







# Parameterized Wasserstein gradient flow

Yijie Jin<sup>a, ,\*</sup>, Shu Liu<sup>b, </sup>, Hao Wu<sup>a</sup>, Xiaojing Ye<sup>c, </sup>, Haomin Zhou<sup>a, </sup>

<sup>a</sup> School of Mathematics, Georgia Institute of Technology, Atlanta, GA, USA

<sup>b</sup> Department of Mathematics, University of California, Los Angeles, CA, USA

<sup>c</sup> Department of Mathematics and Statistics, Georgia State University, Atlanta, GA, USA

## ARTICLE INFO

### Keywords:

Wasserstein gradient flow  
Fokker-Planck equation  
Porous medium equation  
Deep neural networks  
Numerical analysis

## ABSTRACT

We develop a fast and scalable numerical approach to solve Wasserstein gradient flows (WGFs), which is particularly suitable for high-dimensional cases. Our approach is to use general reduced-order models, like deep neural networks, to parameterize the push-forward maps such that they can push a simple reference density to the one solving the given WGF. The new dynamical system is called parameterized WGF (PWGF), and it is defined on the finite-dimensional parameter space equipped with a pullback Wasserstein metric. Our numerical scheme can approximate the solutions of WGFs for general energy functionals effectively, without requiring spatial discretization or nonconvex optimization procedures, thus avoiding some limitations of classical numerical methods and more recent deep learning-based approaches. A comprehensive analysis of the approximation errors measured by Wasserstein distance is also provided in this work. Numerical experiments show promising computational efficiency and verified accuracy on a variety of WGF examples using our approach.

## 1. Introduction

Wasserstein gradient flow (WGF) is a powerful tool for understanding and analyzing density evolution processes. In the seminal work [20] by Jordan, Kinderlehrer and Otto, they showed that Fokker-Planck equation (FPE) is essentially the gradient flow of the relative entropy functional under the Wasserstein metric. Since then, WGFs have shown extensive applications in optimal transport theory, optimization problems, Fokker-Planck equation, porous medium equation, and more [11,29,32]. However, numerical computation of general WGFs remains a challenging problem, especially when the state space is of high dimension. Furthermore, it is often desirable to find a sampler that generates samples following the solution of WGF rather than the actual density function solving WGF in many real-world statistics and machine learning applications.

In this paper, we focus on the numerical computation of WGF to address the aforementioned issues. Let  $\mathcal{M}$  be a smooth manifold without boundary. For simplicity, we assume  $\mathcal{M} = \mathbb{R}^d$  throughout the present work, while generalization to general manifolds is straightforward. We also omit subscript  $\mathcal{M}$  for all integrals unless otherwise noted. Denote the density function space defined on  $\mathcal{M}$  as

$$\mathcal{P}(\mathcal{M}) = \left\{ \rho : \mathcal{M} \rightarrow \mathbb{R} : \int \rho(x) dx = 1, \rho(x) \geq 0, \int |x|^2 \rho(x) dx < \infty \right\}. \quad (1.1)$$

Suppose  $\mathcal{P}(\mathcal{M})$  is equipped with the Wasserstein-2 distance [23,32]:

\* Corresponding author.

E-mail addresses: [yijiejin@gatech.edu](mailto:yijiejin@gatech.edu) (Y. Jin), [shuliu@math.ucla.edu](mailto:shuliu@math.ucla.edu) (S. Liu), [hwwu406@gmail.com](mailto:hwwu406@gmail.com) (H. Wu), [xye@gsu.edu](mailto:xye@gsu.edu) (X. Ye), [hmzhou@math.gatech.edu](mailto:hmzhou@math.gatech.edu) (H. Zhou).

<https://doi.org/10.1016/j.jcp.2024.113660>

Received 27 May 2024; Received in revised form 2 December 2024; Accepted 4 December 2024

$$W(\rho_1, \rho_2) = \left( \inf_{\pi \in \Pi(\rho_1, \rho_2)} \iint |x - y|^2 d\pi(x, y) \right)^{1/2} \tag{1.2}$$

for any  $\rho_1, \rho_2 \in \mathcal{P}(\mathcal{M})$ , where  $\Pi(\rho_1, \rho_2)$  is the joint density space with  $\rho_1$  and  $\rho_2$  as marginals. Then  $\mathcal{P}(\mathcal{M})$  becomes an infinite-dimensional Riemannian manifold with  $W$  inducing its Riemannian metric, and a WGF can be written in the following general form:

$$\frac{\partial \rho}{\partial t} = -\text{grad}_W \mathcal{F}(\rho), \quad \rho(0, x) = \rho_0(x), \tag{1.3}$$

where  $x \in \mathcal{M}$ ,  $\rho_0$  is a given initial probability density,  $\mathcal{F} : \mathcal{P}(\mathcal{M}) \rightarrow \mathbb{R}$  is some energy functional defined on  $\mathcal{P}(\mathcal{M})$ , and  $\text{grad}_W$  stands for the gradient of functionals on  $\mathcal{P}(\mathcal{M})$  with respect to the Wasserstein metric. Many well-known probability evolution equations, such as the Fokker-Planck equation and porous medium equation, are essentially WGF (1.3) with specific choices of  $\mathcal{F}$ .

In recent years, there has been a surge of interest to compute WGF (1.3) numerically. In [20], a proximal point algorithm is used to solve the Fokker-Planck equation as a special case of WGF. Later, various numerical methods have been developed to solve the WGF; see [9,14,30,31,33]. Among them, some use classical methods such as finite difference [25] and finite element method [6]. They are applicable to solving WGF with low dimension  $d$ . This limitation is due to their spatial discretization and hence they suffer the issue known as ‘‘curse-of-dimensionality’’ as the number of unknowns increases exponentially fast in terms of problem dimension  $d$ . Sampling-based approaches [14,30] can be used to solve WGF in high dimension, however, the computation for the minimization and energy functional evaluation used in those algorithms may become time-consuming.

The goal of this work is to develop a fast computational scheme of (1.3) in high-dimensional settings. Similar to [14,29,30], we seek for a time-dependent push-forward map to generate samples whose density function solves (1.3). However, compared to these existing methods, our approach features several notable differences: (i) We use a *new pullback Wasserstein metric* that can be computed efficiently while still maintaining the accuracy up to the same order as the results obtained in [29]; (ii) Our algorithm can be used to compute general WGF with a large class of (possibly nonlinear)  $\mathcal{F}$ ; (iii) We parameterize the push-forward map using neural ordinary differential equation [10] and show that it is computationally convenient to handle functional of probability densities in our case. As a result, we obtain an approximate solution to the WGF as well as a sample generator of the solution distribution; and (iv) Our numerical scheme does not need any spatial discretization and thus is scalable for problems defined on high-dimensional spaces (e.g.,  $d \geq 15$ ). Moreover, our method does not need to solve any nonconvex optimization problem as in typical deep neural network training. Instead, it only requires fast matrix-vector products where the vector size is determined by the number of parameters in the selected push-forward map.

## 2. Related work

The Fokker-Planck equation is first interpreted as a special case of WGF in [20]. It is shown that the Fokker-Planck equation can be viewed as the gradient flow on the Wasserstein manifold of a special energy functional, which consists of a linear potential energy and a relative entropy. For general (possibly nonlinear) energy functionals, one can apply similar derivations to obtain their corresponding WGFs. As WGFs are essentially time-evolution partial differential equations (PDEs), classical numerical methods such as finite difference and finite element methods can be applied with proper modifications [3,8,9,27,33]. However, the application of these classical methods is limited to low-dimensional cases due to the curse of dimensionality.

