathcal{N}(0, I)$, and reference permeability field $a(x, \theta)$ of the Darcy flow problem




FIGURE 10. The estimated posterior covariance matrix through MCMC, in the Darcy flow problem

The maximum stable time step sizes (Δt) for each method, when using either the forward Euler or Euler-Maruyama time integration schemes, are outlined below: 0.030 (Wasserstein GF), 0.030 (Affine invariant Wasserstein GF), 0.990 (Stein GF), 0.090 (Affine Invariant Stein GF), 0.162 (Gaussian Approximated Fisher-Rao GF), 0.002 (Gaussian Approximated GF), and 0.018 (Gaussian Approximated Wasserstein GF). We determine these maximum stable time steps through iterative looping, incrementally stepping by 10^{-3} .

The convergence curves of these approaches (in terms of relative mean errors and relative covariance errors) are presented in Figure 11. The Gaussian approximate Fisher-Rao GF exhibits notably faster convergence, as a consequence of its relatively large time step size and its affine invariance property. Affine invariant modifications of the Wasserstein GF and Stein GF also show improved convergence rates compared to their non-affine invariant counterparts. Both the Wasserstein GF and affine invariant Wasserstein GF oscillate slightly due to the added noise.

Figure 12 shows properties of the approximated posterior distributions by different methods, after the 5000th iteration. We compare them with the result by MCMC and with the true parameters of the permeability field (referred to as “Reference”). Notably, the true parameters fall within the 3- σ confidence intervals determined by MCMC consistently. We observe that most approaches can reproduce the posterior mean and confidence intervals computed by MCMC accurately after 5000 iterations, except for the Gaussian approximate GF and Stein GF which exhibit very slow convergence.




FIGURE 11. The relative mean errors and relative covariance errors of θ obtained by various gradient flow approaches, in the Darcy flow problem. The observed error stagnation arises from (i) finite particle size effects for particle methods, and (ii) approximation errors inherent to Gaussian methods.




FIGURE 12. The estimated parameters $\theta_{(i)}$ (conditional mean) and the associated 3- σ confidence intervals (dashed lines) obtained by various gradient flow approaches and MCMC, in the Darcy flow problem

These experiments demonstrate the potential of gradient flow approaches, designed using the KL energy functional with various metrics, numerically realized by either particle or Gaussian approximations, for efficiently solving inverse problems in scientific computing with a reasonable level of accuracy (often around 10% error in the examples shown in this paper). Particularly noteworthy is the efficient and robust performance of gradient flows equipped with affine invariance, such as the Gaussian approximate Fisher-Rao gradient flow, an observation that aligns closely with our theoretical expectations.

For inverse problems in higher dimensions and in many natural science and engineering applications, computing derivatives of the forward map is costly and may not even be feasible [66]. In such scenario, one can further apply derivative-free approximations to the Gaussian approximate and particle-approximate gradient flows. Specifically for Fisher-Rao gradient flows, the resulting algorithms connect to ensemble Kalman filters and unscented Kalman filters. In fact, the Kalman inversion algorithm in [65] can be interpreted as a derivative-free approximation of Fisher-Rao gradient flows, which are shown to be efficient in several large-scale applications. Moreover, when the effective dimension is low, this structure may be incorporated into the gradient flow design, following approaches in [35, 129, 138] to further accelerate sampling.

7. CONCLUSIONS

In this work, we have studied the design ingredients of continuous-time gradient flows for sampling distributions with unknown normalization constants, focusing on the energy functional and the metric tensor. Regarding the energy functional, we show that the KL divergence has the *unique* property (among all f -divergences) that gradient flows resulting from this energy functional do not depend on the normalization constant of the target distribution. This makes the KL divergence advantageous in terms of numerical implementations. Regarding the metric tensor, we highlight the importance of invariance properties and, in particular, their influence on the convergence rates of the gradient flow. The unique diffeomorphism invariance property of the Fisher-Rao gradient flow allows it to achieve a uniform exponential convergence rate under general conditions; however the particle implementation of the Fisher-Rao gradient flow is highly non-trivial. We introduce a relaxed, affine invariance property for the gradient flows; in particular, we construct various affine invariant Wasserstein and Stein gradient flows. These affine invariant gradient flows are more convenient to approximate numerically via particle methods than the diffeomorphism invariant Fisher-Rao gradient flow, and they behave more favorably compared to their non-affine-invariant versions when sampling highly anisotropic distributions.

In addition, we study Gaussian approximations of the flow that lead to efficient implementable algorithms as alternatives to particle methods. In particular, Gaussian approximations can be readily applied for the diffeomorphism invariant Fisher-Rao gradient flow. Our theory and numerical experiments demonstrate the strengths and potential limitations of the Gaussian approximate Fisher-Rao gradient flow, which is affine invariant, for a wide range of target distributions.

Our study shows that for general posterior distributions beyond the Gaussian and logconcave class, the consideration of affine invariance and Gaussian approximation may not be enough for designing an accurate and efficient sampler. Examples include posteriors that concentrate on manifolds with significant curvatures such as the Rosenbrock function and also multi-modal posteriors. In our future work, we will explore more sophisticated approximations such as Gaussian mixtures [24, 27, 41, 42, 81] and other invariance properties to design samplers for these challenging scenarios. Furthermore, it is possible to incorporate a Metropolization step into interacting particle systems [122] to mitigate finite-particle error, which offers another promising avenue for improving sampling accuracy.

Data and code availability. All codes used to produce the numerical results and figures in this paper are available at <https://github.com/PKU-CMEGroup/InverseProblems.jl/tree/master/Gradient-Flow>.

APPENDIX A. UNIQUE PROPERTY OF THE KL DIVERGENCE ENERGY

A.1. Proof of Theorem 3.1.

Proof. We first note that the KL divergence satisfies the desired property: for any $c \in (0, \infty)$ and $\rho_{\text{post}} \in \mathcal{P}$ it holds that

$$\text{KL}[\rho \| c\rho_{\text{post}}] - \text{KL}[\rho \| \rho_{\text{post}}] = -\log c.$$

Now we establish uniqueness. For any f -divergence with property that $D_f[\rho \| c\rho_{\text{post}}] - D_f[\rho \| \rho_{\text{post}}]$ is independent of ρ , we have

$$(A.1) \quad \lim_{t \rightarrow 0} \frac{(D_f[\rho + t\sigma \| c\rho_{\text{post}}] - D_f[\rho + t\sigma \| \rho_{\text{post}}]) - (D_f[\rho \| c\rho_{\text{post}}] - D_f[\rho \| \rho_{\text{post}}])}{t} = 0,$$

for any bounded, smooth function σ supported in $B^d(0, R)$ that integrates to zero. Here $R > 0$ is a finite parameter that we will choose later, and $d = N_\theta$ is the dimension of θ . In the above formula we have used the fact that for sufficiently small t , one has $\rho + t\sigma \in \mathcal{P}$ since R is finite, so the formula is well-defined. Using the identity

$$\lim_{t \rightarrow 0} \frac{D_f[\rho + t\sigma \| \rho_{\text{post}}] - D_f[\rho \| \rho_{\text{post}}]}{t} = \int f' \left(\frac{\rho}{\rho_{\text{post}}} \right) \sigma d\theta,$$

we can rewrite eq. (A.1) as

$$(A.2) \quad \int_{B^d(0, R)} \left(f' \left(\frac{\rho}{c\rho_{\text{post}}} \right) - f' \left(\frac{\rho}{\rho_{\text{post}}} \right) \right) \sigma d\theta = 0.$$

Since σ is arbitrary, $f' \left(\frac{\rho}{c\rho_{\text{post}}} \right) - f' \left(\frac{\rho}{\rho_{\text{post}}} \right)$ must be a constant function in $B^d(0, R)$.

Because ρ and ρ_{post} integrate to 1 and they are continuous, there exists $\theta^\dagger$ such that $\rho(\theta^\dagger)/\rho_{\text{post}}(\theta^\dagger) = 1$. Choose R sufficiently large such that $\theta^\dagger \in B^d(0, R)$. Then, we obtain

$$(A.3) \quad f' \left(\frac{\rho(\theta)}{\rho_{\text{post}}(\theta)} \right) - f' \left(\frac{\rho(\theta)}{c\rho_{\text{post}}(\theta)} \right) = f' \left(\frac{\rho(\theta^\dagger)}{\rho_{\text{post}}(\theta^\dagger)} \right) - f' \left(\frac{\rho(\theta^\dagger)}{c\rho_{\text{post}}(\theta^\dagger)} \right) = f'(1) - f'(1/c),$$

for any $\theta \in B^d(0, R)$. As R can be arbitrarily large, the above identity holds for all $\theta \in \mathbb{R}^d$. Let $g(1/c) := f'(1) - f'(1/c)$. We have obtained

$$(A.4) \quad f'\left(\frac{\rho(\theta)}{\rho_{\text{post}}(\theta)}\right) - f'\left(\frac{\rho(\theta)}{c\rho_{\text{post}}(\theta)}\right) = g(1/c),$$

where c is an arbitrary positive number. Note, furthermore, that $g(\cdot)$ is continuous since f is continuously differentiable. Since ρ and ρ_{post} are arbitrary, we can write (A.4) equivalently as

$$(A.5) \quad f'(y) - f'(cy) = g(c),$$

for any $y, c \in \mathbb{R}_+$. Let $h : \mathbb{R} \rightarrow \mathbb{R}$ such that $h(z) = f'(\exp(z))$. Then, we can equivalently formulate (A.5) as

$$(A.6) \quad h(z_1) - h(z_2) = r(z_1 - z_2),$$

for any $z_1, z_2 \in \mathbb{R}$ and $r : \mathbb{R} \rightarrow \mathbb{R}$ such that $r(t) = g(\exp(-t))$.

We can show r is linear function. Setting $z_1 = z_2$ in (A.6) shows that $r(0) = 0$. Note also that, again by (A.6),

$$r(z_1 - z_2) + r(z_2 - z_3) = h(z_1) - h(z_3) = r(z_1 - z_3).$$

Hence, since z_1, z_2 and z_3 are arbitrary, we deduce that for any $x, y \in \mathbb{R}$, it holds that

$$(A.7) \quad r(x) + r(y) = r(x + y).$$

Furthermore r is continuous since f is assumed continuously differentiable. With the above conditions, it is a standard result in functional equations that $r(x)$ is linear. Indeed, as a sketch of proof, by (A.7) we can first deduce $r(n) = nr(1)$ for $n \in \mathbb{Z}$. Then, by setting x, y to be dyadic rationals, we can deduce $r(\frac{i}{2^k}) = \frac{i}{2^k}r(1)$ for any $i, k \in \mathbb{Z}$. Finally using the continuity of the function r , we get $r(x) = xr(1)$ for any $x \in \mathbb{R}$. For more details see [54].

Using the fact that r is linear and eq. (A.6), we know that h is an affine function and thus $f'(\exp(z)) = az + b$ for some $a, b \in \mathbb{R}$. Equivalently, $f'(y) = a \log(y) + b$. Using the condition $f(1) = 0$, we get $f(y) = ay \log(y) + (b - a)(y - 1)$. Plugging this f into the formula for D_f , we get

$$D_f[\rho || \rho_{\text{post}}] = a \text{KL}[\rho || \rho_{\text{post}}],$$

noting that the affine term in $f(y)$ has zero contributions in the formula for D_f . The proof is complete. $\square$

APPENDIX B. PROOF FOR THE CONVERGENCE OF FISHER-RAO GRADIENT FLOWS

B.1. Proof of Theorem 4.1.

Proof. The Fisher-Rao gradient flow of the KL divergence eq. (4.2) can be solved analytically using the variation of constants formula as follows. First note that

$$\frac{\partial \log \rho_t(\theta)}{\partial t} = \log \rho_{\text{post}}(\theta) - \log \rho_t(\theta) - \mathbb{E}_{\rho_t}[\log \rho_{\text{post}}(\theta) - \log \rho_t(\theta)],$$

so that

$$\frac{\partial e^t \log \rho_t(\theta)}{\partial t} = e^t \log \rho_{\text{post}}(\theta) - e^t \mathbb{E}_{\rho_t}[\log \rho_{\text{post}}(\theta) - \log \rho_t(\theta)].$$

Thus

$$\log \rho_t(\theta) = (1 - e^{-t}) \log \rho_{\text{post}}(\theta) + e^{-t} \log \rho_0(\theta) - \int_0^t e^{\tau-t} \mathbb{E}_{\rho_\tau} [\log \rho_{\text{post}}(\theta) - \log \rho_\tau(\theta)] d\tau.$$

It follows that

$$(B.1) \quad \rho_t(\theta) = \frac{1}{Z_t} \rho_0(\theta)^{e^{-t}} \rho_{\text{post}}(\theta)^{1-e^{-t}}, \quad \frac{\rho_t(\theta)}{\rho_{\text{post}}(\theta)} = \frac{1}{Z_t} \left(\frac{\rho_0(\theta)}{\rho_{\text{post}}(\theta)} \right)^{e^{-t}},$$

where Z_t is a normalization constant. Indeed, we have the formula

$$Z_t = \int \left(\frac{\rho_0(\theta)}{\rho_{\text{post}}(\theta)} \right)^{e^{-t}} \rho_{\text{post}}(\theta) d\theta$$

by using the condition $\int \rho_t(\theta) d\theta = 1$.

In the following, we first obtain the following lower bound on Z_t :

$$(B.2) \quad Z_t \geq e^{-Ke^{-t}(1+B)},$$

where the constants K, B are from eq. (4.5) and eq. (4.6). In fact, using our assumptions eq. (4.5) and eq. (4.6), we have

$$\begin{aligned} Z_t &= \int \left(\frac{\rho_0(\theta)}{\rho_{\text{post}}(\theta)} \right)^{e^{-t}} \rho_{\text{post}}(\theta) d\theta \geq \int \left(e^{-K(1+|\theta|^2)} \right)^{e^{-t}} \rho_{\text{post}}(\theta) d\theta \\ &= \int e^{-Ke^{-t}(1+|\theta|^2)} \rho_{\text{post}}(\theta) d\theta \geq e^{\int -Ke^{-t}(1+|\theta|^2) \rho_{\text{post}}(\theta) d\theta} \geq e^{-Ke^{-t}(1+B)}, \end{aligned}$$

where in the second to last inequalities, we used Jensen's inequality and the fact that e^x is convex.

By plugging eq. (4.6) and eq. (B.2) into eq. (B.1), we get

$$(B.3) \quad \frac{\rho_t(\theta)}{\rho_{\text{post}}(\theta)} = \frac{1}{Z_t} \left(\frac{\rho_0(\theta)}{\rho_{\text{post}}(\theta)} \right)^{e^{-t}} \leq e^{Ke^{-t}(1+B)} e^{Ke^{-t}(1+|\theta|^2)} = e^{Ke^{-t}(2+B+|\theta|^2)}.$$

Using eq. (B.3), we get the following upper bound on the KL divergence:

$$\begin{aligned} (B.4) \quad & \text{KL}[\rho_t \parallel \rho_{\text{post}}] \\ &= \int \rho_t(\theta) \log \frac{\rho_t(\theta)}{\rho_{\text{post}}(\theta)} d\theta \leq \int \rho_t(\theta) \log(e^{Ke^{-t}(2+B+|\theta|^2)}) d\theta \\ &= \int \rho_t(\theta) Ke^{-t}(2+B+|\theta|^2) d\theta = Ke^{-t} \left((2+B) + \int |\theta|^2 \rho_t(\theta) d\theta \right). \end{aligned}$$

For the last integral in eq. (B.4), using eq. (B.1) and the Hölder inequality, we can rewrite it as

$$\begin{aligned} (B.5) \quad & \frac{1}{Z_t} \int |\theta|^2 \rho_0(\theta)^{e^{-t}} \rho_{\text{post}}(\theta)^{1-e^{-t}} d\theta = \frac{1}{Z_t} \int (|\theta|^2 \rho_0(\theta))^{e^{-t}} (|\theta|^2 \rho_{\text{post}}(\theta))^{1-e^{-t}} d\theta \\ & \leq \frac{1}{Z_t} \left(\int |\theta|^2 \rho_0(\theta) d\theta \right)^{e^{-t}} \left(\int |\theta|^2 \rho_{\text{post}}(\theta) d\theta \right)^{1-e^{-t}} \leq Be^{Ke^{-t}(1+B)}, \end{aligned}$$

where for the last inequality we used eq. (B.2) and our assumption eq. (4.6).