In recent years, sampling-based approaches have been proposed as a promising alternative solution method [14,19,24,29,30]. Many of these methods start from the numerical scheme used in [20] known as the Jordan-Kinderlehrer-Otto (JKO) scheme, which is a proximal algorithm in the optimization context. By utilizing the Benamou-Brenier formula, we know that for any time step, the optimal push-forward map in the JKO scheme can be expressed as the gradient of some convex function. Motivated by this fact, several recent works are devoted to solving the WGF by approximating the optimal push-forward map by neural networks  $T_\theta$ . However, there are two main challenges in this approach: one is the evaluation of potential energy  $\mathcal{F}$ , especially when  $\mathcal{F}$  involves the density function  $\rho$  explicitly. As mentioned in [26], this problem can be addressed by introducing an extra optimization procedure. However, such a computation is generally expensive; the other is the computational challenge for solving the JKO scheme, which is an optimization problem on the density manifold.

In contrast to those existing methods, our method is mostly motivated by [29], in which the authors develop the parameterization to solve FPE. This approach exploits WGFs in the parameter space, in which the parametric functions, such as deep neural networks, are used to parameterize the push-forward maps. We design a numerical scheme for the parameter dynamics. The main challenge in our formulation is the computation of the Wasserstein gradient on the parameter manifold. We show that it can be effectively approximated once introducing the new pullback Wasserstein metric as discussed in (3.17). Furthermore, we employ several powerful neural networks, such as the normalizing flow [21] and continuous normalizing flow [10], which are both convenient and efficient in evaluating the energy functionals involving push-forward densities and their gradients.

In addition to the aforementioned literatures, it is worth mentioning that the study on the Wasserstein information matrix as the metric tensor defined on the parameter manifold has been introduced in [28]; Some numerical analysis results on 1-dimensional neural projected WGFs have been reported in [41]; Studies on WGFs on the manifold of Gaussian distributions have been conducted in [11,40] and the references therein.

Furthermore, a series of deep-learning methods [1,7,13,16] have been composed to compute general time-evolution PDEs. Compared with these references which directly approximate the PDE solutions via neural networks, our algorithm is developed on the space of push-forward maps and utilizes the geometric principle of the WGFs. Our treatment respects the essential properties of WGFs

such as positivity, conservation of mass, and energy dissipation: the first two properties are naturally preserved on push-forward models; the energy dissipation property is justified via our numerical experiments demonstrated in Section 4.

### 3. Parameterization of Wasserstein gradient flow

In this section, we briefly review the derivation of WGF and several examples of WGF including Fokker-Planck and porous media equations, as well as the pullback Wasserstein metric in parameter space. Then we present our approach to solve WGF numerically and provide a comprehensive analysis of the proposed scheme.

#### 3.1. Background on Wasserstein gradient flow

We denote the tangent space of  $\mathcal{P}(\mathcal{M})$  at  $\rho$  by

$$\mathcal{T}_\rho \mathcal{P}(\mathcal{M}) = \left\{ \sigma \in C^\infty(\mathcal{M}) : \int \sigma(x) dx = 0 \right\},$$

which is the same for different  $\rho$  [32]. For any specific  $\rho \in \mathcal{P}(\mathcal{M})$ , the Wasserstein metric tensor  $g^W$  is a positive definite bilinear form on the tangent bundle  $\mathcal{TP}(\mathcal{M}) = \{(\rho, \sigma) : \rho \in \mathcal{P}(\mathcal{M}), \sigma \in \mathcal{T}_\rho \mathcal{P}(\mathcal{M})\}$  defined by:

$$g^W(\rho)(\sigma_1, \sigma_2) = \int \nabla \Phi_1(x) \cdot \nabla \Phi_2(x) \rho(x) dx, \quad \forall \sigma_i \in \mathcal{T}_\rho \mathcal{P}(\mathcal{M}), i = 1, 2$$

where  $\sigma_i = -\nabla \cdot (\rho \nabla \Phi_i)$  for  $i = 1, 2$ . Suppose  $\mathcal{F} : \mathcal{P}(\mathcal{M}) \rightarrow \mathbb{R}$  is a smooth functional defined on  $\mathcal{P}(\mathcal{M})$ , its Riemannian gradient over  $(\mathcal{P}(\mathcal{M}), g^W)$  is given as follows:

$$\text{grad}_W \mathcal{F}(\rho) = g^W(\rho)^{-1} \left( \frac{\delta \mathcal{F}}{\delta \rho} \right) (x) = -\nabla \cdot \left( \rho(x) \nabla \frac{\delta \mathcal{F}}{\delta \rho} (x) \right), \tag{3.1}$$

where  $\frac{\delta \mathcal{F}}{\delta \rho}$  is the first variation of  $\mathcal{F}$  in the  $L^2$  sense. The Wasserstein gradient flow (WGF) of  $\mathcal{F}$  is given by

$$\frac{\partial \rho}{\partial t} = -\text{grad}_W \mathcal{F}(\rho). \tag{3.2}$$

Many well-known equations can be formulated as WGF with selected energy functionals. Here are three examples.

**Example 3.1 (Fokker-Planck equation).** Let  $V \in C^2(M)$  be a given potential function and  $\rho_*(x) = \frac{1}{Z_D} e^{-V(x)/D}$  be the corresponding Gibbs distribution, where  $D > 0$  is a fixed constant and  $Z_D = \int e^{-\frac{V(x)}{D}} dx$  is the normalization constant. Suppose  $\mathcal{F}(\rho)$  is the relative entropy with respect to  $\rho_*$  scaled by  $D$ , i.e.,

$$\mathcal{F}(\rho) = D \mathcal{H}_{KL}(\rho \| \rho_*) \quad \text{where} \quad \mathcal{H}_{KL}(\rho \| \rho_*) := \left( \int \frac{1}{D} V(x) \rho(x) + \rho(x) \log \rho(x) dx \right) + \log Z_D. \tag{3.3}$$

Then the WGF (3.2) with  $\mathcal{F}$  defined in (3.3) becomes the Fokker-Planck equation:

$$\frac{\partial \rho}{\partial t} = -\text{grad}_W \mathcal{F}(\rho) = \nabla \cdot (\rho \nabla V) + D \nabla \cdot (\rho \nabla \log \rho) = \nabla \cdot (\rho \nabla V) + D \Delta \rho. \tag{3.4}$$

The relative entropy  $\mathcal{H}_{KL}(\rho \| \rho_*)$  is closely related to the Fisher information defined as

$$\mathcal{I}(\rho | \rho_*) = \int \left\| \nabla \log \left( \frac{\rho(x)}{\rho_*(x)} \right) \right\|^2 \rho(x) dx. \tag{3.5}$$

It provides an upper bound to  $\mathcal{H}_{KL}$  and further guarantees the uniform convergence of the dynamics (3.4) as shown in the following theorem, which will be used in our error estimation later.

**Theorem 3.2 (Holley–Stroock perturbation [18]).** Suppose the potential  $V$  can be decomposed as  $V = U + \phi$  where  $\nabla^2 U \geq KI$  for some  $K > 0$  and  $\phi \in L^\infty$ . Denote  $\text{osc}(\phi) := \sup \phi - \inf \phi < \infty$ . Then the following logarithmic Sobolev inequality holds

$$\mathcal{H}_{KL}(\rho \| \rho_*) \leq \frac{e^{\text{osc}(\phi)}}{K} \mathcal{I}(\rho | \rho_*) \tag{3.6}$$

for any  $\rho \in \mathcal{P}(M)$ . Assume  $\rho$  solves equation (3.4) with initial value  $\rho(0, \cdot) = \rho_0(x)$ , then

$$\mathcal{H}_{KL}(\rho_t \| \rho_*) \leq \mathcal{H}_{KL}(\rho_0 \| \rho_*) e^{-\frac{DK}{\text{osc}(\phi)} t}, \quad \forall t > 0. \tag{3.7}$$

**Example 3.3 (Porous medium equation).** Consider the energy functional  $\mathcal{F}$  defined by

$$\mathcal{F}(\rho) = \frac{1}{m-1} \int \rho^m(x) dx, \tag{3.8}$$

for an integer  $m > 1$ . The corresponding WGF becomes the porous medium equation (PME), which is a nonlinear heat equation given by

$$\frac{\partial \rho}{\partial t} = \Delta(\rho^m). \tag{3.9}$$

There are many applications of PME in physical problems, for example, heat transfer or diffusion, gas flow [36], and flow of reactive fluids in porous media [22].

**Example 3.4 (Aggregation model).** Consider the interaction energy functional

$$F(\rho) = \frac{1}{2} \iint J(|x - y|) \rho(x) \rho(y) dx dy \tag{3.10}$$

which is used to model the interactive behavior of swarm of particles [15]. Here  $J$  is an interaction kernel consisting of repulsive and attractive parts. A typical choice for  $J$  is

$$J(x) = \frac{|x|^a}{a} - \frac{|x|^b}{b}, \quad x \in \mathcal{M}, \tag{3.11}$$

where  $a > b > 0$  are positive constants. Its corresponding WGF is

$$\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \nabla (J * \rho)). \tag{3.12}$$

We note that Keller-Siegel system is a special case of aggregation model [5]. We remark that the WGF (3.2) can be associated with a particle-level dynamic, as given in the following proposition.

**Proposition 3.5.** [29] Assume  $X \in \mathcal{M}$  is a random process and solves the following equation,

$$\dot{X} = -\nabla_X \frac{\delta}{\delta \rho} F(\rho(X), X) \tag{3.13}$$

where  $\rho$  is the density of  $X$ . Then  $\rho$  solves the WGF (3.2).

The particle-level dynamic (3.13) provides a physical interpretation of WGF. More importantly, we will use it to establish the error bound of our approximation using parameterized WGF below.

### 3.2. Push-forward maps and parameterized Fokker-Planck equation

In [29], the authors developed a method to parameterize push-forward maps in order to approximately solve the FPE (3.4).

Fix any reference probability distribution  $\lambda$  that is absolutely continuous with respect to the standard Lebesgue measure  $\mu$  on  $\mathcal{M}$ , we use  $\rho = d\lambda/d\mu$ , the Radon-Nikodym derivative of  $\lambda$  with respect to  $\mu$ , to denote the reference density determined by  $\lambda$ . Then for any push-forward map  $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ , it induces a new probability measure which is absolutely continuous to the Lebesgue measure  $\mu$ . Hence, there exists a probability density  $T_{\#}\rho$  on  $\mathbb{R}^d$  defined by

$$\int_E T_{\#}\rho(x) d\mu(x) = \int_{T^{-1}(E)} \rho(z) d\mu(z) = \int_{T^{-1}(E)} d\lambda(z) = \lambda(T^{-1}(E)) \quad \text{for all measurable } E \subset \mathbb{R}^d,$$

where  $T^{-1}(E)$  is the pre-image of  $E$ . Hereafter we use  $dx$  instead of  $d\mu(x)$  to reduce notation complexity.

Let us take  $T$  as parameterized map, namely for any  $\theta \in \Theta$ ,  $T_{\theta} : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is a parametric function with parameter  $\theta$ . Here  $\Theta$ , as a subset of  $\mathbb{R}^n$ , is called the *parameter space*, where  $n$  is the number of parameters of  $T_{\theta}$  (i.e., the dimension of  $\theta$ ). Typical examples of  $T_{\theta}$  include Fourier expansion, finite element approximation, and (deep) neural networks.

The map  $T_{\#} : \Theta \rightarrow \mathcal{P}$  given by  $\theta \mapsto T_{\theta\#}\rho$  naturally defines an immersion from  $\Theta$  to the probability manifold  $\mathcal{P}$ . Collecting all parameterized distributions together, i.e.,

$$\mathcal{P}_{\Theta} = \{ \rho_{\theta} = T_{\theta\#}\rho : \theta \in \Theta \},$$

we obtain a finite-dimensional submanifold  $\mathcal{P}_{\Theta}$  of  $\mathcal{P}$ . We can define the tangent space of  $\mathcal{P}_{\Theta}$  at each  $\theta = (\theta_1, \dots, \theta_n) \in \Theta \subset \mathbb{R}^n$  as  $T_{\rho_{\theta}} \mathcal{P}_{\Theta} = \text{span} \{ \frac{\partial \rho_{\theta}}{\partial \theta_1}, \dots, \frac{\partial \rho_{\theta}}{\partial \theta_n} \}$ . The tangent bundle is then  $\mathcal{T} \mathcal{P}_{\Theta} = \cup_{\theta \in \Theta} \{ \rho_{\theta} \} \times T_{\rho_{\theta}} \mathcal{P}_{\Theta}$ . On the other hand, the cotangent space  $T_{\rho_{\theta}}^* \mathcal{P}_{\Theta}$  is the dual space of  $T_{\rho_{\theta}} \mathcal{P}_{\Theta}$ , and the cotangent bundle is  $\mathcal{T}^* \mathcal{P}_{\Theta} = \cup_{\theta \in \Theta} \{ \rho_{\theta} \} \times T_{\rho_{\theta}}^* \mathcal{P}_{\Theta}$ .

Note that  $\mathcal{P}_{\Theta}$  is a finite-dimensional subset of  $\mathcal{P}$ , hence we can pull back the Wasserstein metric  $g^W$  to  $\mathcal{P}_{\Theta}$ . It is shown in [29] that the pullback Wasserstein metric  $G(\theta) = T_{\theta\#}^* g^W$  on  $\mathcal{P}$  is given by

$$G(\theta) = \int \nabla \Psi \circ T_{\theta}(z) \nabla \Psi \circ T_{\theta}(z)^{\top} d\lambda(z), \tag{3.14}$$

where  $\Psi_\theta = (\psi_1, \dots, \psi_n)^\top$ , and  $\psi_j$  is solved from the following elliptic equation,

$$\nabla \cdot (\rho_\theta \nabla \psi_j(x)) = \nabla \cdot \left( \rho_\theta \frac{\partial T_\theta}{\partial \theta_j} \circ T_\theta^{-1}(x) \right). \quad (3.15)$$

The parameterized version of FPE on parameter manifold  $\Theta$  can be expressed as

$$\dot{\theta} = -G(\theta)^{-1} \nabla_\theta F(\theta), \quad (3.16)$$

where  $F(\theta) := \mathcal{F}(\rho_\theta)$  and  $\mathcal{F}$  is the scaled relative entropy defined in (3.3).

If directly computing the parameterized FPE (3.16), the cost is intractable in higher dimensions because it requires solving  $n$  elliptic equations (3.15) to obtain  $G(\theta)$ . To alleviate this difficulty, authors in [29] derived a minimax formulation that can advance the dynamics without solving the elliptic equations. Although the cost becomes manageable, it can still be high if the spatial dimension  $d$  is not small. In this work, we advocate a new, yet naturally derived pullback metric  $\hat{G}$  to replace  $G$ , and develop a computational framework to solve general WGF. Details shall be provided in the next subsection.

### 3.3. Parameterized WGF with a new pullback Wasserstein metric

Following the study in [39], we replace  $G(\theta)$  by a new pullback Wasserstein metric defined as

$$\hat{G}(\theta) = \int \partial_\theta T_\theta(z)^\top \partial_\theta T_\theta(z) d\lambda(z). \quad (3.17)$$

For any energy functional  $\mathcal{F}$  with smooth variation  $\frac{\delta \mathcal{F}}{\delta \rho}$ , the corresponding parameterized WGF is given by

$$\dot{\theta} = -\hat{G}(\theta)^\dagger \nabla_\theta F(\theta), \quad (3.18)$$

where  $\hat{G}(\theta)^\dagger$  is the Penrose-Moore pseudo inverse of  $\hat{G}(\theta)$ .