Combining eq. (B.4) and eq. (B.5) together, for $t \geq \log((1+B)K)$, we have

$$(B.6) \quad \text{KL}[\rho_t \parallel \rho_{\text{post}}] \leq K e^{-t} (2 + B + B e^{K e^{-t}(1+B)}) \leq (2 + B + eB) K e^{-t}.$$

This completes the proof of Theorem 4.1. $\square$

APPENDIX C. INTUITIONS REGARDING THE DEFINITIONS OF VARIOUS METRIC TENSORS

C.1. The Fisher-Rao Riemannian metric. Writing tangent vectors on a multiplicative scale, by setting $\sigma = \rho \psi_\sigma$, we see that this metric may be written as $g_\rho^{\text{FR}}(\sigma_1, \sigma_2) = \int \psi_{\sigma_1} \psi_{\sigma_2} \rho d\theta$, and hence that in the ψ_σ variable the metric is described via the L_ρ^2 inner-product. That is, the Fisher-Rao Riemannian metric measures the multiplicative factor via the L_ρ^2 energy. In Appendix C.2, we will see that another important metric, the Wasserstein Riemannian metric, may also be understood as an L_ρ^2 measurement, but of the velocity field instead.

C.2. The Wasserstein Riemannian metric. The Wasserstein metric [70, 88, 105] has been widely used to construct gradient flows. Define ψ_σ to be the solution of the PDE

$$(C.1) \quad -\nabla_\theta \cdot (\rho \nabla_\theta \psi_\sigma) = \sigma.$$

This definition requires specification of function spaces to ensure unique invertibility of the divergence form elliptic operator; here we use it formally to derive flow equations. The Wasserstein Riemannian metric has the form

$$(C.2) \quad g_\rho^{\text{W}}(\sigma_1, \sigma_2) = \int \rho(\theta) \nabla_\theta \psi_{\sigma_1}(\theta)^T \nabla_\theta \psi_{\sigma_2}(\theta) d\theta$$

which defines the Wasserstein Riemannian metric tensor in (4.14).

The Wasserstein Riemannian metric has a transport interpretation. To understand this fix $\sigma \in T_\rho \mathcal{P}$ and consider the family of *velocity fields* v related to σ via the constraint $\sigma = -\nabla_\theta \cdot (\rho v)$. Then define $v_\sigma = \arg \min_v \int \rho |v|^2$ in which the minimization is over all v satisfying the constraint. A formal Lagrange multiplier argument can be used to deduce that $v_\sigma = \nabla_\theta \psi_\sigma$ for some ψ_σ . This motivates the relationship appearing in eq. (C.1) as well as the form of the Wasserstein Riemannian metric appearing in eq. (C.2) which may then be viewed as measuring the kinetic energy $\int \rho |v_\sigma|^2 d\theta$. We emphasize that, for ease of understanding, our discussion on the Riemannian structure of the Wasserstein metric is purely formal; for rigorous treatments, the reader can consult [3].

C.3. The Stein Riemannian metric. The Stein gradient flow [48, 83, 89] is the continuous limit of the Stein variational gradient descent [84]. Let ψ_σ solve the integro-partial differential equation

$$(C.3) \quad -\nabla_\theta \cdot \left(\rho(\theta) \int \kappa(\theta, \theta', \rho) \rho(\theta') \nabla_{\theta'} \psi_\sigma(\theta') d\theta' \right) = \sigma(\theta).$$

Here $\kappa : \mathbb{R}^{N_\theta} \times \mathbb{R}^{N_\theta} \times \mathcal{P} \rightarrow \mathbb{R}$ is a positive definite kernel for any fixed ρ . As before suitable function spaces need to be specified to ensure that this equation is

uniquely solvable; here we will treat it formally to derive flow equations. The Stein Riemannian metric is

$$(C.4) \quad g_\rho^S(\sigma_1, \sigma_2) = \int \int \kappa(\theta, \theta', \rho) \rho(\theta) \nabla_{\theta'} \psi_{\sigma_1}(\theta)^T \nabla_{\theta'} \psi_{\sigma_2}(\theta') \rho(\theta') d\theta d\theta'.$$

which defines the Stein Riemannian metric tensor in (4.19).

As in the Wasserstein setting, the Stein Riemannian metric [83] also has a transport interpretation. The Stein Riemannian metric identifies, for each $\sigma \in T_\rho \mathcal{P}$, the set of velocity fields v satisfying the constraint $\sigma = -\nabla_\theta \cdot (\rho v)$. Then $v_\sigma = \arg \min_v \|v\|_{\mathcal{H}_\kappa}^2$, with minimization over all v satisfying the constraint, and where $\mathcal{H}_\kappa$ is a Reproducing Kernel Hilbert Space (RKHS) with kernel κ . A formal Lagrangian multiplier argument shows that

$$v_\sigma = \int \kappa(\theta, \theta', \rho) \rho(\theta') \nabla_{\theta'} \psi_\sigma(\theta') d\theta'$$

for some ψ_σ . The Stein metric measures this transport change via the RKHS norm $\|v_\sigma\|_{\mathcal{H}_\kappa}^2$, leading to the interpretation that the Stein Riemannian metric can be written in the form

$$g_\rho^S(\sigma_1, \sigma_2) = \langle v_{\sigma_1}, v_{\sigma_2} \rangle_{\mathcal{H}_\kappa}.$$

APPENDIX D. PROOFS FOR AFFINE INVARIANCE OF GRADIENT FLOWS

D.1. Preliminaries. In this section, we present a lemma concerning the change-of-variable formula under the gradient and divergence operators, which is useful for our proofs in later subsections.

Lemma D.1. *Consider any invertible affine mapping from $\theta \in \mathbb{R}^{N_\theta}$ to $\tilde{\theta} \in \mathbb{R}^{N_\theta}$ defined by $\tilde{\theta} = \varphi(\theta) = A\theta + b$. For any differentiable scalar field $f : \mathbb{R}^{N_\theta} \rightarrow \mathbb{R}$ and vector field $g : \mathbb{R}^{N_\theta} \rightarrow \mathbb{R}^{N_\theta}$, we have*

$$\nabla_\theta f(\theta) = A^T \nabla_{\tilde{\theta}} \tilde{f}(\tilde{\theta}) \quad \text{and} \quad \nabla_\theta \cdot g(\theta) = \nabla_{\tilde{\theta}} \cdot (A\tilde{g}(\tilde{\theta})),$$

where $\tilde{f}(\tilde{\theta}) := f(A^{-1}(\tilde{\theta} - b))$ and $\tilde{g}(\tilde{\theta}) := g(A^{-1}(\tilde{\theta} - b))$.

Proof. Note that $\tilde{f}(\tilde{\theta}) = f(\theta)$ and $\tilde{g}(\tilde{\theta}) = g(\theta)$. By direct calculations, we get

$$\begin{aligned} [\nabla_\theta f(\theta)]_i &= \frac{\partial f(\theta)}{\partial \theta_i} = \sum_j \frac{\partial f(\theta)}{\partial \tilde{\theta}_j} \frac{\partial \tilde{\theta}_j}{\partial \theta_i} = \sum_j \frac{\partial \tilde{f}(\tilde{\theta})}{\partial \tilde{\theta}_j} A_{ji} = [A^T \nabla_{\tilde{\theta}} \tilde{f}(\tilde{\theta})]_i, \\ \nabla_\theta \cdot g(\theta) &= \sum_i \frac{\partial g_i(\theta)}{\partial \theta_i} = \sum_i \sum_j \frac{\partial g_i(\theta)}{\partial \tilde{\theta}_j} \frac{\partial \tilde{\theta}_j}{\partial \theta_i} = \sum_i \sum_j \frac{\partial \tilde{g}_i(\tilde{\theta})}{\partial \tilde{\theta}_j} A_{ji} = \nabla_{\tilde{\theta}} \cdot (A\tilde{g}(\tilde{\theta})). \end{aligned}$$

This completes the proof. $\square$

Remark D.2. Since $\tilde{f}(\tilde{\theta}) = f(\theta)$ and $\tilde{g}(\tilde{\theta}) = g(\theta)$, we can also summarize the result in Lemma D.1 as $\nabla_\theta f = A^T \nabla_{\tilde{\theta}} \tilde{f}$ and $\nabla_\theta \cdot g = \nabla_{\tilde{\theta}} \cdot (A\tilde{g})$.

D.2. Proof of Theorem 4.6.

Proof. We write down the form of the affine invariant Wasserstein gradient flow as follows:

$$(D.1) \quad \frac{\partial \rho_t(\theta)}{\partial t} = \nabla_\theta \cdot \left(\rho_t P(\theta, \rho_t) (\nabla_\theta \log \rho_t - \nabla_\theta \log \rho_{\text{post}}) \right).$$

Consider $\tilde{\theta} = \varphi(\theta) = A\theta + b$ and $\tilde{\rho}_t = \varphi\#\rho_t$ for an invertible affine transformation φ . Then, it suffices to show that under the assumption $P(\tilde{\theta}, \tilde{\rho}) = AP(\theta, \rho)A^T$, one has

$$(D.2) \quad \frac{\partial \tilde{\rho}_t(\tilde{\theta})}{\partial t} = \nabla_{\tilde{\theta}} \cdot \left(\tilde{\rho}_t P(\tilde{\theta}, \tilde{\rho}_t) (\nabla_{\tilde{\theta}} \log \tilde{\rho}_t - \nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}}) \right).$$

First, for the left hand side, by definition, we have

$$(D.3) \quad \tilde{\rho}_t(\tilde{\theta}) = \rho_t(\varphi^{-1}(\tilde{\theta})) |\nabla_{\tilde{\theta}} \varphi^{-1}(\tilde{\theta})| = \rho_t(\theta) |A^{-1}|.$$

For the right hand side, the chain rule leads to that

$$(D.4) \quad \nabla_{\tilde{\theta}} \log \tilde{\rho}_t = A^{-T} \nabla_{\theta} \log \rho_t \quad \nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}} = A^{-T} \nabla_{\theta} \log \rho_{\text{post}}.$$

Therefore, we can write the right hand side of eq. (D.2) as

$$(D.5) \quad \begin{aligned} & \nabla_{\tilde{\theta}} \cdot \left(\tilde{\rho}_t P(\tilde{\theta}, \tilde{\rho}_t) (\nabla_{\tilde{\theta}} \log \tilde{\rho}_t - \nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}}) \right) \\ &= \nabla_{\tilde{\theta}} \cdot \left(\tilde{\rho}_t P(\tilde{\theta}, \tilde{\rho}_t) A^{-T} (\nabla_{\theta} \log \rho_t - \nabla_{\theta} \log \rho_{\text{post}}) \right) \\ &= \nabla_{\theta} \cdot \left(\tilde{\rho}_t A^{-1} P(\tilde{\theta}, \tilde{\rho}_t) A^{-T} (\nabla_{\theta} \log \rho_t - \nabla_{\theta} \log \rho_{\text{post}}) \right) \\ &= \nabla_{\theta} \cdot \left(\rho_t A^{-1} P(\tilde{\theta}, \tilde{\rho}_t) A^{-T} (\nabla_{\theta} \log \rho_t - \nabla_{\theta} \log \rho_{\text{post}}) \right) \cdot |A^{-1}|, \end{aligned}$$

where in the second equality, we used eq. (D.4), and in the third equality, we used Lemma D.1. Based on eq. (D.3), eq. (D.5) and eq. (D.1), a sufficient condition for eq. (D.2) to hold is $A^{-1}P(\tilde{\theta}, \tilde{\rho})A^{-T} = P(\theta, \rho)$, or equivalently, $P(\tilde{\theta}, \tilde{\rho}) = AP(\theta, \rho)A^T$. This completes the proof. $\square$

D.3. Proof for the Fokker-Planck equation. By definition, the Fokker-Planck equation associated with (4.17) follows

$$(D.6) \quad \frac{\partial \rho_t(\theta)}{\partial t} = -\nabla_{\theta} \cdot (\rho_t b_t) + \nabla_{\theta} \cdot (\nabla_{\theta} \cdot (D \rho_t)),$$

where

$$\begin{aligned} b_t &= P(\theta, \rho_t) \nabla_{\theta} \log \rho_{\text{post}} + \left((D(\theta, \rho_t) - P(\theta, \rho_t)) \nabla_{\theta} \log \rho_t(\theta) + d(\theta, \rho_t) \right), \\ D &= \frac{1}{2} h(\theta, \rho_t) h(\theta, \rho_t)^T. \end{aligned}$$

Note that in the above, we use the notation where for a matrix-valued function $A(\theta) \in \mathbb{R}^{N_{\theta} \times N_{\theta}}$, the divergence $v = \nabla \cdot A \in \mathbb{R}^{N_{\theta}}$ is defined such that for each component $v_i = \sum_{j=1}^{N_{\theta}} \partial_j A_{ij}$, where $1 \leq i \leq N_{\theta}$.

Using the definitions of b_t , D and $d = \nabla_{\theta} \cdot D$, we have

$$(D.7) \quad \begin{aligned} -\nabla_{\theta} \cdot (\rho_t b_t) &= \nabla_{\theta} \cdot \left(\rho_t P(\theta, \rho_t) (\nabla_{\theta} \log \rho_t - \nabla_{\theta} \log \rho_{\text{post}}) \right) \\ &\quad - \nabla_{\theta} \cdot \left(\rho_t D(\theta, \rho_t) \nabla_{\theta} \log \rho_t(\theta) + \rho_t d(\theta, \rho_t) \right), \\ \nabla_{\theta} \cdot (\nabla_{\theta} \cdot (D \rho_t)) &= \nabla_{\theta} \cdot \left(\nabla_{\theta} \cdot D(\theta, \rho_t) \rho_t + D(\theta, \rho_t) \nabla_{\theta} \rho_t \right) \\ &= \nabla_{\theta} \cdot \left(\nabla_{\theta} \cdot D(\theta, \rho_t) \rho_t + D(\theta, \rho_t) \rho_t \nabla_{\theta} \log \rho_t \right) \\ &= \nabla_{\theta} \cdot \left(d(\theta, \rho_t) \rho_t + D(\theta, \rho_t) \rho_t \nabla_{\theta} \log \rho_t \right). \end{aligned}$$

Bringing eq. (D.7) into eq. (D.6) leads to (4.16).

D.4. Proof of Theorem 4.7.