The motivation of introducing  $\hat{G}(\theta)$  is three-folded:

- The definition is directly inspired by (3.14) and (3.15). In the 1-dimensional case, the new metric  $\hat{G}$  coincides with the exact matrix  $G$ , and the proof is shown in [29].
- A deeper motivation is influenced by [32]. The geometry of  $(\Theta, \hat{G}(\theta))$  is an isometric embedding into the flat Riemannian space  $(\mathcal{O}, g^\mathcal{O})$  introduced in [32]. Specifically, let  $\rho \in P(\mathcal{M})$  be a fixed density function and  $\mathcal{O}$  the set of diffeomorphisms  $T : \mathcal{M} \rightarrow \mathcal{M}$ , and define the Riemannian metric  $g^\mathcal{O}$  on  $\mathcal{O}$  at  $T$  by  $g^\mathcal{O}(T)(v_1, v_2) := \int v_1 v_2 \rho dx$  for any  $v_1, v_2 \in \mathcal{T}_T \mathcal{O}$ . Here  $\mathcal{T}_T \mathcal{O}$  is the tangent space of  $\mathcal{O}$  at  $T$ , which is set to be the space of all vector fields on  $\mathcal{M}$  for every  $T$  in [32]. Then it is shown that the mapping  $\Pi : \theta \mapsto \rho_\theta$  yields an isometric submersion from  $(\Theta, g^\Theta)$  to  $(P(\mathcal{M}), g)$ , where  $g$  is the 2-Wasserstein metric. Consider a (deep) neural network structure  $T_\theta$  with parameter  $\theta$  such that  $T_\theta : \mathcal{M} \rightarrow \mathcal{M}$  is a diffeomorphism for every  $\theta \in \Theta$  and  $T_\theta(\cdot) \in C^1(\Theta \times \mathcal{M}; \mathbb{R})$ . Consider the map  $\mathcal{T} : \Theta \rightarrow \mathcal{O}$  defined by  $\theta \mapsto \mathcal{T}(\theta) := T_\theta$ , we define the pullback metric on  $\Theta$  as

$$\hat{G}(\theta) = \mathcal{T}^* g^\mathcal{O}(T_\theta) \quad (3.19)$$

where  $\mathcal{T}^*$  is the pullback operation induced by  $\mathcal{T}$ . To obtain  $\hat{G}(\theta)$ , consider any curve  $\{\theta(t)\}_{- \epsilon \leq t \leq \epsilon}$  for some  $\epsilon > 0$  passing through  $\theta(0)$  at  $t = 0$  and denote  $\dot{\theta}(0) = \frac{d}{dt} \theta(t)|_{t=0}$ . By the definitions of the pullback operation and the metric  $g^\mathcal{O}$  above, we have

$$\begin{aligned} \hat{G}(\theta(0))(\dot{\theta}(0), \dot{\theta}(0)) &= g^\mathcal{O}(T_\theta) \left( \frac{d}{dt} T_{\theta(t)}|_{t=0}, \frac{d}{dt} T_{\theta(t)}|_{t=0} \right) \\ &= \int (\dot{\theta}(0)^\top \partial_\theta T_{\theta(0)}(z)^\top \partial_\theta T_{\theta(0)}(z) \dot{\theta}(0)) \rho(z) dz \\ &= \dot{\theta}(0)^\top \left( \int \partial_\theta T_{\theta(0)}(z)^\top \partial_\theta T_{\theta(0)}(z) d\lambda(z) \right) \dot{\theta}(0), \end{aligned}$$

which implies (3.17) since  $\theta(0)$  and  $\dot{\theta}(0)$  are arbitrary. The matrix  $\hat{G}(\theta) = (\hat{G}(\theta)_{ij})_{1 \leq i, j \leq m}$  has components

$$\hat{G}(\theta)_{ij} = \sum_{k=1}^d \int \partial_{\theta_i} T_\theta^{(k)}(z) \cdot \partial_{\theta_j} T_\theta^{(k)}(z) d\lambda(z) \quad (3.20)$$

where  $T_\theta^{(k)} : \mathcal{M} \rightarrow \mathbb{R}$  is the  $k$ -th component of  $T_\theta : \mathcal{M} \rightarrow \mathcal{M}$ . This is also discussed in [38].

- Compared with (3.14), the new metric is more computationally efficient. Directly evaluating the metric tensor  $G$  requires solving  $n$  different elliptic PDEs where  $n$  is the number of the parameters in the push-forward map  $T_\theta$ , and  $m$  can be large if we choose  $T_\theta$  to be a deep neural network. To circumvent the computation complexity challenge, a bi-level minimization scheme was proposed in [29]. However, it may still be computationally expensive to solve such optimization problems in general. We demonstrate the difference in the computation time in the experiment for computing the solution of the Fokker-Planck equation in Section 5.1.

The gradient  $\nabla_\theta F(\theta)$  can be evaluated through the following formulation as derived in [39],

$$\nabla_\theta F(\theta) = \int \partial_\theta T_\theta(z)^\top \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_\theta(z) \, d\lambda(z). \quad (3.21)$$

Following [39], the formula (3.21) suggests that  $\nabla_\theta F(\theta)$  is in the range of  $\widehat{G}(\theta)$  even when  $\widehat{G}$  is not invertible. Therefore, the right hand side of (3.18) is always well defined. For convenience, we still call (3.18) the parameterized WGF, or again PWGF for short.

It is interesting to consider the particle level dynamics corresponding to (3.18). For a fix  $z_0 \in \mathcal{M}$ ,  $Y = T_\theta(z_0)$  is the push-forward point in  $\mathcal{M}$ . When  $\theta$  varies as a function of  $t$  according to (3.18),  $Y(t)$  forms a curve in  $\mathcal{M}$ , and it satisfies

$$\dot{Y} = \partial_\theta T_\theta(z_0) \dot{\theta} = -\partial_\theta T_\theta(z_0) \widehat{G}(\theta)^\dagger \nabla_\theta F(\theta), \quad Y_0 = T_{\theta(0)}(z_0). \quad (3.22)$$

To further investigate the properties of parameterized particle dynamics (3.22), we introduce the kernel operator  $\mathcal{K}$ : for any  $f \in L^2(\mathcal{M}; \mathcal{M}, \lambda)$ ,  $\mathcal{K}_\theta[f] \in L^2(\mathcal{M}; \mathcal{M}, \lambda)$  is defined as

$$\mathcal{K}_\theta[f](\cdot) = \partial_\theta T_\theta(\cdot) \widehat{G}(\theta)^\dagger \int \partial_\theta T_\theta(z)^\top f(z) \, d\lambda(z) = \int K_\theta(\cdot, z) f(z) \, d\lambda(z), \quad (3.23)$$

where the kernel matrix  $K_\theta(z', z) \in \mathbb{R}^{d \times d}$  is defined by

$$K_\theta(z', z) := \partial_\theta T_\theta(z') \widehat{G}(\theta)^\dagger \partial_\theta T_\theta(z)^\top. \quad (3.24)$$

It is shown that  $\mathcal{K}_\theta[f]$  is the orthogonal projection of  $f$  onto the tangent space spanned by  $\partial_\theta T_\theta$  [39]. Combining (3.21), (3.22) and (3.23), we obtain

$$\begin{aligned} \dot{Y} &= -\partial_\theta T_\theta(z_0) \widehat{G}(\theta)^\dagger \int \partial_\theta T_\theta(z)^\top \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_\theta(z) \, d\lambda(z) \\ &= -\mathcal{K}_\theta \left[ \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_\theta \right] (z_0) \\ &= -\mathcal{K}_\theta \left[ \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_\theta \right] (T_\theta^{-1}(Y)). \end{aligned} \quad (3.25)$$

The particle-level dynamics (3.25) is determined by the time-evolving push-forward map  $T_\theta$ . We will use it to derive an upper bound on the error of PWGF (3.18) in the next section.

### 3.4. Error bounds in Wasserstein metric

We provide an error analysis of the proposed PWGF formulation. Specifically, we establish an upper bound on the difference, measured by Wasserstein distance, between the approximation  $\rho_{\theta(t)}(\cdot) = T_{\theta(t)\#} \rho(\cdot)$  where  $\theta(t)$  solves PWGF and the true solution  $\rho(t, \cdot)$  to the original WGF. We consider the error analysis in two scenarios: the first variation  $\frac{\delta \mathcal{F}}{\delta \rho}$  is Lipschitz continuous and the case of Fokker-Planck equation where  $\frac{\delta \mathcal{F}}{\delta \rho}$  is not Lipschitz continuous.