Proof. Consider the invertible affine transformation $\tilde{\theta} = \varphi(\theta) = A\theta + b$ and correspondingly $\tilde{\rho} = \varphi\#\rho$ and $\tilde{\rho}_{\text{post}} = \varphi\#\rho_{\text{post}}$. We aim to verify that the mean-field equation is affine invariant according to Definition 4.3, where

$$f(\theta, \rho, \rho_{\text{post}}) = P(\theta, \rho)\nabla_{\theta} \log \rho_{\text{post}}(\theta) + (D(\theta, \rho) - P(\theta, \rho))\nabla_{\theta} \log \rho(\theta) + d(\theta, \rho).$$

Using eq. (4.11b), we get

$$D(\tilde{\theta}, \tilde{\rho}) = \frac{1}{2}h(\tilde{\theta}, \tilde{\rho})h(\tilde{\theta}, \tilde{\rho})^T = \frac{1}{2}Ah(\theta, \rho)h(\theta, \rho)^T A^T = AD(\theta, \rho)A^T.$$

Similarly, it holds that $d(\tilde{\theta}, \tilde{\rho}) = Ad(\theta, \rho)$. Based on these relations, we can calculate as follows:

$$\begin{aligned} & (D.8) \\ & Af(\theta, \rho, \rho_{\text{post}}) \\ &= AP(\theta, \rho)\nabla_{\theta} \log \rho_{\text{post}}(\theta) + A(D(\theta, \rho) - P(\theta, \rho))\nabla_{\theta} \log \rho(\theta) + Ad(\theta, \rho) \\ &= AP(\theta, \rho)A^T\nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}}(\tilde{\theta}) + (AD(\theta, \rho) - AP(\theta, \rho))A^T\nabla_{\tilde{\theta}} \log \tilde{\rho}(\tilde{\theta}) + Ad(\theta, \rho) \\ &= P(\tilde{\theta}, \tilde{\rho})\nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}}(\tilde{\theta}) + (D(\tilde{\theta}, \tilde{\rho}) - P(\tilde{\theta}, \tilde{\rho}))\nabla_{\tilde{\theta}} \log \tilde{\rho}(\tilde{\theta}) + d(\tilde{\theta}, \tilde{\rho}) \\ &= f(\tilde{\theta}, \tilde{\rho}, \tilde{\rho}_{\text{post}}). \end{aligned}$$

The first equality is by definition. In the second equality, we used $A^T\nabla_{\tilde{\theta}} \log \tilde{\rho}(\tilde{\theta}) = \nabla_{\theta} \log \rho(\theta)$. In the third equality, we used the condition in Theorem 4.6. With this result, the mean-field equation is affine invariant and the proof is complete. $\square$

D.5. Proof of Theorem 4.10.

Proof. The proof is similar to that in Appendix D.2. The affine invariant Stein gradient flow has the form

$$(D.9) \quad \frac{\partial \rho_t(\theta)}{\partial t} = \nabla_{\theta} \cdot \left[\rho_t(\theta) \int \kappa(\theta, \theta', \rho_t) \rho_t(\theta') P(\theta, \theta', \rho_t) \nabla_{\theta'} (\log \rho_t(\theta') - \log \rho_{\text{post}}(\theta')) d\theta' \right].$$

Consider $\tilde{\theta} = \varphi(\theta) = A\theta + b$ and $\tilde{\rho}_t = \varphi\#\rho_t$ for an invertible affine transformation φ . Then, it suffices to show that under the assumed condition, one has

$$(D.10) \quad \frac{\partial \tilde{\rho}_t(\tilde{\theta})}{\partial t} = \nabla_{\tilde{\theta}} \cdot \left[\tilde{\rho}_t(\tilde{\theta}) \int \kappa(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}_t) \tilde{\rho}_t(\tilde{\theta}') P(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}_t) \nabla_{\tilde{\theta}'} (\log \tilde{\rho}_t(\tilde{\theta}') - \log \tilde{\rho}_{\text{post}}(\tilde{\theta}')) d\tilde{\theta}' \right].$$

For the right hand side of eq. (D.10), we have

$$\begin{aligned} & (D.11) \\ & \nabla_{\tilde{\theta}} \cdot \left[\tilde{\rho}_t(\tilde{\theta}) \int \kappa(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}_t) \tilde{\rho}_t(\tilde{\theta}') P(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}_t) \nabla_{\tilde{\theta}'} (\log \tilde{\rho}_t(\tilde{\theta}') - \log \tilde{\rho}_{\text{post}}(\tilde{\theta}')) d\tilde{\theta}' \right] \\ &= \nabla_{\tilde{\theta}} \cdot \left[\tilde{\rho}_t(\tilde{\theta}) \int \kappa(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}_t) \tilde{\rho}_t(\tilde{\theta}') P(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}_t) A^{-T} \nabla_{\theta'} (\log \rho_t(\theta') - \log \rho_{\text{post}}(\theta')) d\tilde{\theta}' \right] \\ &= \nabla_{\tilde{\theta}} \cdot \left[\tilde{\rho}_t(\tilde{\theta}) \int \kappa(\tilde{\theta}, \varphi(\theta'), \tilde{\rho}_t) \rho_t(\theta') P(\tilde{\theta}, \varphi(\theta'), \tilde{\rho}_t) A^{-T} \nabla_{\theta'} (\log \rho_t(\theta') - \log \rho_{\text{post}}(\theta')) d\theta' \right] \\ &= \nabla_{\theta} \cdot \left[\tilde{\rho}_t(\tilde{\theta}) \int \kappa(\tilde{\theta}, \varphi(\theta'), \tilde{\rho}_t) \rho_t(\theta') A^{-1} P(\tilde{\theta}, \varphi(\theta'), \tilde{\rho}_t) A^{-T} \nabla_{\theta'} (\log \rho_t(\theta') - \log \rho_{\text{post}}(\theta')) d\theta' \right] \\ &= (\nabla_{\theta} \cdot f) \cdot |A^{-1}|, \\ & f = \rho_t(\theta) \int \kappa(\tilde{\theta}, \varphi(\theta'), \tilde{\rho}_t) \rho_t(\theta') A^{-1} P(\tilde{\theta}, \varphi(\theta'), \tilde{\rho}_t) A^{-T} \nabla_{\theta'} (\log \rho_t(\theta') - \log \rho_{\text{post}}(\theta')) d\theta', \end{aligned}$$

where in the first equality, we used eq. (D.4); in the second equality, we changed coordinates in the integral $\tilde{\theta}' = \varphi(\theta')$; in the third equality, we used Lemma D.1 about g , and in the last equality, we used eq. (D.3).

By eq. (D.3), eq. (D.11) and eq. (D.9), a sufficient condition for eq. (D.10) to hold is

$$\kappa(\tilde{\theta}, \varphi(\theta'), \tilde{\rho})A^{-1}P(\tilde{\theta}, \varphi(\theta'), \tilde{\rho})A^{-T} = \kappa(\theta, \theta', \rho)P(\theta, \theta', \rho),$$

or equivalently,

$$\kappa(\tilde{\theta}, \tilde{\theta}', \tilde{\rho})P(\tilde{\theta}, \tilde{\theta}', \tilde{\rho}) = \kappa(\theta, \theta', \rho)AP(\theta, \theta', \rho)A^T,$$

where $\tilde{\theta}' = \varphi(\theta')$. This completes the proof. $\square$

D.6. Proof of Theorem 4.11.

Proof. Consider the invertible affine transformation $\tilde{\theta} = \varphi(\theta) = A\theta + b$ and correspondingly $\tilde{\rho} = \varphi\#\rho$, $\tilde{\rho}_{\text{post}} = \varphi\#\rho_{\text{post}}$. By direct calculations, we get

$$\begin{aligned} & Af(\theta, \rho, \rho_{\text{post}}) \\ &= \int \kappa(\theta, \theta', \rho)AP(\theta, \theta', \rho)\nabla_{\theta'}(\log \rho(\theta') - \log \rho_{\text{post}}(\theta'))\rho(\theta')d\theta' \\ (D.12) \quad &= \int \kappa(\theta, \theta', \rho)AP(\theta, \theta', \rho)A^T(\nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}}(\tilde{\theta}') - \nabla_{\tilde{\theta}} \log \tilde{\rho}(\tilde{\theta}'))\rho(\theta')d\theta' \\ &= \int \kappa(\tilde{\theta}, \tilde{\theta}', \tilde{\rho})P(\tilde{\theta}, \tilde{\theta}', \tilde{\rho})(\nabla_{\tilde{\theta}} \log \tilde{\rho}_{\text{post}}(\tilde{\theta}') - \nabla_{\tilde{\theta}} \log \tilde{\rho}(\tilde{\theta}'))\tilde{\rho}(\tilde{\theta}')d\tilde{\theta}' \\ &= f(\tilde{\theta}, \tilde{\rho}, \tilde{\rho}_{\text{post}}). \end{aligned}$$

The first equality is by definition. In the second equality, we used $A^T\nabla_{\tilde{\theta}} \log \tilde{\rho}(\tilde{\theta}) = \nabla_{\theta} \log \rho(\theta)$. In the third equality, we used the relation $\rho(\theta') = \tilde{\rho}(\tilde{\theta}')|A|$ due to eq. (D.3), and $d\tilde{\theta}' = |A|d\theta'$; we also used the condition in Theorem 4.10. With this result, the mean-field equation is affine invariant (Definition 4.3). The proof is complete. $\square$

APPENDIX E. PROOFS FOR GAUSSIAN APPROXIMATE GRADIENT FLOWS

E.1. Preliminaries. We start with the following Stein's identities concerning the Gaussian density function ρ_a .

Lemma E.1. Assume $\theta \sim \mathcal{N}(m, C)$ with density $\rho_a(\theta) = \rho_a(\theta; m, C)$, we have

$$(E.1) \quad \nabla_m \rho_a(\theta) = -\nabla_{\theta} \rho_a(\theta) \quad \text{and} \quad \nabla_C \rho_a(\theta) = \frac{1}{2} \nabla_{\theta} \nabla_{\theta} \rho_a(\theta).$$

Furthermore, for any scalar field $f : \mathbb{R}^{N_{\theta}} \rightarrow \mathbb{R}$ and vector field $g : \mathbb{R}^{N_{\theta}} \rightarrow \mathbb{R}^{N_{\theta}}$, we have

$$\begin{aligned} (E.2) \quad & \mathbb{E}_{\rho_a}[\nabla_{\theta} g(\theta)] = \nabla_m \mathbb{E}_{\rho_a}[g(\theta)] = \text{Cov}[g(\theta), \theta]C^{-1}, \\ & \mathbb{E}_{\rho_a}[\nabla_{\theta} \nabla_{\theta} f(\theta)] = \text{Cov}[\nabla_{\theta} f(\theta), \theta]C^{-1} = -C^{-1} \mathbb{E}_{\rho_a}[(C - (\theta - m)(\theta - m)^T)f]C^{-1}. \end{aligned}$$

Proof. For Gaussian density function ρ_a , we have

$$\begin{aligned}\nabla_m \rho_a(\theta) &= \nabla_m \frac{1}{\sqrt{|2\pi C|}} \exp\left\{-\frac{1}{2}(\theta - m)C^{-1}(\theta - m)\right\} \\ &= C^{-1}(\theta - m)\rho_a(\theta) = -\nabla_\theta \rho_a(\theta), \\ \nabla_C \rho_a(\theta) &= \rho_a(\theta) \left(-\frac{1}{2} \frac{\partial \log |C|}{\partial C} - \frac{1}{2} \frac{\partial(\theta - m)^T C^{-1}(\theta - m)}{\partial C} \right) \\ &= -\frac{1}{2} \rho_a(\theta) \left(C^{-1} - C^{-1}(\theta - m)(\theta - m)^T C^{-1} \right) = \frac{1}{2} \nabla_\theta \nabla_\theta \rho_a(\theta).\end{aligned}$$

For any scalar field $f(\theta)$ and vector field $g(\theta)$, we have
(E.3)

$$\begin{aligned}\mathbb{E}_{\rho_a}[\nabla_\theta g(\theta)] &= \int \nabla_\theta g(\theta) \rho_a(\theta) d\theta = - \int g(\theta) \nabla_\theta \rho_a(\theta)^T d\theta = \nabla_m \mathbb{E}_{\rho_a}[g(\theta)] \\ &= \int g(\theta)(\theta - m)^T C^{-1} \rho_a(\theta) d\theta = \text{Cov}[g(\theta), \theta] C^{-1}, \\ \mathbb{E}_{\rho_a}[\nabla_\theta \nabla_\theta f(\theta)] &= \int \nabla_\theta \nabla_\theta f(\theta) \rho_a(\theta) d\theta = - \int \nabla_\theta f(\theta) \nabla_\theta \rho_a(\theta)^T d\theta \\ &= \int \nabla_\theta f(\theta)(\theta - m)^T C^{-1} \rho_a(\theta) d\theta = \text{Cov}[\nabla_\theta f(\theta), \theta] C^{-1} \\ &= \int f(\theta) \nabla_\theta \nabla_\theta \rho_a(\theta) d\theta \\ &= -C^{-1} \mathbb{E}_{\rho_a} \left[(C - (\theta - m)(\theta - m)^T) f \right] C^{-1}.\end{aligned}$$

□

Lemma E.2 is proved in [77, Theorem 1]:

Lemma E.2. *Consider the KL divergence*

$$\text{KL}[\rho_a \parallel \rho_{\text{post}}] = -\frac{1}{2} \log |C| - \int \rho_a(\theta) \log \rho_{\text{post}}(\theta) d\theta + \text{const.}$$

For a fixed ρ_{post} , any minimizer of this divergence over the space $\mathcal{P}^G$, denoted by $\rho_{a_\star}(\theta) = \mathcal{N}(\theta, m_\star, C_\star)$ with $a_\star = (m_\star, C_\star)$, satisfies

$$\mathbb{E}_{\rho_{a_\star}}[\nabla_\theta \log \rho_{\text{post}}(\theta)] = 0 \quad \text{and} \quad C_\star^{-1} = -\mathbb{E}_{\rho_{a_\star}}[\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)].$$

E.2. Equivalence of two Gaussian approximation approaches in Section 5.1.

Here we prove Theorem 5.1 and Theorem 5.2, which identifies specific gradient flows, within the paper, that satisfy the assumptions required for application of Theorem 5.1.

Proof of Theorem 5.1. Any element in the tangent space of the Gaussian density space $\mathcal{P}^G$, $T_{\rho_a} \mathcal{P}^G$ is given in the form $\nabla_a \rho_a \cdot \sigma$, where $\sigma \in \mathbb{R}^{N_a}$.

The density evolution equation of eq. (5.5) is

$$(E.4) \quad \frac{\partial \rho_{a_t}}{\partial t} = -\nabla_a \rho_{a_t} \cdot \mathfrak{M}(a_t)^{-1} \frac{\partial \mathcal{E}(\rho_a; \rho_{\text{post}})}{\partial a} \Big|_{a=a_t},$$

where $\rho_{a_t} = \mathcal{N}(m_t, C_t)$. We will prove that its evolution equations of m_t and C_t are eq. (5.6).

Using assumption (eq. (5.8)), we know that for any $f(\theta) \in \text{span}\{\theta_i, \theta_i \theta_j, 1 \leq i, j \leq N_\theta\}$, there exists a $\sigma \in \mathbb{R}^{N_a}$, such that

$$f(\theta) = M(\rho_a) \nabla_a \rho_a \cdot \sigma.$$

Then we have

$$\begin{aligned} \left\langle M(\rho_a)^{-1} \frac{\delta \mathcal{E}(\rho)}{\delta \rho} \Big|_{\rho=\rho_a}, f(\theta) \right\rangle &= \left\langle M(\rho_a)^{-1} \frac{\delta \mathcal{E}(\rho)}{\delta \rho} \Big|_{\rho=\rho_a}, M(\rho_a) \nabla_a \rho_a \cdot \sigma \right\rangle \\ (E.5) \quad &= \left\langle \frac{\delta \mathcal{E}(\rho)}{\delta \rho} \Big|_{\rho=\rho_a}, \nabla_a \rho_a \cdot \sigma \right\rangle \end{aligned}$$

and

$$\begin{aligned} &\left\langle \nabla_a \rho_a \cdot \mathfrak{M}(a)^{-1} \frac{\partial \mathcal{E}(\rho_a; \rho_{\text{post}})}{\partial a}, f(\theta) \right\rangle \\ &= \left\langle \nabla_a \rho_a \cdot \mathfrak{M}(a)^{-1} \frac{\partial \mathcal{E}(\rho_a; \rho_{\text{post}})}{\partial a}, M(\rho_a) \nabla_a \rho_a \cdot \sigma \right\rangle \\ (E.6) \quad &= \left\langle M(\rho_a) \nabla_a \rho_a \cdot \mathfrak{M}(a)^{-1} \frac{\partial \mathcal{E}(\rho_a; \rho_{\text{post}})}{\partial a}, \nabla_a \rho_a \cdot \sigma \right\rangle \\ &= \left\langle \frac{\partial \mathcal{E}(\rho_a; \rho_{\text{post}})}{\partial a}, \sigma \right\rangle \\ &= \left\langle \frac{\delta \mathcal{E}(\rho)}{\delta \rho} \Big|_{\rho=\rho_a}, \nabla_a \rho_a \cdot \sigma \right\rangle. \end{aligned}$$

Here we used the definition of $\mathfrak{M}(a)$ in eq. (5.4) for the third equality. Combining eq. (E.5) and eq. (E.6), we have that the mean and covariance evolution equations of eq. (E.4) are eq. (5.7). $\square$

Proof of Theorem 5.2. Using the calculation in the proof of Lemma E.1, the tangent space of the Gaussian density manifold at ρ_a with $a = (m, C)$ is

$$\begin{aligned} T_{\rho_a} \mathcal{P}^G &= \text{span}\{\rho_a [C^{-1}(\theta - m)]_i, \rho_a [C^{-1}(\theta - m)(\theta - m)^T C^{-1} - C^{-1}]_{ij}\} \\ (E.7) \quad &= \text{span}\{\rho_a(\theta_i - \mathbb{E}_{\rho_a}[\theta_i]), \rho_a(\theta_i \theta_j - \mathbb{E}_{\rho_a}[\theta_i \theta_j])\}, \end{aligned}$$

and $1 \leq i, j \leq N_\theta$.

For the Fisher-Rao metric, we have

$$(E.8) \quad M^{\text{FR}}(\rho_a)^{-1} \text{span}\{\theta_i, \theta_i \theta_j\} = \text{span}\{\rho_a(\theta_i - \mathbb{E}_{\rho_a}[\theta_i]), \rho_a(\theta_i \theta_j - \mathbb{E}_{\rho_a}[\theta_i \theta_j])\} = T_{\rho_a} \mathcal{P}^G.$$

For the affine invariant Wasserstein metric with preconditioner P independent of θ , we have

$$\begin{aligned} (E.9) \quad &M^{\text{AIW}}(\rho_a)^{-1} \text{span}\{\theta_i, \theta_i \theta_j\} \\ &= \text{span}\{\nabla_\theta \cdot (\rho_a(\theta) P(\rho_a) e_i), \nabla_\theta \cdot (\rho_a(\theta) P(\rho_a) e_i \theta_j)\} \\ &= \text{span}\{\nabla_\theta \cdot (\rho_a(\theta)(b' + A'\theta)) \quad \forall b' \in \mathbb{R}^{N_\theta} \quad A' \in \mathbb{R}^{N_\theta \times N_\theta}\} \\ &= \text{span}\{\rho_a(\theta)[\text{tr}(A') - (\theta - m)^T C^{-1}(b' + A'\theta)] \quad \forall b' \in \mathbb{R}^{N_\theta} \quad A' \in \mathbb{R}^{N_\theta \times N_\theta}\} \\ &= T_{\rho_a} \mathcal{P}^G. \end{aligned}$$

Here e_i is the i -th unit vector.

For the affine invariant Stein metric with preconditioner P independent of θ and with a bilinear kernel $\kappa(\theta, \theta', \rho) = (\theta - m)^T A(\rho)(\theta' - m) + b(\rho)$ ($b \neq 0$, A non-singular), we have

$$\begin{aligned}
 & (E.10) \\
 & M^{\text{AIS}}(\rho_a)^{-1} \text{span}\{\theta_i, \theta_i \theta_j\} \\
 &= \text{span}\left\{ \nabla_\theta \cdot \left(\rho_a(\theta) P(\rho_a) \int \kappa \rho_a(\theta') e_i d\theta' \right), \nabla_\theta \cdot \left(\rho_a(\theta) P(\rho_a) \int \kappa \rho_a(\theta') e_i \theta'_j d\theta' \right) \right\} \\
 &= \text{span}\left\{ \nabla_\theta \cdot \left(\rho_a(\theta) P(\rho_a) b e_i \right), \nabla_\theta \cdot \left(\rho_a(\theta) P(\rho_a) [(\theta - m)^T A e_j c_{jj} e_i + b e_i m_j] \right) \right\} \\
 &= \text{span}\left\{ \nabla_\theta \cdot (\rho_a(\theta)(b' + A'\theta)) \quad \forall b' \in \mathbb{R}^{N_\theta} \quad A' \in \mathbb{R}^{N_\theta \times N_\theta} \right\} \\
 &= T_{\rho_a} \mathcal{P}^G.
 \end{aligned}$$

□

E.3. Analytical solutions for the Gaussian approximate Fisher-Rao gradient flow. The analytical solution for the Gaussian approximate Fisher-Rao gradient flow (5.9) to sample Gaussian distribution is in Lemma E.3:

Lemma E.3. *Assume the posterior distribution is Gaussian $\rho_{\text{post}}(\theta) \sim \mathcal{N}(m_\star, C_\star)$. Then the Gaussian approximate Fisher-Rao gradient flow (5.9) has the analytical solution*

$$(E.11a) \quad m_t = m_\star + e^{-t} \left((1 - e^{-t}) C_\star^{-1} + e^{-t} C_0^{-1} \right)^{-1} C_0^{-1} (m_0 - m_\star),$$

$$(E.11b) \quad C_t^{-1} = C_\star^{-1} + e^{-t} (C_0^{-1} - C_\star^{-1}).$$

The proof is in Appendix E.4. We remark that both mean and covariance converge exponentially fast to $m_\star$ and $C_\star$ with convergence rate $\mathcal{O}(e^{-t})$. This rate is independent of $m_\star$ and $C_\star$. The uniform convergence rate $\mathcal{O}(e^{-t})$ of the Gaussian approximate affine invariant Wasserstein gradient flow (5.13) is obtained in [57, Lemma 3.2] [18].

E.4. Proof of Lemma E.3.

Proof. Under the Gaussian posterior assumption and the fact that

$$\frac{dC_t^{-1}}{dt} = -C_t^{-1} \frac{dC_t}{dt} C_t^{-1},$$

the Gaussian approximate Fisher-Rao gradient flow eq. (5.9) becomes

$$\begin{aligned}
 (E.12) \quad & \frac{dm_t}{dt} = C_t C_\star^{-1} (m_\star - m_t), \\
 & \frac{dC_t^{-1}}{dt} = -C_t^{-1} + C_\star^{-1}.
 \end{aligned}$$

The covariance update equation has an analytical solution

$$(E.13) \quad C_t^{-1} = (1 - e^{-t}) C_\star^{-1} + e^{-t} C_0^{-1}.$$

We reformulate eq. (E.11a) as

$$m_\star - m_t = e^{-t} C_t C_0^{-1} (m_\star - m_0).$$

Computing its time derivative leads to

$$\begin{aligned}
 \frac{dm_t}{dt} &= -e^{-t}C_tC_0^{-1}(m_0 - m_\star) + e^{-t}\frac{dC_t}{dt}C_0^{-1}(m_0 - m_\star) \\
 &= -e^{-t}C_tC_0^{-1}(m_0 - m_\star) + e^{-t}(C_t - C_tC_\star^{-1}C_t)C_0^{-1}(m_0 - m_\star) \\
 &= C_tC_\star^{-1}e^{-t}C_tC_0^{-1}(m_\star - m_0) \\
 &= C_tC_\star^{-1}(m_\star - m_t).
 \end{aligned}
 \tag{E.14}$$

□

E.5. Proof of convergence for Gaussian posterior (Theorem 5.6).

Proof. Under the Gaussian posterior assumption, for $\theta_t \sim \mathcal{N}(m_t, C_t)$ with density ρ_{a_t} , we have

$$\begin{aligned}
 \mathbb{E}_{\rho_{a_t}}[\nabla_\theta \log \rho_{\text{post}}(\theta_t)] &= -C_\star^{-1}(m_t - m_\star), \quad \mathbb{E}_{\rho_{a_t}}[\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta_t)] = -C_\star^{-1}.
 \end{aligned}
 \tag{E.15}$$

Here $m_\star$ and $C_\star$ are the posterior mean and covariance given in eq. (5.2). We have explicit expressions for the Gaussian approximate gradient flows eqs. (5.9), (5.12) and (5.18):

$$\begin{aligned}
 \text{Gaussian approximate gradient flow:} \\
 \partial_t m_t &= -C_\star^{-1}(m_t - m_\star), \quad \partial_t C_t = \frac{1}{2}C_t^{-1} - \frac{1}{2}C_\star^{-1}.
 \end{aligned}
 \tag{E.16}$$

Gaussian approximate Fisher-Rao gradient flow:

$$\partial_t m_t = -C_tC_\star^{-1}(m_t - m_\star), \quad \partial_t C_t = C_t - C_tC_\star^{-1}C_t.$$

Gaussian approximate Wasserstein gradient flow:

$$\partial_t m_t = -C_\star^{-1}(m_t - m_\star), \quad \partial_t C_t = 2I - C_tC_\star^{-1} - C_\star^{-1}C_t.$$

For the dynamics of m_t for both Gaussian approximate gradient flow and Gaussian approximate Wasserstein gradient flow, we have

$$m_t - m_\star = e^{-tC_\star^{-1}}(m_0 - m_\star).
 \tag{E.17}$$

By taking the 2-norm on both sides, using $\|\cdot\|_2$ to denote both the vector and induced matrix norms, and recalling that the largest eigenvalue of $C_\star$ is $\lambda_{\star, \max}$, we obtain

$$\|m_t - m_\star\|_2 \leq \|e^{-tC_\star^{-1}}\|_2 \|m_0 - m_\star\|_2 \leq e^{-t/\lambda_{\star, \max}} \|m_0 - m_\star\|_2 = \mathcal{O}(e^{-t/\lambda_{\star, \max}}).
 \tag{E.18}$$

The bound can be achieved, when $m_0 - m_\star$ has non-zero component in the $C_\star$ eigenvector direction corresponding to $\lambda_{\star, \max}$.

For the Gaussian approximate Fisher-Rao gradient flow, thanks to the explicit formula eq. (E.11a), we find that

$$\|m_t - m_\star\|_2 \leq e^{-t} \max\{\|C_\star\|_2, \|C_0\|_2\} \|C_0^{-1}\|_2 \|m_0 - m_\star\|_2 = \mathcal{O}(e^{-t}).
 \tag{E.19}$$

The bound can be achieved when $m_0 - m_\star$ is non-zero.

Next, we analyze the dynamics of the covariance matrix C_t . Our initialization $C_0 = \lambda_0 I$ commutes with $C_\star, C_\star^{-1}$. It follows that C_t commutes with $C_\star, C_\star^{-1}$ for any $t \geq 0$ and all gradient flows in eq. (E.16), since 0 is the unique solution of the evolution ordinary differential equation of $C_tC_\star - C_\star C_t$. So we can diagonalize

$C_t, C_t^{-1}, C_*, C_*^{-1}$ simultaneously, and write down the dynamics of the eigenvalues of C_t . For any eigenvalue λ_t of C_t , it satisfies the following differential equations,

$$(E.20) \quad \begin{aligned} &\text{Gaussian approximate gradient flow: } \partial_t \lambda_t = \frac{1}{2\lambda_t} - \frac{1}{2\lambda_*}, \\ &\text{Gaussian approximate Fisher-Rao gradient flow: } \partial_t \lambda_t = \lambda_t - \lambda_t^2 \lambda_*^{-1}, \\ &\text{Gaussian approximate Wasserstein gradient flow: } \partial_t \lambda_t = 2 - \frac{2\lambda_t}{\lambda_*}, \end{aligned}$$

where λ_* is the corresponding eigenvalue of C_* . From eq. (E.20), we know that λ_t is bounded between λ_0 and λ_* . Moreover, the ordinary differential equations in eq. (E.20) can be solved explicitly.

For the Gaussian approximate gradient flow

$$(E.21) \quad \lambda_t - \lambda_* = (\lambda_0 - \lambda_*) e^{-\frac{t}{2\lambda_*^2} - \frac{\lambda_t - \lambda_0}{\lambda_*}}, \quad |\lambda_t - \lambda_*| = \mathcal{O}(e^{-t/2\lambda_*^2}).$$

Since the largest eigenvalue of C_* is $\lambda_{*,\max}$, we conclude that

$$\|C_t - C_*\|_2 = \mathcal{O}(e^{-t/2\lambda_{*,\max}^2}).$$

For the Gaussian approximate Fisher-Rao gradient flow,

$$(E.22) \quad \lambda_t = \frac{\lambda_*}{1 + \left(\frac{\lambda_*}{\lambda_0} - 1\right) e^{-t}}, \quad |\lambda_t - \lambda_*| = \mathcal{O}(e^{-t}).$$

It follows that

$$\|C_t - C_*\|_2 = \mathcal{O}(e^{-t}).$$

For the Gaussian approximate Wasserstein gradient flow

$$(E.23) \quad \lambda_t = \lambda_* + e^{-2t/\lambda_*}(\lambda_0 - \lambda_*), \quad |\lambda_t - \lambda_0| = \mathcal{O}(e^{-2t/\lambda_*}).$$

Since the large eigenvalue of C_* is $\lambda_{*,\max}$, we conclude that

$$\|C_t - C_*\|_2 = \mathcal{O}(e^{-2t/\lambda_{*,\max}}).$$

□

E.6. Proof of Theorem 5.7. Let ρ_{a_*} be $\mathcal{N}(m_*, C_*)$, the unique minimizer of eq. (5.2), noting that this is also the unique stationary point of Gaussian approximate Fisher-Rao gradient flow eq. (5.9); see Lemma E.2 for definition of a_* . It satisfies

$$(E.24) \quad \mathbb{E}_{\rho_{a_*}} [\nabla_\theta \log \rho_{\text{post}}(\theta)] = 0 \quad \text{and} \quad \mathbb{E}_{\rho_{a_*}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)] = C_*^{-1}.$$

For $N_\theta = 1$, we can calculate the linearized Jacobian matrix of the ODE system eq. (5.9) around (m_*, C_*) :

$$(E.25) \quad \begin{aligned} &\frac{\partial \text{RHS}}{\partial(m, C)} \Big|_{(m, C) = (m_*, C_*)} \\ &= \begin{bmatrix} -1 & -\frac{1}{2} \mathbb{E}_{\rho_{a_*}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)(\theta - m_*)] \\ -\mathbb{E}_{\rho_{a_*}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)(\theta - m_*)] C_* & -\frac{1}{2} - \frac{1}{2} \mathbb{E}_{\rho_{a_*}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)(\theta - m_*)^2] \end{bmatrix}, \end{aligned}$$

where RHS denotes the right-hand side of the ODE system eq. (5.9).