#### 3.4.1. Error analysis with Lipschitz continuity assumption

In this subsection, we assume the following condition on  $\mathcal{F}$  holds.

**Assumption 1.** There exists a constant  $C_{\mathcal{F}} > 0$  such that for any two push-forward maps  $T$  and  $\tilde{T}$  there is

$$\int \left| \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\#} \rho(\cdot), \cdot) \circ T(z) - \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(\tilde{T}_{\#} \rho(\cdot), \cdot) \circ \tilde{T}(z) \right|^2 \, d\lambda(z) \leq C_{\mathcal{F}} \int |T(z) - \tilde{T}(z)|^2 \, d\lambda(z). \quad (3.26)$$

The accuracy of our approach depends on the representation power of the parameterized push-forward map  $T_\theta$ . We use  $\delta_0$  to characterize the representation error as below.

**Definition 3.6.** Define the projection error as

$$\begin{aligned} \delta_0 &= \sup_{\theta \in \Theta} \min_{\xi \in T_\theta^\dagger \Theta} \left\{ \int \left| \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_\theta(z) - \partial_\theta T_\theta(z) \xi \right|^2 \, d\lambda(z) \right\} \\ &= \sup_{\theta \in \Theta} \left\{ \int \left| \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_\theta(z) - \mathcal{K}_\theta \left[ \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_\theta \right] (z) \right|^2 \, d\lambda(z) \right\}. \end{aligned} \quad (3.27)$$

This error is essentially the difference between  $\frac{\delta \mathcal{F}(\rho_\theta)}{\delta \rho}$  and its orthogonal projection onto the tangent space  $T_{\rho_\theta} \mathcal{P}(\mathcal{M})$ . In the definition, we use the fact that the operator  $\mathcal{K}_\theta$  is the orthogonal projection. Under Assumption 1, we can show that the density  $\rho_\theta$  approximates the true solution with guaranteed Wasserstein-2 error stated in the next theorem.

**Theorem 3.7.** Suppose Assumption 1 holds for  $\mathcal{F}$ ,  $\theta$  is solved from (3.18) with initial value  $\theta(0)$ , and  $\rho$  is solved from (1.3). Then the Wasserstein-2 distance between the push-forward density  $\rho_{\theta(t)(\cdot)}$  obtained by PWGF and the true density  $\rho(t, \cdot)$  satisfies

$$W_2^2(\rho_{\theta(t)(\cdot)}, \rho(t, \cdot)) \leq e^{(1+2C_{\mathcal{F}})t} \epsilon_{\rho(0)} + \frac{2\delta_0}{1+2C_{\mathcal{F}}} (e^{(1+2C_{\mathcal{F}})t} - 1), \tag{3.28}$$

where  $\epsilon_{\rho(0)} = W_2^2(\rho_{\theta(0)}, \rho_0)$  is the initial approximation error.

**Proof.** Assume  $\mathbf{X}$  is solved from (3.13) and  $\mathbf{Y}$  is solved from (3.22). Suppose the Monge map from  $\rho_{\theta(0)}$  to  $\rho_0$  is given by  $\omega$  and we assume the random variables  $\mathbf{X}, \mathbf{Y}$  are coupled via  $\mathbf{X}(0) = \omega(\mathbf{Y}(0))$ . We continue from the calculation in (3.25) We continue from the calculation in (3.25),

$$\begin{aligned} \dot{\mathbf{Y}} &= -\mathcal{K}_{\theta} \left[ \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_{\theta} \right] (\mathbf{z}) \\ &= -\nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_{\theta} (\mathbf{z}) + \left( \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_{\theta} (\mathbf{z}) - \mathcal{K}_{\theta} \left[ \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_{\theta} \right] (\mathbf{z}) \right). \end{aligned}$$

Denote  $E(t) = \mathbb{E} \|\mathbf{X} - \mathbf{Y}\|^2$ , we compute,

$$\frac{d}{dt} E(t) = 2\mathbb{E} [(\mathbf{X} - \mathbf{Y}) \cdot (\dot{\mathbf{X}} - \dot{\mathbf{Y}})] \leq 2\sqrt{\mathbb{E} \|\mathbf{X} - \mathbf{Y}\|^2} \sqrt{\mathbb{E} \|\dot{\mathbf{X}} - \dot{\mathbf{Y}}\|^2} \leq \mathbb{E} \|\mathbf{X} - \mathbf{Y}\|^2 + \mathbb{E} \|\dot{\mathbf{X}} - \dot{\mathbf{Y}}\|^2. \tag{3.29}$$

Notice that  $\mathbf{X}$  is a push-forward of  $\mathbf{X}(0)$  through the dynamics (3.13), we have

$$\begin{aligned} \mathbb{E} \|\dot{\mathbf{X}} - \dot{\mathbf{Y}}\|^2 &= \int \left| - \left( \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(\rho(x), x) - \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_{\theta} (\mathbf{z}) \right) \right. \\ &\quad \left. - \left( \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_{\theta} (\mathbf{z}) - \mathcal{K}_{\theta} \left[ \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_{\theta} \right] (\mathbf{z}) \right) \right|^2 d\lambda(\mathbf{z}) \\ &\leq 2\delta_0 + 2C_{\mathcal{F}} \mathbb{E} \|\mathbf{X} - \mathbf{Y}\|^2. \end{aligned} \tag{3.30}$$

Plugging (3.30) into (3.29), we obtain

$$\frac{d}{dt} E(t) \leq 2\delta_0 + (1 + 2C_{\mathcal{F}})E(t). \tag{3.31}$$

By Grönwall's inequality, we have

$$E(t) \leq e^{(1+2C_{\mathcal{F}})t} E(0) + \frac{2\delta_0}{1+2C_{\mathcal{F}}} (e^{(1+2C_{\mathcal{F}})t} - 1). \tag{3.32}$$

On the other hand, we have

$$W_2^2(T_{\theta\#} \rho, \rho) = W_2^2(\text{Law}(\mathbf{Y}), \text{Law}(\mathbf{X})) \leq \mathbb{E} \|\mathbf{Y} - \mathbf{X}\|^2 = E(t). \tag{3.33}$$

Combining (3.32) and (3.33) yields (3.28).  $\square$

**Remark 3.8.** In [29], global error estimates in Wasserstein metric were obtained for Fokker-Planck equation. It's different from Theorem 3.7 which applies to a broader class of PDEs.

### 3.4.2. Asymptotic analysis for parametrized Wasserstein gradient flow

For many examples, finding  $C_{\mathcal{F}}$  could be a challenging task. However, if there exists a Gibbs solution  $\rho_*$  to the WGF satisfying Polyak-Łojasiewicz inequality, then we can show uniform convergence of  $\rho_{\theta}$  to the target density  $\rho_*$ , as stated in the following theorem:

**Theorem 3.9.** Assume that  $\rho_*$  is the Gibbs distribution, and the functional  $\mathcal{F}(\rho)$  satisfies certain Polyak-Łojasiewicz inequality as follows

$$\frac{1}{\zeta} \int_{\mathbb{R}^d} \left| \nabla_X \frac{\delta \mathcal{F}}{\delta \rho}(\rho, x) \right|^2 \rho(x) dx \geq \mathcal{F}(\rho) - \mathcal{F}(\rho_*), \tag{3.34}$$

where  $\zeta$  is a positive constant. Then

$$\mathcal{F}(\rho_{\theta(t)}) - \mathcal{F}(\rho_*) \leq \frac{\delta_0}{\zeta} (1 - e^{-\zeta t}) + \mathcal{F}(\rho_{\theta(0)}) e^{-\zeta t}, \tag{3.35}$$

where  $\delta_0$  is defined in (3.27).