We further define

$$(E.26) \quad \begin{aligned} A_1 &:= \mathbb{E}_{\rho_{a_*}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)(\theta - m_*)], \\ A_2 &:= \mathbb{E}_{\rho_{a_*}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)(\theta - m_*)^2] \geq 0. \end{aligned}$$

Using the Cauchy-Schwarz inequality and eq. (E.24), we have

$$(E.27) \quad A_2 C_\star^{-1} = A_2 \mathbb{E}_{\rho_{a_\star}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)] \geq A_1^2.$$

By direct calculations, the two eigenvalues of eq. (E.25) satisfy

$$(E.28a) \quad \lambda_1 = \frac{-(\frac{3}{2} + \frac{1}{2}A_2) - \sqrt{(\frac{1}{2} - \frac{1}{2}A_2)^2 + 2A_1^2 C_\star}}{2} \leq -1,$$

$$(E.28b) \quad \lambda_2 = \frac{-(\frac{3}{2} + \frac{1}{2}A_2) + \sqrt{(\frac{1}{2} - \frac{1}{2}A_2)^2 + 2A_1^2 C_\star}}{2} \\ = -\frac{1 + A_2 - A_1^2 C_\star}{(\frac{3}{2} + \frac{1}{2}A_2) + \sqrt{(\frac{1}{2} - \frac{1}{2}A_2)^2 + 2A_1^2 C_\star}} \leq -\frac{1}{3 + A_2},$$

where in the last inequality, we have used (E.27). In the following, we prove bounds on λ_2 .

Step 1 (Upper bound). Since the upper bounds (E.28) of the two eigenvalues depend on A_2 , we will first prove that

$$(E.29) \quad A_2 \leq (4 + \frac{4}{\sqrt{\pi}}) \left(\log\left(\frac{\beta}{\alpha}\right) + 1 \right).$$

Without loss of generality, we assume $m_\star = 0$; otherwise we can always achieve this through a change of variable. Considering now only the right half of the integration (i.e. integration from 0 to $+\infty$) defining A_2 , we have

$$(E.30) \quad \int_0^{+\infty} -\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta) \theta^2 \frac{1}{\sqrt{2\pi C_\star}} e^{-\frac{1}{2} \frac{\theta^2}{C_\star}} d\theta \\ = \int_0^A -\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta) \theta^2 \frac{1}{\sqrt{2\pi C_\star}} e^{-\frac{1}{2} \frac{\theta^2}{C_\star}} d\theta \\ + \int_A^{+\infty} -\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta) \theta^2 \frac{1}{\sqrt{2\pi C_\star}} e^{-\frac{1}{2} \frac{\theta^2}{C_\star}} d\theta \\ \leq \frac{A^2}{C_\star} + \int_A^{+\infty} -\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta) \theta^2 \frac{1}{\sqrt{2\pi C_\star}} e^{-\frac{1}{2} \frac{\theta^2}{C_\star}} d\theta \quad (\text{Using eq. (E.24) and } \theta \leq A) \\ \leq \frac{A^2}{C_\star} + \beta \int_A^{+\infty} \frac{\theta^3}{A} \frac{1}{\sqrt{2\pi C_\star}} e^{-\frac{1}{2} \frac{\theta^2}{C_\star}} d\theta \quad (\text{Using } \theta \geq A \text{ and } \beta\text{-smoothness of } \log \rho_{\text{post}}) \\ = \frac{A^2}{C_\star} + \frac{2\beta C_\star^2}{A\sqrt{2\pi C_\star}} \left(\frac{A^2}{2C_\star} + 1 \right) e^{-\frac{A^2}{2C_\star}} \quad (\text{Direct calculation}) \\ \leq \frac{A^2}{C_\star} + \frac{\beta}{\sqrt{\pi}\alpha} \left(\frac{A}{\sqrt{2C_\star}} + \frac{\sqrt{2C_\star}}{A} \right) e^{-\frac{A^2}{2C_\star}} \quad (\text{Using } C_\star \leq \frac{1}{\alpha}) \\ \leq (2 + \frac{2}{\sqrt{\pi}}) \left(\log\left(\frac{\beta}{\alpha}\right) + 1 \right),$$

where in the last inequality, we have chosen A such that $\frac{A^2}{2C_\star} = \max\{\log(\frac{\beta}{\alpha}), 1\}$ since the previous derivations work for any positive A . We can get a similar bound for the left half of the integration defining A_2 (i.e. integration from $-\infty$ to 0). Combining

these two bounds leads to eq. (E.29). Bringing eq. (E.29) into eq. (E.28) leads to

$$(E.31) \quad \lambda_2 \leq -\frac{1}{(7 + \frac{4}{\sqrt{\pi}})(1 + \log(\frac{\beta}{\alpha}))}.$$

Therefore, we finish the proof of (5.19) in Theorem 5.7.

Step 2 (Lower bound). Next, we will construct an example to show the bound is sharp. The basic idea is to construct a sequence of triplets $\rho_{\text{post},n}$, β_n , and α_n , where $\lim_{n \rightarrow \infty} \frac{\beta_n}{\alpha_n} = \infty$, and the corresponding $-\lambda_{2,n} = \mathcal{O}(1/\log(\frac{\beta_n}{\alpha_n}))$. In the following proof, we ignore the subscript n for simplicity.

We consider the following sequence of posterior density functions ρ_{post} , such that (E.32)

$$\begin{aligned} -\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta) &= \int H(x) \frac{e^{-\frac{(\theta-x)^2}{2\sigma^2}}}{\sigma\sqrt{2\pi}} dx \quad H(x) = \begin{cases} \beta & \gamma - \sigma < x < \gamma + \sigma \\ \alpha & \text{Otherwise} \end{cases}, \\ \nabla_{\theta} \log \rho_{\text{post}}(\theta) &= \int_{-\infty}^{\theta} -\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta') d\theta' + c, \end{aligned}$$

where $-\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta)$ is a smoothed bump function containing four parameters $\gamma, \sigma > 0$ and $0 < \alpha < \beta$. Clearly, we have

$$\alpha I \preceq -\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta) \preceq \beta I.$$

Moreover, we will have another parameter c to determine $\nabla_{\theta} \log \rho_{\text{post}}(\theta)$. It is worth mentioning that for such α -strongly logconcave posterior, the Gaussian variational inference (5.2) has a unique minimizer (m^*, C^*) , which is determined by the stationary point condition in eq. (E.24); see also [78].

The **intuition** behind the construction of the bump function is as follows. Our objective is to ensure that the dominant eigenvalue, denoted as λ_2 in equation (E.25), as large as possible (thus leading to a lower bound). Since λ_2 satisfies

$$\begin{aligned} \lambda_2 &= -\frac{1 + A_2 - A_1^2 C_{\star}}{\frac{3}{2} + \frac{1}{2}A_2 + \sqrt{(\frac{1}{2}A_2 - \frac{1}{2})^2 + 2A_1^2 C_{\star}}} \geq -\frac{1 + A_2 - A_1^2 C_{\star}}{1 + A_2} \\ &= -\frac{\frac{1}{A_2} + (1 - \frac{A_1^2 C_{\star}}{A_2})}{1 + \frac{1}{A_2}}, \end{aligned}$$

we require A_2 to be as large as possible, whilst ensuring that the expression $(1 - \frac{A_1^2 C_{\star}}{A_2})$ as small as possible. Recall the definitions of A_1 and A_2 in eq. (E.26); for the latter term, we have

$$(E.33) \quad A_2 C_{\star}^{-1} = A_2 \mathbb{E}_{\rho_{a_{\star}}} [-\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta)]$$

$$(E.34) \quad = \mathbb{E}_{\rho_{a_{\star}}} [-\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta)(\theta - m_{\star})^2] \mathbb{E}_{\rho_{a_{\star}}} [-\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta)]$$

$$(E.35) \quad \geq (\mathbb{E}_{\rho_{a_{\star}}} [-\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta)(\theta - m_{\star})])^2 = A_1^2$$

due to the Cauchy-Schwarz inequality. Thus we get $(1 - \frac{A_1^2 C_{\star}}{A_2}) \geq 0$. To make this term as close to zero as possible, we consider when the Cauchy-Schwarz inequality can become equality. In fact, we need $-\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta)$ to take the form of a delta function. As in the assumption we have $\log \rho_{\text{post}}(\theta) \in C^2$ so this is not achievable. To approximate this condition, we can construct $-\nabla_{\theta} \nabla_{\theta} \log \rho_{\text{post}}(\theta)$ as a bump function $H(x)$ and gradually narrow the width of the bump to approach

zero. The Gaussian kernel is employed to smooth the bump function and simplify the subsequent calculations; this is why the form of H in (E.32) is constructed.

Now, we will provide a detailed construction. Instead of specifying β and c , we can specify $m_\star$ and $C_\star$ since there is a one-to-one correspondence between them. We specify⁹ that

$$(E.36) \quad \sigma = \gamma^{1.5} \quad m_\star = 0 \quad C_\star = -\frac{\gamma^2}{2\log \gamma} - \gamma^3 \quad \text{and} \quad \alpha = \frac{1}{(-\log \gamma)C_\star}.$$

Then, β and c are determined by the stationary point condition (E.24), namely

$$(E.37) \quad \begin{aligned} C_\star^{-1} &= \mathbb{E}_{\rho_{a_\star}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)] = \int H(x) \frac{1}{\sqrt{2\pi(\sigma^2 + C_\star)}} e^{-\frac{(x-m_\star)^2}{2(\sigma^2 + C_\star)}} dx \\ &= \alpha + (\beta - \alpha) \int_{\gamma-\sigma}^{\gamma+\sigma} \frac{1}{\sqrt{2\pi(\sigma^2 + C_\star)}} e^{-\frac{(x-m_\star)^2}{2(\sigma^2 + C_\star)}} dx. \\ 0 &= \mathbb{E}_{\rho_{a_\star}} [\nabla_\theta \log \rho_{\text{post}}(\theta)] = \int_{-\infty}^{\infty} \rho_{a_\star}(\theta) \int_{-\infty}^{\theta} -\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta') d\theta' d\theta + c. \end{aligned}$$

We will let $\gamma \rightarrow 0$ later. Note that in the above system, when γ is close to zero, for any $C_\star = \frac{1}{-\alpha \log \gamma} \leq \frac{1}{\alpha}$, we will have $\beta > \alpha$, so the above lead to valid specification of parameters. We now only have one free parameter γ . We can now also view the triplet ρ_{post} , β , and α as functions parameterized by γ .

In the following, we will let $\gamma \rightarrow 0$ and estimate the leading order of A_1 , A_2 , α and β in terms of γ . Let $m_{\star\sigma} = \frac{x C_\star + m_\star \sigma^2}{\sigma^2 + C_\star}$ and $C_{\star\sigma} = \frac{\sigma^2 C_\star}{\sigma^2 + C_\star}$, we have

$$-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta) \mathcal{N}(m_\star, C_\star) = \int H(x) \frac{e^{-\frac{(\theta-m_\star\sigma)^2}{2C_{\star\sigma}}}}{\sqrt{2\pi C_{\star\sigma}}} \frac{1}{\sqrt{2\pi(\sigma^2 + C_\star)}} e^{-\frac{(x-m_\star)^2}{2(\sigma^2 + C_\star)}} dx.$$

Bringing this to eq. (E.26) leads to

$$(E.38) \quad \begin{aligned} A_1 &= \mathbb{E}_{\rho_{a_\star}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)(\theta - m_\star)] \\ &= \int H(x) \frac{m_{\star\sigma} - m_\star}{\sqrt{2\pi(\sigma^2 + C_\star)}} e^{-\frac{(x-m_\star)^2}{2(\sigma^2 + C_\star)}} dx, \\ A_2 &= \mathbb{E}_{\rho_{a_\star}} [-\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)(\theta - m_\star)^2] \\ &= \int H(x) \frac{C_{\star\sigma} + (m_{\star\sigma} - m_\star)^2}{\sqrt{2\pi(\sigma^2 + C_\star)}} e^{-\frac{(x-m_\star)^2}{2(\sigma^2 + C_\star)}} dx. \end{aligned}$$

With the definitions in eq. (E.36), we have the following estimation about Gaussian integration

$$(E.39) \quad \int_{\gamma-\sigma}^{\gamma+\sigma} x^n e^{-\frac{x^2}{2(\sigma^2 + C_\star)}} dx = 2\sqrt{\gamma} \gamma^{n+2} (1 + \Theta(\log \gamma \sqrt{\gamma})).$$

Here we use Θ to denote the leading order term, as $\gamma \rightarrow 0$. Bringing the definition of σ and $C_\star$ and using change-of-variable with $y = \frac{x}{\gamma}$, the left hand side of eq. (E.39) becomes

$$(E.40) \quad \int_{\gamma-\sigma}^{\gamma+\sigma} x^n e^{-\frac{x^2}{2(\sigma^2 + C_\star)}} dx = \gamma^{n+1} \int_{1-\sqrt{\gamma}}^{1+\sqrt{\gamma}} y^n e^{y^2 \log \gamma} dy.$$

⁹These choices of parameters are motivated by the subsequent calculations.

Bringing the following inequalities into eq. (E.40) leads to eq. (E.39)

$$\begin{aligned}\gamma^{n+1} \int_{1-\sqrt{\gamma}}^{1+\sqrt{\gamma}} y^n e^{y^2 \log \gamma} dy &\geq 2\sqrt{\gamma} \gamma^{n+1} (1 - \sqrt{\gamma})^n e^{(1+\sqrt{\gamma})^2 \log \gamma} \\ &= 2\sqrt{\gamma} \gamma^{n+2} (1 + \Theta(2 \log \gamma \sqrt{\gamma})), \\ \gamma^{n+1} \int_{1-\sqrt{\gamma}}^{1+\sqrt{\gamma}} y^n e^{y^2 \log \gamma} dy &\leq 2\sqrt{\gamma} \gamma^{n+1} (1 + \sqrt{\gamma})^n e^{(1-\sqrt{\gamma})^2 \log \gamma} \\ &= 2\sqrt{\gamma} \gamma^{n+2} (1 - \Theta(2 \log \gamma \sqrt{\gamma})).\end{aligned}$$