**Proof.** Denote  $f(\mathbf{z}) = \nabla_X \frac{\delta}{\delta \rho} \mathcal{F}(T_{\theta\#} \rho(\cdot), \cdot) \circ T_{\theta} (\mathbf{z})$ . Then by the property of orthogonal projection operator  $\mathcal{K}_{\theta}$  we have  $\int (I - \mathcal{K}_{\theta}) [f](\mathbf{z})^{\top} \mathcal{K}_{\theta} [f](\mathbf{z}) d\lambda(\mathbf{z}) = 0$ . Furthermore,



$$\begin{aligned}
 \int f(z)^\top \mathcal{K}_\theta[f](z) d\lambda(z) &= \int \mathcal{K}_\theta[f](z)^\top \mathcal{K}_\theta[f](z) d\lambda(z) + \int (I - \mathcal{K}_\theta)[f](z)^\top \mathcal{K}_\theta[f](z) d\lambda(z) \\
 &= \int \mathcal{K}_\theta[f](z)^\top \mathcal{K}_\theta[f](z) d\lambda(z) \\
 &= \int |f(z)|^2 d\lambda(z) - \int |(I - \mathcal{K}_\theta)[f](z)|^2 d\lambda(z),
 \end{aligned} \tag{3.36}$$

where the last equality is due to the Pythagoras theorem. Following the definition of  $\delta_0$  in (3.27), we have

$$\int |(I - \mathcal{K}_\theta)[f](z)|^2 d\lambda(z) \leq \delta_0. \tag{3.37}$$

By inequality (3.34), we have

$$\int |f(z)|^2 d\lambda(z) \geq \zeta F(\rho). \tag{3.38}$$

We deduce that

$$\begin{aligned}
 \frac{d}{dt} F(\rho_\theta) &= \nabla_\theta F(\theta) \cdot \dot{\theta} \\
 &= -\nabla_\theta F(\theta)^\top \widehat{G}(\theta)^\dagger \nabla_\theta F(\theta) \\
 &= -\int \nabla_X \frac{\delta}{\delta \rho} F(T_{\theta^\#} \lambda(\cdot), \cdot) \circ T_\theta(z)^\top \partial_\theta T_\theta(z) \widehat{G}(\theta)^\dagger \nabla_\theta F(\theta) d\lambda(z) \\
 &= -\int \nabla_X \frac{\delta}{\delta \rho} F(T_{\theta^\#} \lambda(\cdot), \cdot) \circ T_\theta(z)^\top \mathcal{K}_\theta[\nabla_X \frac{\delta}{\delta \rho} F(T_{\theta^\#} \lambda(\cdot), \cdot) \circ T_\theta](z) d\lambda(z) \\
 &= -\int f(z)^\top \mathcal{K}_\theta[f](z) d\lambda(z),
 \end{aligned} \tag{3.39}$$

where the first equality is due to  $F(\theta) := F(\rho_\theta)$ , the second equality is obtained by the formulation of PWGF in (3.18), the third equality uses the substitution of first  $\nabla_\theta F(\theta)$  using (3.21), the fourth equality is because of the definition of  $\mathcal{K}_\theta$  in (3.23), and the fifth equality is by the definition of  $f$  at the beginning of this proof.

We also observe that

$$-\int f(z)^\top \mathcal{K}_\theta[f](z) d\lambda(z) \leq \delta_0 - \int |f(z)|^2 d\lambda(z) \leq \delta_0 - \zeta F(\rho_\theta), \tag{3.40}$$

where the first inequality is obtained by combining (3.36) and (3.37), the second equality is due to (3.38). Combining (3.39) and (3.40) yields

$$\frac{d}{dt} F(\theta) \leq \delta_0 - \zeta F(\theta). \tag{3.41}$$

Applying the Grönwall inequality to (3.41) yields the claimed estimate (3.35).  $\square$

We would like to mention that the bound obtained in Theorem 3.9 is similar to the results reported in [29] except that they use different pullback metrics, namely we use  $\widehat{G}(\theta)$  while  $G(\theta)$  is considered in [29].

#### 4. Numerical method

The PWGF (3.18) can be solved based on standard numerical integrators. We showcase how to implement this using the forward Euler scheme, whereas the idea can be easily generated to other methods such as Runge-Kutta 4th-order method, predictor-corrector method, and those with variable step sizes. In forward Euler scheme, we can discretize the time  $t$  as  $\{hl : l = 0, 1, \dots\}$ . Let  $\theta^l$  be the approximation of  $\theta(hl)$  of PWGF (3.18), then we can compute  $\theta^{l+1}$  given  $\theta^l$  by solving

$$\frac{\theta^{l+1} - \theta^l}{h} = -\widehat{G}(\theta^l)^\dagger \nabla_\theta F(\theta^l), \tag{4.1}$$

for  $l = 0, 1, \dots$ . The key to solving (4.1) is finding the minimum norm solution to the linear system  $\widehat{G}(\theta^l)\eta = \nabla_\theta F(\theta^l)$  for  $\eta$ . Minimum norm solutions of a linear system can be done by numerous existing methods, such as MINRES [35]. In these methods, it is instrumental to implement the matrix-vector product with the matrix  $\widehat{G}(\theta^l)$ .

In our experiments, we use deep neural networks as the push-forward map  $T_\theta$ . There are several choices of the neural network architecture. We can use the invertible neural networks (e.g., normalizing flow [34], Real NVP [12] and neural ODE [10]) or non-invertible neural networks (e.g., the multi-layer perceptron or ResNet [17]), both has its own advantages. Normalizing flow and continuous normalizing flow simplify the computation of log determinant of Jacobian matrix of the map, so we can easily compute the density function. For the experiment of the Fokker-Planck equation, we use a 60-layer normalizing flow as the push-forward map  $T_\theta$ . For the experiments of the Porous-medium equation and the aggregation model, we pick the residual neural network as push-forward map:



**Algorithm 4.1** Parameterized Wasserstein gradient flow solver.

---

```

Initialize the neural network  $T_\theta$ 
for  $l = 0, \dots, K - 1$  do
  Sample  $\{X_1, \dots, X_N\}$ , evaluate  $\nabla_\theta F(\theta^l)$ 
  Solve  $\eta$  as the minimum norm solution to the linear system  $\widehat{G}(\theta^l)\eta = -\nabla_\theta F(\theta^l)$ 
  Set  $\theta^{l+1} = \theta^l + h\eta$ 
end for
return  $\theta^K, T_{\theta^K}$ 

```

---

$$T_\theta = Id + R_\theta$$

where  $Id$  is the identity map and  $R_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is a standard multilayer perceptron. We take the hyperbolic tangent function as the activation function since we require the second-order derivative in the computation of numerical integrals. The bias for the output layer in  $f_\theta$  is set to be None. To solve the linear system  $\widehat{G}(\theta)\theta = -\nabla_\theta F(\theta)$ , we utilized MINRES with tolerance  $3 \cdot 10^{-4}$  in our experiments.

To initialize  $\theta^0 := \theta(0)$ , it depends on the given information of the initial distribution. If the initial probability density function is given and easy to generate samples from, such as Gaussian distribution or Gaussian mixture distribution, then one can initialize the  $\theta(0)$  such that  $T_{\theta(0)}$  is an identity map. If the initial samples are given but the initial probability density function is unknown, one can still initialize the  $\theta(0)$  to make  $T_{\theta(0)}$  an identity map. If the given initial density function is hard to sample from, one may use Gaussian distribution as a reference distribution and utilize methods such as MCMC [4] and Wasserstein GAN [2] to enforce the initial distribution.

If the probability density function is desired, one can consult with density estimation methods, such as kernel density estimation, while it's not the focus of this paper. In our experiments, we're given the initial density functions that are easy to sample, hence we initialize the push-forward map as an identity map, i.e.,  $\rho_{\theta(0)} dx = d(T_{\theta(0)\#}\lambda)(\cdot) = d\lambda(\cdot) = \rho(0, \cdot) dx$ .

The sample size and number of layers in normalizing flow are determined empirically. It's a critical issue to develop systematic criteria for the performance depending on the sample size and number of layers. These issues to be addressed are different from Theorem 3.7 and Theorem 3.9 as they're on density level.

The numerical scheme is summarized in Algorithm 4.1.

### 5. Numerical examples

In this section, we test the proposed PWGF method, Algorithm 4.1, on three WGFs. Depending on the energy functional  $\mathcal{F}$ , we choose push-forward maps as ResNet [17], or normalizing flow [21]. Details are shown below.

#### 5.1. Fokker-Planck equation

We first test Algorithm 4.1 on a 30-dimensional Fokker-Planck equation. We set the coefficient  $D = 1$  and choose the potential function  $V$  to be the Styblinski-Tang function in (3.4):

$$V(x) = \frac{3}{50} \left( \sum_{i=1}^d x_i^4 - 16x_i^2 + 5x_i \right). \tag{5.1}$$

The initial condition is set to be a Gaussian distribution with mean  $\mu_0 = \mathbf{0}$  and variance  $\sigma_0 = \mathbf{1}$ . We use an  $M$ -layer normalizing flow ( $M = 40$ ) as the push-forward map  $T_\theta$ :

$$T_\theta = f_M \circ f_{M-1} \circ \dots \circ f_2 \circ f_1, \tag{5.2}$$

where  $f_j$  ( $1 \leq j \leq M$ ) is:

$$f_j(x) = x + \tanh(w_j^\top x + b_j)u_j$$

Here  $w_j, u_j \in \mathbb{R}^d$  and  $b_j \in \mathbb{R}$ , and  $\theta = \{w_j, b_j, u_j : 1 \leq j \leq M\}$ .

We run Algorithm 4.1 with time step size  $h = 0.005$ . 2,000 samples are generated from  $\rho_\theta$  and their projection of these samples onto the 2-dimensional subspace of  $(x_1, x_2)$  are plotted. Notice that  $V(x)$  has 4 centers, and the push-forward samples follow a density that captures this structure as shown in Fig. 1. The KL divergence (3.3) is evaluated empirically with all samples by

$$\begin{aligned} \mathcal{H}_{\text{KL}}(\rho \parallel \rho_*) &= \int \frac{1}{D} V(x)\rho(x) + \rho(x) \log \rho(x) dx + \log Z_D \\ &\approx \frac{1}{N} \sum_{i=1}^N \frac{1}{D} V(x^{(i)}) + \log \rho(x^{(i)}) + \log Z_D, \quad \text{where } \{x^{(i)}\}_{i=1}^N \sim \rho. \end{aligned} \tag{5.3}$$

We plot the empirically evaluated shifted KL divergence  $\mathcal{H}_{\text{KL}}(\rho \parallel \rho_*) - \log Z_D$  following the solution of PWGF in Fig. 2, which shows a clear decreasing trend.




Fig. 1. Sample plots of computed  $\rho_\theta$  at different time  $t$  for Fokker-Planck equation with the Styblinski-Tang function as the potential function  $V(x)$ .




Fig. 2. The decay of KL divergence along the solution of PWGF for the Fokker-Planck equation with Styblinski-Tang potential in the 30-dimensional example. The x-axis is the number of iterations and the y-axis is the value of the empirical evaluation of the shifted KL-divergence  $\mathcal{H}_{\text{KL}}(\rho||\rho_*) - \log Z_D$ .

To evaluate the improvement in its computation efficiency with the new metric  $\hat{G}(\theta)$ , we compare the running time of PWGF Algorithm 4.1 with the PFPE algorithm proposed in [29]. We take the 2-dimensional FPE with  $V(x)$  defined in (5.1). For both algorithms, we use the same neural network structure (40-layer normalizing flow) and the same sample size (12000) to evaluate the KL divergence. The  $\psi$  neural network in PFPE, as well as the training parameters, are set as its default values as reported in [29]. We run both algorithms with the same step size. It takes 14.26 seconds for PWGF to finish 10 time steps, while 6093.58 seconds for PFPE to finish 10 time steps, indicating a more than 400 times speed up. Both codes are run on an NVIDIA RTX-3070 GPU with CUDA enabled.

### 5.2. Porous medium equation

We apply Algorithm 4.1 to solve the porous medium equation (3.9). In recent years, the porous medium equation has been solved by several other methods, such as physics-informed neural networks (PINN) [37] which computes the density function directly in high dimensions and neural network-based implicit particle methods with JKO scheme [24] which computes samples of the solution. For all three examples of porous-medium equation, we use neural ODE (NODE) as the push-forward map

$$T_\theta = \text{NODE}(R_\theta) \tag{5.4}$$

where NODE indicates the neural ODE structure,  $R_\theta$  is a multi-layer perceptron with 3 hidden layers and 100 neurons in each hidden layer. We use tanh as the activation function in all layers.

An example with exact weak solution: Suppose the initial condition  $\rho(x, 0) = \delta(x)$  is the Dirac delta function.

$$\frac{\partial \rho(x, t)}{\partial t} = \Delta \rho^m, \quad x \in \mathbb{R}^d \tag{5.5}$$

$$\rho(x, 0) = \delta(x) \tag{5.6}$$

The exact weak solution to (5.5) is

$$\rho(x, t) = t^{-\alpha} F\left(\frac{x}{t^\beta}\right), \quad \text{where } F(\xi) = (C - k\xi^2)_+^{\frac{1}{m-1}}, \tag{5.7}$$

and  $(s)_+ = \max\{s, 0\}$ . The constants are

$$\alpha = \frac{d}{d(m-1)+2}, \quad \beta = \frac{\alpha}{d}, \quad k = \frac{(m-1)\alpha}{2md}$$

The constant  $C > 0$  in (5.7) can be uniquely determined by normalizing the mass as