Here we used the Taylor expansions of $(1-\sqrt{\gamma})^n e^{(\gamma+2\sqrt{\gamma}) \log \gamma} = \Theta((1-n\sqrt{\gamma})(1+(\gamma+2\sqrt{\gamma}) \log \gamma))$ and $(1+\sqrt{\gamma})^n e^{(\gamma-2\sqrt{\gamma}) \log \gamma} = \Theta((1+n\sqrt{\gamma})(1+(\gamma-2\sqrt{\gamma}) \log \gamma))$. Then the estimation for A_1, A_2 from eq. (E.38) and the covariance condition in eq. (E.24) become

$$\begin{aligned}\text{(E.41a)} \quad A_1 &= \frac{(\beta - \alpha)C_\star}{\sigma^2 + C_\star} \int_{\gamma-\sigma}^{\gamma+\sigma} x \frac{1}{\sqrt{2\pi(\sigma^2 + C_\star)}} e^{-\frac{x^2}{2(\sigma^2 + C_\star)}} dx \\ &= \frac{(\beta - \alpha)C_\star}{\sigma^2 + C_\star} \frac{2\gamma^{2.5} \sqrt{-\log \gamma}}{\sqrt{\pi}} (1 + \Theta(\log \gamma \sqrt{\gamma})) \quad (\text{Using (E.39)}) \\ &= \frac{2(\beta - \alpha)}{\sqrt{\pi}} \gamma^{2.5} (-\log \gamma)^{0.5} + \Theta\left(\frac{2(\beta - \alpha)}{\sqrt{\pi}} \gamma^3 (-\log \gamma)^{1.5}\right),\end{aligned}$$

$$\begin{aligned}\text{(E.41b)} \quad A_2 &= \alpha C_\star + (\beta - \alpha) \int_{\gamma-\sigma}^{\gamma+\sigma} \frac{C_\star \sigma + m_{\star\sigma}^2}{\sqrt{2\pi(\sigma^2 + C_\star)}} e^{-\frac{x^2}{2(\sigma^2 + C_\star)}} dx \\ \text{(E.41c)} \quad &= \alpha C_\star + (\beta - \alpha) \frac{2\sqrt{-\log \gamma} \gamma^{3.5} (1 + 2\gamma \log \gamma) (1 + \gamma + 2\gamma \log \gamma)}{\sqrt{\pi}} \\ &\quad \times (1 + \Theta(\log \gamma \sqrt{\gamma})) \quad (\text{Using (E.39)}) \\ &= \alpha C_\star + \frac{2(\beta - \alpha)}{\sqrt{\pi}} \gamma^{3.5} (-\log \gamma)^{0.5} + \Theta\left(\frac{2(\beta - \alpha)}{\sqrt{\pi}} \gamma^4 (-\log \gamma)^{1.5}\right),\end{aligned}$$

$$\begin{aligned}\text{(E.41d)} \quad C_\star^{-1} &= \alpha + (\beta - \alpha) \int_{\gamma-\sigma}^{\gamma+\sigma} \frac{1}{\sqrt{2\pi(\sigma^2 + C_\star)}} e^{-\frac{x^2}{2(\sigma^2 + C_\star)}} dx \\ &= \alpha + (\beta - \alpha) 2\gamma^{1.5} \frac{\sqrt{-\log \gamma}}{\sqrt{\pi}} (1 + \Theta(\log \gamma \sqrt{\gamma})) \quad (\text{Using (E.39)}) \\ &= \alpha + \frac{2(\beta - \alpha)}{\sqrt{\pi}} \gamma^{1.5} (-\log \gamma)^{0.5} + \Theta\left(\frac{2(\beta - \alpha)}{\sqrt{\pi}} \gamma^2 (-\log \gamma)^{1.5}\right).\end{aligned}$$

From eq. (E.36), we have

$$\text{(E.42)} \quad \alpha = \Theta\left(\frac{1}{\gamma^2}\right).$$

Combining eq. (E.42), the definition of $C_\star$ in eq. (E.36), and eq. (E.41d) leads to the estimation about β , as follows:

$$\text{(E.43)} \quad \frac{\beta}{\alpha} = \Theta\left(\frac{(-\log \gamma)^{0.5}}{\gamma^{1.5}}\right) \quad \text{and} \quad \beta - \alpha = \Theta\left(\frac{(-\log \gamma)^{0.5}}{\gamma^{3.5}}\right).$$

Combining eq. E.41–E.43 leads to the estimations about A_1 and A_2 , as follows

$$(E.44) \quad A_1 = \Theta\left(\frac{-\log \gamma}{\gamma}\right) \quad A_2 = \Theta(-\log \gamma),$$

$$(E.45) \quad 1 - \frac{A_1^2 C_\star}{A_2} = \frac{A_2 C_\star^{-1} - A_1^2}{A_2 C_\star^{-1}} = \frac{\Theta\left((- \log \gamma)/\gamma^2\right)}{\Theta\left((- \log \gamma)^2/\gamma^2\right)} = \Theta\left(\frac{1}{-\log \gamma}\right).$$

Finally, we can bound the large eigenvalue of eq. (E.25) by

$$-\lambda_2 = \frac{1 + A_2 - A_1^2 C_\star}{\frac{3}{2} + \frac{1}{2}A_2 + \sqrt{(\frac{1}{2}A_2 - \frac{1}{2})^2 + 2A_1^2 C_\star}} \leq \frac{\frac{1}{A_2} + (1 - \frac{A_1^2 C_\star}{A_2})}{1 + \frac{1}{A_2}} = \Theta\left(\frac{1}{-\log \gamma}\right).$$

Here we use $A_2 \geq 0$ and $1 - \frac{A_1^2 C_\star}{A_2} \geq 0$. Eq. (E.43) $\log(\frac{\beta}{\alpha}) = \Theta(-\log \gamma)$ indicates that for the constructed logconcave density, the local convergence rate is not faster than $-\mathcal{O}(1/\log(\frac{\beta}{\alpha}))$.

E.7. Proof of Theorem 5.8. Consider the following example, where $\theta \in \mathbb{R}$ and $\Phi_R(\theta) = \sum_{k=0}^{2K+1} a_{2k} \theta^{2k}$ with $a_{4K+2} > 0$. We will choose a_{2k} later so that the convergence of these dynamics is $\Theta(t^{-\frac{1}{2K}})$. Recall the Gaussian approximate Fisher-Rao gradient flow is

$$(E.46) \quad \begin{aligned} \frac{dm_t}{dt} &= C_t \mathbb{E}_{\rho_{a_t}} [\nabla_\theta \log \rho_{\text{post}}], \\ \frac{dC_t}{dt} &= C_t + C_t \mathbb{E}_{\rho_{a_t}} [\nabla_\theta \nabla_\theta \log \rho_{\text{post}}] C_t. \end{aligned}$$

We first calculate the explicit formula of the dynamics. For the mean part, we have

$$(E.47) \quad \begin{aligned} \mathbb{E}_{\rho_a} [\nabla_\theta \log \rho_{\text{post}}(\theta)] &= - \sum_{k=1}^{2K+1} 2ka_{2k} \mathbb{E}_{\rho_a} [\theta^{2k-1}] \\ &= - \sum_{k=1}^{2K+1} 2ka_{2k} \sum_{i=0}^{k-1} \binom{2k-1}{2i+1} m^{2i+1} C^{k-i-1} \frac{(2k-2i-2)!}{2^{k-i-1}(k-i-1)!}. \end{aligned}$$

In the above we have used the explicit formula for the moments of Gaussian distributions. By (E.47), we know that when $m = 0$, $\mathbb{E}_{\rho_a} [\nabla_\theta \log \rho_{\text{post}}(\theta)] = 0$. Later, we will initialize the dynamics at $m_0 = 0$; as a consequence, $m_t = 0$ so we only need to consider the covariance dynamics given $m = 0$.

For the covariance part (assuming $m = 0$), using Stein's identity, we get

$$(E.48) \quad \mathbb{E}_{\rho_a} [\nabla_\theta \nabla_\theta \log \rho_{\text{post}}(\theta)]_{m=0} = \frac{\partial \mathbb{E}_{\rho_a} [\nabla_\theta \log \rho_{\text{post}}(\theta)]}{\partial m} \Big|_{m=0} = -f(C),$$

where

$$(E.49) \quad f(C) = \sum_{k=1}^{2K+1} 2k(2k-1)a_{2k} C^{k-1} \frac{(2k-2)!}{2^{k-1}(k-1)!}.$$

We choose $\{a_{2k}\}_{k=1}^{2K+1}$ such that

$$2k(2k-1) \frac{(2k-2)!}{2^{k-1}(k-1)!} a_{2k} = \binom{2K+1}{k} (-1)^{2K+1-k},$$

which leads to the identity

$$1 - f(C)C = -(C - 1)^{2K+1}.$$

Now, we calculate the explicit form of the dynamics, with the above choice of Φ_R . We initialize the dynamics with $m_0 = 0$. For the Gaussian approximate Fisher-Rao gradient flow eq. (5.9), we have

$$(E.50) \quad \begin{aligned} \frac{dm_t}{dt} &= 0, \\ \frac{dC_t}{dt} &= C_t(1 - f(C_t)C_t) = -C_t(C_t - 1)^{2K+1}. \end{aligned}$$

For the vanilla Gaussian approximate gradient flow eq. (5.18), we have

$$(E.51) \quad \begin{aligned} \frac{dm_t}{dt} &= 0, \\ \frac{dC_t}{dt} &= \frac{1}{2C_t}(1 - f(C_t)C_t) = -\frac{(C_t - 1)^{2K+1}}{2C_t}. \end{aligned}$$

For the Gaussian approximate Wasserstein gradient flow eq. (5.12), we have

$$(E.52) \quad \begin{aligned} \frac{dm_t}{dt} &= 0, \\ \frac{dC_t}{dt} &= 2(1 - f(C_t)C_t) = -2(C_t - 1)^{2K+1}. \end{aligned}$$

For these ODE systems given in eqs. (E.50) to (E.52), when C_0 is close to 1, their behavior is well approximated by the following equation:

$$(E.53) \quad \frac{dC_t}{dt} = -\alpha(C_t - 1)^{2K+1},$$

where $\alpha = 1, -\frac{1}{2}, -2$ for eqs. (E.50) to (E.52) respectively. This system has a stationary point at $C_t = 1$. When $C_0 \neq 1$, the solution to eq. (E.53) is given by

$$(E.54) \quad C_t = 1 + \left(2K\alpha t + (C_0 - 1)^{-2K}\right)^{-\frac{1}{2K}},$$

which converges to 1 at the rate $\mathcal{O}(t^{-\frac{1}{2K}})$.

APPENDIX F. DETAILS OF NUMERICAL INTEGRATION

In this section, we discuss how to compute the reference values of $\mathbb{E}[\theta]$, $\text{Cov}[\theta]$, and $\mathbb{E}[\cos(\omega^T \theta + b)]$ for the logconcave posterior and Rosenbrock posterior in Section 4.5. First, for any Gaussian distribution, we have

$$\int \cos(\omega^T \theta + b) \mathcal{N}(\theta; m, C) d\theta = \exp\left(-\frac{1}{2}\omega^T C \omega\right) \cos(\omega^T m + b).$$

We can rewrite the logconcave function as

$$\Phi_R(\theta) = \frac{1}{2} \frac{(\theta^{(1)} - \theta^{(2)}/\sqrt{\lambda})^2}{10/\lambda} + \frac{\theta^{(2)^4}{20}.$$

The following integration formula holds:

$$\begin{aligned}
\int e^{-\Phi_R(\theta)} d\theta^{(1)} &= \sqrt{20\pi/\lambda} e^{-\theta^{(2)4}/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(1)} d\theta^{(1)} &= \sqrt{20\pi/\lambda} \frac{\theta^{(2)}}{\sqrt{\lambda}} e^{-\theta^{(2)4}/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(2)} d\theta^{(1)} &= \sqrt{20\pi/\lambda} \theta^{(2)} e^{-\theta^{(2)4}/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(1)2} d\theta^{(1)} &= \sqrt{20\pi/\lambda} \left(\frac{\theta^{(2)2}}{\lambda} + \frac{10}{\lambda} \right) e^{-\theta^{(2)4}/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(1)} \theta^{(2)} d\theta^{(1)} &= \sqrt{20\pi/\lambda} \frac{\theta^{(2)2}}{\sqrt{\lambda}} e^{-\theta^{(2)4}/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(2)2} d\theta^{(1)} &= \sqrt{20\pi/\lambda} \theta^{(2)2} e^{-\theta^{(2)4}/20}, \\
\int e^{-\Phi_R(\theta)} \cos(\omega^{(1)} \theta^{(1)} + \omega^{(2)} \theta^{(2)} + b) d\theta^{(1)} \\
&= \sqrt{20\pi/\lambda} e^{-\frac{5}{\lambda} \omega^{(1)2}} \cos(\omega^{(1)} \theta^{(2)} / \sqrt{\lambda} + \omega^{(2)} \theta^{(2)} + b) e^{-\theta^{(2)4}/20}.
\end{aligned}$$

Other 2D integrations can be addressed by first performing 1D integration with respect to $\theta^{(1)}$ analytically, and then the second 1D integration with respect to $\theta^{(2)}$ is computed numerically with 10^7 uniform points.

We can rewrite the Rosenbrock function as

$$\Phi_R(\theta) = \frac{1}{2} \frac{(\theta^{(2)} - \theta^{(1)2})^2}{10/\lambda} + \frac{(1 - \theta^{(1)})^2}{20}.$$

We have the following integration formula:

$$\begin{aligned}
\int e^{-\Phi_R(\theta)} d\theta^{(2)} &= \sqrt{20\pi/\lambda} e^{-(1-\theta^{(1)})^2/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(1)} d\theta^{(2)} &= \sqrt{20\pi/\lambda} \theta^{(1)} e^{-(1-\theta^{(1)})^2/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(2)} d\theta^{(2)} &= \sqrt{20\pi/\lambda} \theta^{(1)2} e^{-(1-\theta^{(1)})^2/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(1)2} d\theta^{(2)} &= \sqrt{20\pi/\lambda} \theta^{(1)2} e^{-(1-\theta^{(1)})^2/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(2)2} d\theta^{(2)} &= \sqrt{20\pi/\lambda} \left(\theta^{(1)4} + \frac{10}{\lambda} \right) e^{-(1-\theta^{(1)})^2/20}, \\
\int e^{-\Phi_R(\theta)} \theta^{(1)} \theta^{(2)} d\theta^{(2)} &= \sqrt{20\pi/\lambda} \theta^{(1)3} e^{-(1-\theta^{(1)})^2/20}, \\
\int e^{-\Phi_R(\theta)} \cos(\omega^{(2)} \theta^{(2)} + \omega^{(1)} \theta^{(1)} + b) d\theta^{(2)} \\
&= \sqrt{20\pi/\lambda} e^{-\frac{5}{\lambda} \omega^{(2)2}} \cos(\omega^{(2)} \theta^{(1)2} + \omega^{(1)} \theta^{(1)} + b) e^{-(1-\theta^{(1)})^2/20}.
\end{aligned}$$

Moreover, we have

$$\int \int e^{-\Phi_R(\theta)} d\theta^{(1)} d\theta^{(2)} = \frac{20\pi}{\sqrt{\lambda}} \quad \mathbb{E}[\theta] = \begin{bmatrix} 1 \\ 11 \end{bmatrix} \quad \text{Cov}[\theta] = \begin{bmatrix} 10 & 20 \\ 20 & \frac{10}{\lambda} + 240 \end{bmatrix}.$$

Other 2D integrations can be addressed by first performing 1D integration with respect to $\theta^{(2)}$ analytically, and then the second 1D integration with respect to $\theta^{(1)}$ is computed numerically with 10^7 uniform points.

REFERENCES

- [1] S.-I. Amari, *Natural gradient works efficiently in learning*, Neural Comput., **10** (1998), 251–276.
- [2] S.-i. Amari, *Information Geometry and its Applications*, Applied Mathematical Sciences, vol. 194, Springer, [Tokyo], 2016, DOI 10.1007/978-4-431-55978-8. MR3495836
- [3] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient Flows in Metric Spaces and in the Space of Probability Measures*, Lectures in Mathematics ETH Zürich, Birkhäuser Verlag, Basel, 2005. MR2129498
- [4] L. Ambrosio and G. Savaré, *Gradient flows of probability measures*, Handbook of Differential Equations: Evolutionary Equations. Vol. III, Handb. Differ. Equ., Elsevier/North-Holland, Amsterdam, 2007, pp. 1–136, DOI 10.1016/S1874-5717(07)80004-1. MR2549368
- [5] N. Ay, J. Jost, H. V. Lê, and L. Schwachhöfer, *Information geometry and sufficient statistics*, Probab. Theory Related Fields **162** (2015), no. 1-2, 327–364, DOI 10.1007/s00440-014-0574-8. MR3350047
- [6] N. Ay, J. Jost, H. V. Lê, and L. Schwachhöfer, *Information Geometry*, Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics], vol. 64, Springer, Cham, 2017, DOI 10.1007/978-3-319-56478-4. MR3701408
- [7] D. Bakry and M. Émery, *Diffusions hypercontractives* (French), Séminaire de probabilités, XIX, 1983/84, Lecture Notes in Math., vol. 1123, Springer, Berlin, 1985, pp. 177–206, DOI 10.1007/BFb0075847. MR889476
- [8] D. Bakry, I. Gentil, and M. Ledoux, *Analysis and Geometry of Markov Diffusion Operators*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 348, Springer, Cham, 2014, DOI 10.1007/978-3-319-00227-9. MR3155209
- [9] M. Bauer, M. Bruveris, and P. W. Michor, *Uniqueness of the Fisher-Rao metric on the space of smooth densities*, Bull. Lond. Math. Soc. **48** (2016), no. 3, 499–506, DOI 10.1112/blms/bdw020. MR3509909
- [10] B. M. Bell, *The iterated Kalman smoother as a Gauss-Newton method*, SIAM J. Optim. **4** (1994), no. 3, 626–636, DOI 10.1137/0804035. MR1287820
- [11] B. M. Bell and F. W. Cathey, *The iterated Kalman filter update as a Gauss-Newton method*, IEEE Trans. Automat. Control **38** (1993), no. 2, 294–297, DOI 10.1109/9.250476. MR1206815
- [12] R. Bhatia, T. Jain, and Y. Lim, *On the Bures-Wasserstein distance between positive definite matrices*, Expo. Math. **37** (2019), no. 2, 165–191, DOI 10.1016/j.exmath.2018.01.002. MR3992484
- [13] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, *Variational inference: a review for statisticians*, J. Amer. Statist. Assoc. **112** (2017), no. 518, 859–877, DOI 10.1080/01621459.2017.1285773. MR3671776
- [14] N. M. Boffi and E. Vanden-Eijnden, *Probability flow solution of the Fokker-Planck equation*, Mac. Learn. Sci. Techol. **4** (2023), 035012, DOI 10.1088/2632-2153/ace2aa.
- [15] S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng (eds.), *Handbook of Markov Chain Monte Carlo*, Chapman & Hall/CRC Handbooks of Modern Statistical Methods, CRC Press, Boca Raton, FL, 2011, DOI 10.1201/b10905. MR2742422
- [16] M. Burger, M. Erbar, F. Hoffmann, D. Matthes, and A. Schlichting, *Covariance-modulated optimal transport and gradient flows*, Arch. Ration. Mech. Anal. **249** (2025), no. 1, Paper No. 7, 95, DOI 10.1007/s00205-024-02065-w. MR4841737
- [17] S. Cao and D. Z. Huang, *Bayesian calibration for large-scale fluid structure interaction problems under embedded/immersed boundary framework*, Internat. J. Numer. Methods Engrg. **123** (2022), no. 8, 1791–1812, DOI 10.1002/nme.6916. MR4403854
- [18] J. A. Carrillo and U. Vaes, *Wasserstein stability estimates for covariance-preconditioned Fokker-Planck equations*, Nonlinearity **34** (2021), no. 4, 2275–2295, DOI 10.1088/1361-6544/abbe62. MR4246458

- [19] J. A. Carrillo, Y. Chen, D. Z. Huang, J. Huang, and D. Wei, *Fisher-Rao gradient flow: geodesic convexity and functional inequalities*, Preprint, [arXiv:2407.15693](#), 2024.
- [20] J. A. Carrillo, Y.-P. Choi, C. Totzeck, and O. Tse, *An analytical framework for consensus-based global optimization method*, Math. Models Methods Appl. Sci. **28** (2018), no. 6, 1037–1066, DOI 10.1142/S0218202518500276. MR3804923
- [21] N. N. Čencov, *Statistical Decision Rules and Optimal Inference*, Translations of Mathematical Monographs, vol. 53, American Mathematical Society, Providence, RI, 1982. Translation from the Russian edited by Lev J. Leifman. MR645898
- [22] N. K. Chada and X. T. Tong, *Convergence acceleration of ensemble Kalman inversion in nonlinear settings*, Math. Comp. **91** (2022), no. 335, 1247–1280, DOI 10.1090/mcom/3709. MR4405495
- [23] N. K. Chada, Y. Chen, and D. Sanz-Alonso, *Iterative ensemble Kalman methods: a unified perspective with some new variants*, Found. Data Sci. **3** (2021), no. 3, 331–369, DOI 10.3934/fods.2021011. MR4619955
- [24] B. Che, Y. Chen, Z. Huan, D. Z. Huang, and W. Wang, *Stable derivative free Gaussian mixture variational inference for Bayesian inverse problems*, SIAM J. Sci. Comput. **47** (2025), no. 5, A2583–A2608, DOI 10.1137/25M1722548. MR4959980
- [25] Y. Chen, B. Hosseini, H. Owhadi, and A. M. Stuart, *Solving and learning nonlinear PDEs with Gaussian processes*, J. Comput. Phys. **447** (2021), Paper No. 110668, 29, DOI 10.1016/j.jcp.2021.110668. MR4311012
- [26] Y. Chen, D. Z. Huang, J. Huang, S. Reich, and A. M. Stuart, *Gradient flows for sampling: Mean-field models, Gaussian approximations and affine invariance*, Preprint, [arXiv:2302.11024](#), (2023).
- [27] Y. Chen, D. Z. Huang, J. Huang, S. Reich, and A. M. Stuart, *Efficient, multimodal, and derivative-free Bayesian inference with Fisher-Rao gradient flows*, Inverse Problems **40** (2024), no. 12, Paper No. 125001, 39, DOI 10.1088/1361-6420/ad847b. MR4814274
- [28] Y. Chen and W. Li, *Optimal transport natural gradient for statistical manifolds with continuous sample space*, Inf. Geom. **3** (2020), no. 1, 1–32, DOI 10.1007/s41884-020-00028-0. MR4119003
- [29] Y. Chen and D. S. Oliver, *Ensemble randomized maximum likelihood method as an iterative ensemble smoother*, Math. Geosci. **44** (2012), 1–26.
- [30] Y. Chen, H. Owhadi, and A. M. Stuart, *Consistency of empirical Bayes and kernel flow for hierarchical parameter estimation*, Math. Comp. **90** (2021), no. 332, 2527–2578, DOI 10.1090/mcom/3649. MR4305361
- [31] S. Chewi, M. A. Erdogdu, M. Li, R. Shen, and M. S. Zhang, *Analysis of Langevin Monte Carlo from Poincaré to log-Sobolev*, Found. Comput. Math. **25** (2025), no. 4, 1345–1395, DOI 10.1007/s10208-024-09667-6. MR4951529
- [32] S. Chewi, T. Le Gouic, C. Lu, T. Maunu, and P. Rigollet, *SVGD as a kernelized Wasserstein gradient flow of the chi-squared divergence*, Adv. Neural Inf. Process. Syst. **33** (2020), 2098–2109.
- [33] S. Chewi, T. Le Gouic, C. Lu, T. Maunu, P. Rigollet, and A. Stromme, *Exponential ergodicity of mirror-Langevin diffusions*, Adv. Neural Inf. Process. Syst. **33** (2020), 19573–19585.
- [34] L. Chizat and F. Bach, *On the global convergence of gradient descent for over-parameterized models using optimal transport*, Adv. Neural Inf. Process. Syst. **31** (2018).
- [35] P. G. Constantine, E. Dow, and Q. Wang, *Active subspace methods in theory and practice: applications to kriging surfaces*, SIAM J. Sci. Comput. **36** (2014), no. 4, A1500–A1524, DOI 10.1137/130916138. MR3233940
- [36] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White, *MCMC methods for functions: modifying old algorithms to make them faster*, Statist. Sci. **28** (2013), no. 3, 424–446, DOI 10.1214/13-STS421. MR3135540
- [37] J. Coullon and R. J. Webber, *Ensemble sampler for infinite-dimensional inverse problems*, Stat. Comput. **31** (2021), no. 3, Paper No. 28, 9, DOI 10.1007/s11222-021-10004-y. MR4234336
- [38] T. Cui, K. J. H. Law, and Y. M. Marzouk, *Dimension-independent likelihood-informed MCMC*, J. Comput. Phys. **304** (2016), 109–137, DOI 10.1016/j.jcp.2015.10.008. MR3422405
- [39] A. S. Dalalyan and L. Riou-Durand, *On sampling from a log-concave density using kinetic Langevin diffusions*, Bernoulli **26** (2020), no. 3, 1956–1988, DOI 10.3150/19-BEJ1178. MR4091098

- [40] M. Dashti and A. M. Stuart, *The Bayesian approach to inverse problems*, Preprint, [arXiv:1302.6989](#), (2013).
- [41] K. Daudel, R. Douc, and F. Portier, *Infinite-dimensional gradient-based descent for alpha-divergence minimisation*, *Ann. Statist.* **49** (2021), no. 4, 2250–2270, DOI 10.1214/20-aos2035. MR4319249
- [42] K. Daudel, et al., *Mixture weights optimisation for alpha-divergence variational inference*, *Adv. Neural Inf. Process. Syst.* **34** (2021), 4397–4408.
- [43] P. Del Moral, A. Doucet, and A. Jasra, *Sequential Monte Carlo samplers*, *J. R. Stat. Soc. Ser. B Stat. Methodol.* **68** (2006), no. 3, 411–436, DOI 10.1111/j.1467-9868.2006.00553.x. MR2278333
- [44] G. Detommaso, T. Cui, Y. Marzouk, A. Spantini, and R. Scheichl, *A Stein variational Newton method*, *Adv. Neural Inf. Process. Syst.* **31** (2018).
- [45] M. P. Do Carmo and J. Flaherty Francis, *Riemannian Geometry*, vol. 6, Springer, 1992.
- [46] C. Domingo-Enrich and A.-A. Pooladian, *An explicit expansion of the Kullback-Leibler divergence along its Fisher-Rao gradient flow*, Preprint, [arXiv:2302.12229](#), 2023.
- [47] A. Doucet and A. M. Johansen, *A tutorial on particle filtering and smoothing: fifteen years later*, *The Oxford Handbook of Nonlinear Filtering*, Oxford Univ. Press, Oxford, 2011, pp. 656–704. MR2884612
- [48] A. Duncan, N. Nüsken, and L. Szpruch, *On the geometry of Stein variational gradient descent*, *J. Mach. Learn. Res.* **24** (2023), Paper No. [56], 39. MR4582478
- [49] M. M. Dunlop and G. Stadler, *A gradient-free subspace-adjusting ensemble sampler for infinite-dimensional Bayesian inverse problems*, Preprint, [arXiv:2202.11088](#), 2022.
- [50] A. Durmus, S. Majewski, and B. Miasojedow, *Analysis of Langevin Monte Carlo via convex optimization*, *J. Mach. Learn. Res.* **20** (2019), Paper No. 73, 46. MR3960927
- [51] A. A. Emerick and A. C. Reynolds, *Investigation of the sampling performance of ensemble-based methods with a simple reservoir model*, *Comput. Geosci.* **17** (2013), no. 2, 325–350, DOI 10.1007/s10596-012-9333-z. MR3036193
- [52] G. Evensen, *Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics*, *J. Geophys. Res. Oceans* **99** (1994), 10143–10162.
- [53] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *emcee: the MCMC hammer*, *Publ. Astron. Soc. Pacific* **125** (2013), 306.
- [54] D. Friedman, *Mathematical Notes: The Functional Equation $f(x+y) = g(x) + h(y)$* , *Amer. Math. Monthly* **69** (1962), no. 8, 769–772, DOI 10.2307/2310775. MR1531838
- [55] T. Friedrich, *Die fisher-information und symplektische strukturen*, *Math. Nachr.* **153** (1991), 273–296, DOI 10.1002/mana.19911530125.
- [56] T. Galy-Fajou, V. Perrone, and M. Opper, *Flexible and efficient inference with particles for the variational Gaussian approximation*, *Entropy* **23** (2021), no. 8, Paper No. 990, 34, DOI 10.3390/e23080990. MR4303454
- [57] 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), no. 1, 412–441, DOI 10.1137/19M1251655. MR4059375
- [58] A. Garbuno-Inigo, N. Nüsken, and S. Reich, *Affine invariant interacting Langevin dynamics for Bayesian inference*, *SIAM J. Appl. Dyn. Syst.* **19** (2020), no. 3, 1633–1658, DOI 10.1137/19M1304891. MR4123680
- [59] N. García Trillos, B. Hosseini, and D. Sanz-Alonso, *From optimization to sampling through gradient flows*, *Notices Amer. Math. Soc.* **70** (2023), no. 6, 905–917, DOI 10.1090/noti2717. MR4601361
- [60] N. Garcia Trillos and D. Sanz-Alonso, *The Bayesian update: variational formulations and gradient flows*, *Bayesian Anal.* **15** (2020), no. 1, 29–56, DOI 10.1214/18-BA1137. MR4050876
- [61] J. Goodman and J. Weare, *Ensemble samplers with affine invariance*, *Commun. Appl. Math. Comput. Sci.* **5** (2010), no. 1, 65–80, DOI 10.2140/camcos.2010.5.65. MR2600822
- [62] L. Gross, *Logarithmic Sobolev inequalities*, *Amer. J. Math.* **97** (1975), no. 4, 1061–1083, DOI 10.2307/2373688. MR420249
- [63] Y. He, K. Balasubramanian, B. K. Sriperumbudur, and J. Lu, *Regularized Stein variational gradient flow*, *Found. Comput. Math.* **25** (2025), no. 4, 1199–1257, DOI 10.1007/s10208-024-09663-w. MR4951526

- [64] Y.-P. Hsieh, A. Kavis, P. Rolland, and V. Cevher, *Mirrored Langevin dynamics*, Adv. Neural Inf. Process. Syst. **31** (2018).
- [65] D. Z. Huang, J. Huang, S. Reich, and A. M. Stuart, *Efficient derivative-free Bayesian inference for large-scale inverse problems*, Inverse Problems **38** (2022), no. 12, Paper No. 125006, 40, DOI 10.1088/1361-6420/ac99fa. MR4500907
- [66] D. Z. Huang, T. Schneider, and A. M. Stuart, *Iterated Kalman methodology for inverse problems*, J. Comput. Phys. **463** (2022), Paper No. 111262, 40, DOI 10.1016/j.jcp.2022.111262. MR4418985
- [67] M. A. Iglesias, K. J. H. Law, and A. M. Stuart, *Ensemble Kalman methods for inverse problems*, Inverse Problems **29** (2013), no. 4, 045001, 20, DOI 10.1088/0266-5611/29/4/045001. MR3041539
- [68] T. Isaac, N. Petra, G. Stadler, and O. Ghattas, *Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet*, J. Comput. Phys. **296** (2015), 348–368, DOI 10.1016/j.jcp.2015.04.047. MR3348506
- [69] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, *An introduction to variational methods for graphical models*, Mach. Learn. **37** (1999), 183–233.
- [70] R. Jordan, D. Kinderlehrer, and F. Otto, *The variational formulation of the Fokker-Planck equation*, SIAM J. Math. Anal. **29** (1998), no. 1, 1–17, DOI 10.1137/S0036141096303359. MR1617171
- [71] S. J. Julier and J. K. Uhlmann, *New extension of the Kalman filter to nonlinear systems*, Signal Processing, Sensor Fusion, and Target Recognition VI, vol. 3068, International Society for Optics and Photonics, 1997, 182–193.
- [72] S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte, *A new approach for filtering nonlinear systems*, Proceedings of 1995 American Control Conference-ACC'95, vol. 3, IEEE, 1995, 1628–1632.
- [73] R. E. Kalman, *A new approach to linear filtering and prediction problems*, Trans. ASME Ser. D. J. Basic Engrg. **82** (1960), no. 1, 35–45. MR3931993
- [74] M. Khan and W. Lin, *Conjugate-computation variational inference: converting variational inference in non-conjugate models to inferences in conjugate models*, Artificial Intelligence and Statistics, PMLR, 2017, 878–887.
- [75] H. Kim, D. Sanz-Alonso, and A. Strang, *Hierarchical ensemble Kalman methods with sparsity-promoting generalized gamma hyperpriors*, Found. Data Sci. **5** (2023), no. 3, 366–388, DOI 10.3934/fods.2023003. MR4622925
- [76] M. Lambert, S. Bonnabel, and F. Bach, *The continuous-discrete variational Kalman filter (CD-VKF)*, 2022 IEEE 61st Conference on Decision and Control (CDC), 2022, 6632–6639, DOI 10.1109/CDC51059.2022.9992993.
- [77] M. Lambert, S. Bonnabel, and F. Bach, *The recursive variational Gaussian approximation (R-VGA)*, Stat. Comput. **32** (2022), no. 1, Paper No. 10, 24, DOI 10.1007/s11222-021-10068-w. MR4354710
- [78] M. Lambert, S. Chewi, F. Bach, S. Bonnabel, and P. Rigollet, *Variational inference via Wasserstein gradient flows*, Adv. Neural Inf. Process. Syst. **35** (2022), 14434–14447.
- [79] B. Leimkuhler, C. Matthews, and J. Weare, *Ensemble preconditioning for Markov chain Monte Carlo simulation*, Stat. Comput. **28** (2018), no. 2, 277–290, DOI 10.1007/s11222-017-9730-1. MR3747563
- [80] W. Li and J. Zhao, *Wasserstein information matrix*, Inf. Geom. (2023), DOI 10.1007/s41884-023-00099-9.
- [81] W. Lin, M. E. Khan, and M. Schmidt, *Fast and simple natural-gradient variational inference with mixture of exponential-family approximations*, International Conference on Machine Learning, PMLR, 2019, 3992–4002.
- [82] M. Lindsey, J. Weare, and A. Zhang, *Ensemble Markov chain Monte Carlo with teleporting walkers*, SIAM/ASA J. Uncertain. Quantif. **10** (2022), no. 3, 860–885, DOI 10.1137/21M1425062. MR4535507
- [83] Q. Liu, *Stein variational gradient descent as gradient flow*, Adv. Neural Inf. Process. Syst. **30** (2017).
- [84] Q. Liu and D. Wang, *Stein variational gradient descent: a general purpose Bayesian inference algorithm*, Adv. Neural Inf. Process. Syst. **29** (2016).

- [85] T. Liu, P. Ghosal, K. Balasubramanian, and N. Pillai, *Towards understanding the dynamics of Gaussian-Stein variational gradient descent*, Adv. Neural Inf. Process. Syst. **36** (2023), 61234–61291.
- [86] Z. Liu, A. Stuart, and Y. Wang, *Second order ensemble Langevin method for sampling and inverse problems*, Commun. Math. Sci. **23** (2025), no. 5, 1299–1317, DOI 10.4310/cms.250517001811. MR4907170
- [87] I. Lopez-Gomez, C. Christopoulos, H. L. Langeland Ervik, O. R. Dunbar, Y. Cohen, and T. Schneider, *Training physics-based machine-learning parameterizations with gradient-free ensemble Kalman methods*, J. Adv. Model. Earth Syst. **14** (2022), e2022MS003105.
- [88] J. Lott, *Some geometric calculations on Wasserstein space*, Comm. Math. Phys. **277** (2008), 423–437, DOI 10.1007/s00220-007-0367-3.
- [89] J. Lu, Y. Lu, and J. Nolen, *Scaling limit of the Stein variational gradient descent: the mean field regime*, SIAM J. Math. Anal. **51** (2019), no. 2, 648–671, DOI 10.1137/18M1187611. MR3919409
- [90] Y. Lu, J. Lu, and J. Nolen, *Accelerating Langevin sampling with birth-death*, Preprint, [arXiv:1905.09863](https://arxiv.org/abs/1905.09863), 2019.
- [91] Y. Lu, D. Slepčev, and L. Wang, *Birth-death dynamics for sampling: global convergence, approximations and their asymptotics*, Nonlinearity **36** (2023), no. 11, 5731–5772, DOI 10.1088/1361-6544/acf988. MR4646041
- [92] L. Malagò, L. Montrucchio, and G. Pistone, *Wasserstein Riemannian geometry of Gaussian densities*, Inf. Geom. **1** (2018), no. 2, 137–179, DOI 10.1007/s41884-018-0014-4. MR3993926
- [93] L. Malagò and G. Pistone, *Information geometry of the Gaussian distribution in view of stochastic optimization*, FOGA’15—Proceedings of the 2015 ACM Conference on Foundations of Genetic Algorithms XIII, ACM, New York, 2015, pp. 150–162, DOI 10.1145/2725494.2725510. MR3388148
- [94] D. Maoutsa, S. Reich, and M. Opper, *Interacting particle solutions of Fokker-Planck equations through gradient-log-density estimation*, Entropy **22** (2020), no. 8, Paper No. 802, 35, DOI 10.3390/e22080802. MR4220298
- [95] P. A. Markowich and C. Villani, *On the trend to equilibrium for the Fokker-Planck equation: an interplay between physics and functional analysis*, Mat. Contemp. **19** (2000), 1–29. VI Workshop on Partial Differential Equations, Part II (Rio de Janeiro, 1999). MR1812873
- [96] J. Martens, *New insights and perspectives on the natural gradient method*, J. Mach. Learn. Res. **21** (2020), Paper No. 146, 76. MR4209432
- [97] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas, *A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion*, SIAM J. Sci. Comput. **34** (2012), no. 3, A1460–A1487, DOI 10.1137/110845598. MR2970260
- [98] J. C. Mattingly, A. M. Stuart, and D. J. Higham, *Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise*, Stochastic Process. Appl. **101** (2002), no. 2, 185–232, DOI 10.1016/S0304-4149(02)00150-3. MR1931266
- [99] A. Mielke, D. M. Renger, and M. A. Peletier, *A generalization of Onsager’s reciprocity relations to gradient flows with nonlinear mobility*, J. Non-Equilibrium Thermodyn. **41** (2016), 141–149.
- [100] A. Mielke and J.-J. Zhu, *Hellinger-Kantorovich gradient flows: global exponential decay of entropy functionals*, Preprint, [arXiv:2501.17049](https://arxiv.org/abs/2501.17049), 2025.
- [101] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
- [102] J. A. Nelder and R. Mead, *A simplex method for function minimization*, Comput. J. **7** (1965), no. 4, 308–313, DOI 10.1093/comjnl/7.4.308. MR3363409
- [103] L. Onsager, *Reciprocal relations in irreversible processes. I*, Phys. Rev. **37** (1931), 405.
- [104] L. Onsager, *Reciprocal relations in irreversible processes. II*, Phys. Rev. **38** (1931), 2265.
- [105] F. Otto, *The geometry of dissipative evolution equations: the porous medium equation*, Comm. Partial Differential Equations **26** (2001), 101–174, DOI 10.1081/PDE-100002243.
- [106] G. A. Pavliotis, *Stochastic processes and applications*, Texts in Applied Mathematics, vol. 60, Springer, New York, 2014. Diffusion processes, the Fokker-Planck and Langevin equations, DOI 10.1007/978-1-4939-1323-7. MR3288096
- [107] J. Pidstrigach and S. Reich, *Affine-invariant ensemble transform methods for logistic regression*, Found. Comput. Math. **23** (2023), no. 2, 675–708, DOI 10.1007/s10208-022-09550-2. MR4568201

- [108] H. Poincare, *Sur les Equations aux Derivees Partielles de la Physique Mathematique* (French), Amer. J. Math. **12** (1890), no. 3, 211–294, DOI 10.2307/2369620. MR1505534
- [109] M. Quiroz, D. J. Nott, and R. Kohn, *Gaussian variational approximations for high-dimensional state space models*, Bayesian Anal. **18** (2023), no. 3, 989–1016, DOI 10.1214/22-ba1332. MR4626365
- [110] C. Radhakrishna Rao, *Information and the accuracy attainable in the estimation of statistical parameters*, Bull. Calcutta Math. Soc. **37** (1945), 81–91. MR15748
- [111] C. E. Rasmussen, *Gaussian Processes in Machine Learning*, Summer school on machine learning, Springer, 2003, 63–71.
- [112] S. Reich and C. Cotter, *Probabilistic Forecasting and Bayesian Data Assimilation*, Cambridge University Press, New York, 2015, DOI 10.1017/CBO9781107706804. MR3242790
- [113] J. Revels, M. Lubin, and T. Papamarkou, *Forward-mode automatic differentiation in Julia*, [arXiv:1607.07892](https://arxiv.org/abs/1607.07892) [cs.MS], 2016.
- [114] T. Salimans, H. Zhang, A. Radford, and D. Metaxas, *Improving GANs using optimal transport*, Preprint, [arXiv:1803.05573](https://arxiv.org/abs/1803.05573), 2018.
- [115] F. Santambrogio, *{Euclidean, metric, and Wasserstein} gradient flows: an overview*, Bull. Math. Sci. **7** (2017), no. 1, 87–154, DOI 10.1007/s13373-017-0101-1. MR3625852
- [116] S. Särkkä, *On unscented Kalman filtering for state estimation of continuous-time nonlinear systems*, IEEE Trans. Automat. Control **52** (2007), no. 9, 1631–1641, DOI 10.1109/TAC.2007.904453. MR2352439
- [117] T. Schneider, S. Lan, A. Stuart, and J. Teixeira, *Earth system modeling 2.0: a blueprint for models that learn from observations and targeted high-resolution simulations*, Geophys. Res. Lett. **44** (2017), 12–396.
- [118] R. Shen and Y. T. Lee, *The randomized midpoint method for log-concave sampling*, Adv. Neural Inf. Process. Syst. **32** (2019).
- [119] Z. Shen, Z. Wang, S. Kale, A. Ribeiro, A. Karbasi, and H. Hassani, *Self-consistency of the Fokker-Planck equation*, Preprint, [arXiv:2206.00860](https://arxiv.org/abs/2206.00860), 2022.
- [120] Y. Song, S. Garg, J. Shi, and S. Ermon, *Sliced score matching: a scalable approach to density and score estimation*, Uncertainty in Artificial Intelligence, PMLR, 2020, 574–584.
- [121] H. W. Sorenson (ed.), *Kalman Filtering: Theory and Application*, IEEE Press Selected Reprint Series, IEEE Press, New York, 1985. MR847460
- [122] B. Sprungk, S. Weissmann, and J. Zech, *Metropolis-adjusted interacting particle sampling*, Stat. Comput. **35** (2025), no. 3, Paper No. 64, 31, DOI 10.1007/s11222-025-10595-w. MR4877495
- [123] A. Srivastava, I. Jermyn, and S. Joshi, *Riemannian analysis of probability density functions with applications in vision*, 2007 IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2007, 1–8.
- [124] A. Takatsu, *Wasserstein geometry of Gaussian measures*, Osaka J. Math. **48** (2011), no. 4, 1005–1026. MR2871291
- [125] S. Vempala and A. Wibisono, *Rapid convergence of the unadjusted Langevin algorithm: isoperimetry suffices*, Adv. Neural Inf. Process. Syst. **32** (2019).
- [126] C. Villani, *Topics in Optimal Transportation*, Graduate Studies in Mathematics, vol. 58, American Mathematical Society, Providence, RI, 2003, DOI 10.1090/gsm/058. MR1964483
- [127] M. J. Wainwright, M. I. Jordan, et al., *Graphical models, exponential families, and variational inference*, Found. Trends Signal Process. **1** (2008), 1–305.
- [128] E. A. Wan and R. Van Der Merwe, *The unscented Kalman filter for nonlinear estimation*, Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373), IEEE, 2000, 153–158.
- [129] Y. Wang, P. Chen, and W. Li, *Projected Wasserstein gradient descent for high-dimensional Bayesian inference*, SIAM/ASA J. Uncertain. Quantif. **10** (2022), no. 4, 1513–1532, DOI 10.1137/21M1454018. MR4512980
- [130] Y. Wang, P. Chen, M. Pilanci, and W. Li, *Optimal neural network approximation of Wasserstein gradient direction via convex optimization*, SIAM J. Math. Data Sci. **6** (2024), no. 4, 978–999, DOI 10.1137/23M1573173. MR4807008
- [131] Y. Wang and W. Li, *Information Newton’s flow: second-order optimization method in probability space*, Preprint, [arXiv:2001.04341](https://arxiv.org/abs/2001.04341), 2020.
- [132] Y. Wang and W. Li, *Accelerated information gradient flow*, J. Sci. Comput. **90** (2022), no. 1, Paper No. 11, 47, DOI 10.1007/s10915-021-01709-3. MR4342674

- [133] S. Weissmann, N. K. Chada, C. Schillings, and X. T. Tong, *Adaptive Tikhonov strategies for stochastic ensemble Kalman inversion*, Inverse Problems **38** (2022), no. 4, Paper No. 045009, 34, DOI 10.1088/1361-6420/ac5729. MR4399530
- [134] A. Wibisono, *Sampling as optimization in the space of measures: the Langevin dynamics as a composite optimization problem*, Conference on Learning Theory, PMLR, 2018, 2093–3027.
- [135] K. Wu, S. Schmidler, and Y. Chen, *Minimax mixing time of the metropolis-adjusted Langevin algorithm for log-concave sampling*, J. Mach. Learn. Res. **23** (2022), Paper No. [270], 63. MR4577709
- [136] K.-V. Yuen, *Bayesian Methods for Structural Dynamics and Civil Engineering*, John Wiley & Sons, 2010.
- [137] H. Yumei, W. Xuezi, P. Quan, H. Zhentao, and B. Moran, *Variational Bayesian Kalman filter using natural gradient*, Chinese J. Aeronautics **35** (2022), 1–10.
- [138] O. Zahm, T. Cui, K. Law, A. Spantini, and Y. Marzouk, *Certified dimension reduction in nonlinear Bayesian inverse problems*, Math. Comp. **91** (2022), no. 336, 1789–1835, DOI 10.1090/mcom/3737. MR4435948
- [139] G. Zhang, J. Martens, and R. B. Grosse, *Fast convergence of natural gradient descent for over-parameterized neural networks*, Adv. Neural Inf. Process. Syst., **32** (2019).
- [140] K. S. Zhang, G. Peyré, J. Fadili, and M. Pereyra, *Wasserstein control of mirror Langevin Monte Carlo*, Conference on Learning Theory, PMLR, 2020, 3814–3841.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, CALIFORNIA; AND
APPLIED AND COMPUTATIONAL MATHEMATICS, CALIFORNIA INSTITUTE OF TECHNOLOGY,
PASADENA, CALIFORNIA

Email address: `yifanchen@math.ucla.edu`

BEIJING INTERNATIONAL CENTER FOR MATHEMATICAL RESEARCH, PEKING UNIVERSITY, BEI-
JING, PEOPLE’S REPUBLIC OF CHINA

Email address: `huangdz@bicmr.pku.edu.cn`

DEPARTMENT OF STATISTICS AND DATA SCIENCE, UNIVERSITY OF PENNSYLVANIA, PHILADEL-
PHIA, PENNSYLVANIA

Email address: `huangjy@wharton.upenn.edu`

INSTITUT FÜR MATHEMATIK, UNIVERSITÄT POTSDAM, POTSDAM, GERMANY

Email address: `sebastian.reich@uni-potsdam.de`

APPLIED AND COMPUTATIONAL MATHEMATICS, CALIFORNIA INSTITUTE OF TECHNOLOGY,
PASADENA, CALIFORNIA

Email address: `astuart@caltech.edu`