$$\int \rho dx = 1,$$

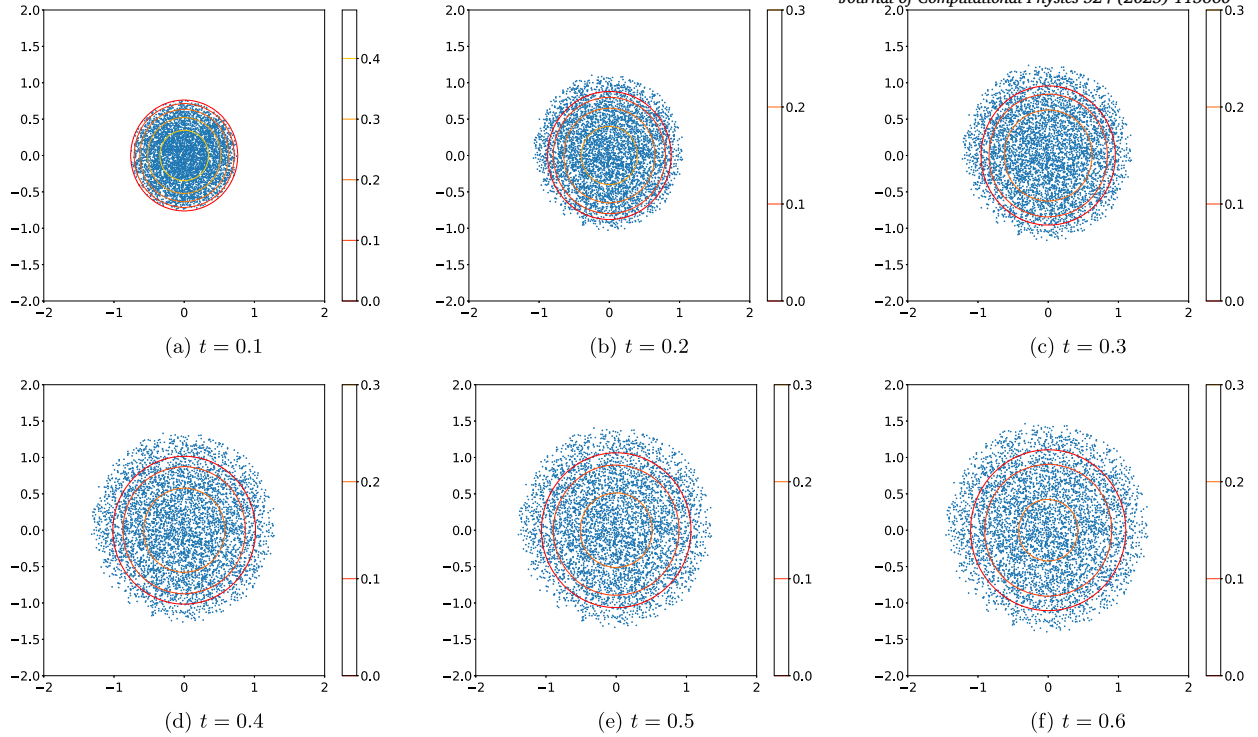


Fig. 3. Sample plots of computed  $\rho_\theta$  at different time  $t$  for Porous Medium equation with Dirac Delta function as the initial condition for  $d = 2$  and  $l = 2.4$ . The figures are plotted with 5000 samples. In the level curves, darker colors correspond to smaller values to emphasize the support.

which gives

$$C = \left[ \frac{\left(\frac{k}{\pi}\right)^{\frac{d}{2}} \Gamma\left(\frac{d}{2}\right)}{d \cdot B\left(\frac{d}{2}, \frac{m}{m-1}\right)} \right]^{\frac{1}{\gamma}},$$

where  $B(\cdot, \cdot)$  and  $\Gamma(\cdot)$  are beta and gamma functions respectively and  $\gamma = \frac{d}{2(m-1)\alpha}$  [36].

For the experiments of  $d = 2$ , the parameter  $m$  is set to 2.4, for the experiment of  $d = 5$ , the parameter  $m$  is set to 3, and for the experiment of  $d = 15$ , the parameter  $m$  is set to 2. Since the initial condition is a Dirac delta mass which is impossible to be represented by deep neural networks, we set the initial time to a positive value. More precisely, we compute the dynamics of the samples generated from the probability density  $\rho$  from  $t = 0.1$  to  $t = 0.6$  for  $d = 2$  and from  $t = 1$  to  $t = 1.5$  for both  $d = 5$  and  $d = 15$ . In all experiments, the step size  $h = 10^{-4}$ , and the number of time steps is  $T = 5000$ . The sample size to evaluate  $\hat{G}(\theta)$  and  $F(\theta)$  are both 30,000. The graph of the samples with  $d = 2$ ,  $d = 5$  and  $d = 15$  are shown in Fig. 3, Fig. 4 and Fig. 5 respectively. In the plots, the blue points are the computed samples and the colored curves are the level sets of the exact solutions. We select darker colors representing smaller values when plotting the level curves to clearly show the boundary of the finite support of the exact solution. In the example with  $d = 2$ , though the computed samples run a bit more quickly than the exact solution, the numerical solution still behaves the same as the exact solution in the sense that they are compactly supported. In the example with  $d = 15$ , all samples are still within the compact support, despite that sampling from a density function in high dimensions remains a challenging problem, especially since samples near the boundary of the compact support are hard to obtain. All experiments for Porous-Medium equation run on an NVIDIA A100 GPU with 40GB GPU memory.

*An Example with Gaussian Mixture Distribution as the Initial Condition:* In this example, we apply Algorithm 4.1 to the porous medium equation with  $d = 10$ ,  $m = 3$  and a Gaussian mixture distribution as the initial condition. The equations are

$$\frac{\partial \rho}{\partial t} = \Delta(\rho^m) \tag{5.8}$$

$$\rho(0, \cdot) = 0.2 \cdot \mathcal{N}(\mu_1, \sigma_1) + 0.8 \cdot \mathcal{N}(\mu_2, \sigma_2) \tag{5.9}$$

where  $\mu_1 = (0, 0, \dots, 0)^T \in \mathbb{R}^{10}$ ,  $\mu_2 = (2, 2, \dots, 2)^T \in \mathbb{R}^{10}$ ,  $\sigma_1 = \text{diag}(0.1, \dots, 0.1)^T$  and  $\sigma_2 = \text{diag}(0.2, \dots, 0.2)^T$ . The sample size to evaluate  $\hat{G}(\theta)$  and  $F(\theta)$  are both 15,000. The computed samples are shown in Fig. 6 plotted with 5000 samples. The exact solution to this initial value problem remains open but our method can provide a simulated sample dynamics.

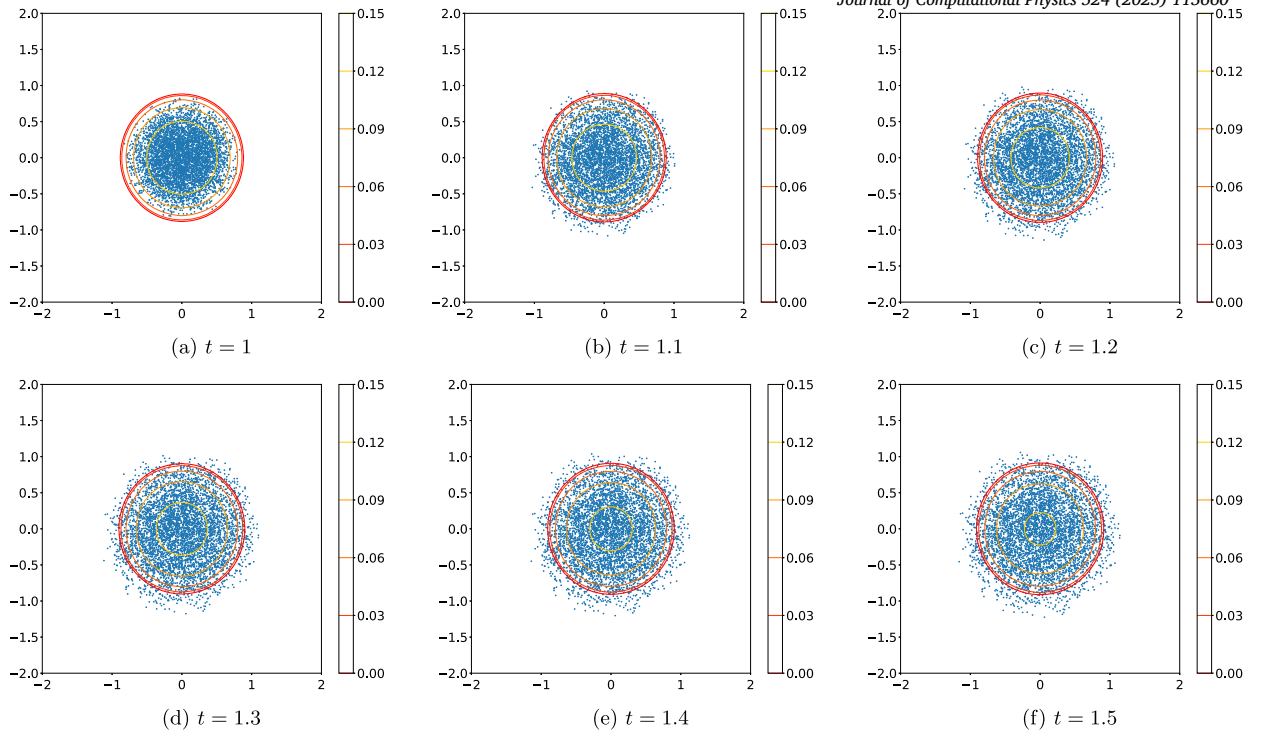


Fig. 4. Sample plots of computed  $\rho_\theta$  at different time  $t$  for the porous medium equation with Dirac Delta function as the initial condition for  $d = 5$  and  $l = 3$ . The figures are plotted with 5000 samples. In the level curves, darker colors correspond to smaller values to emphasize the support.

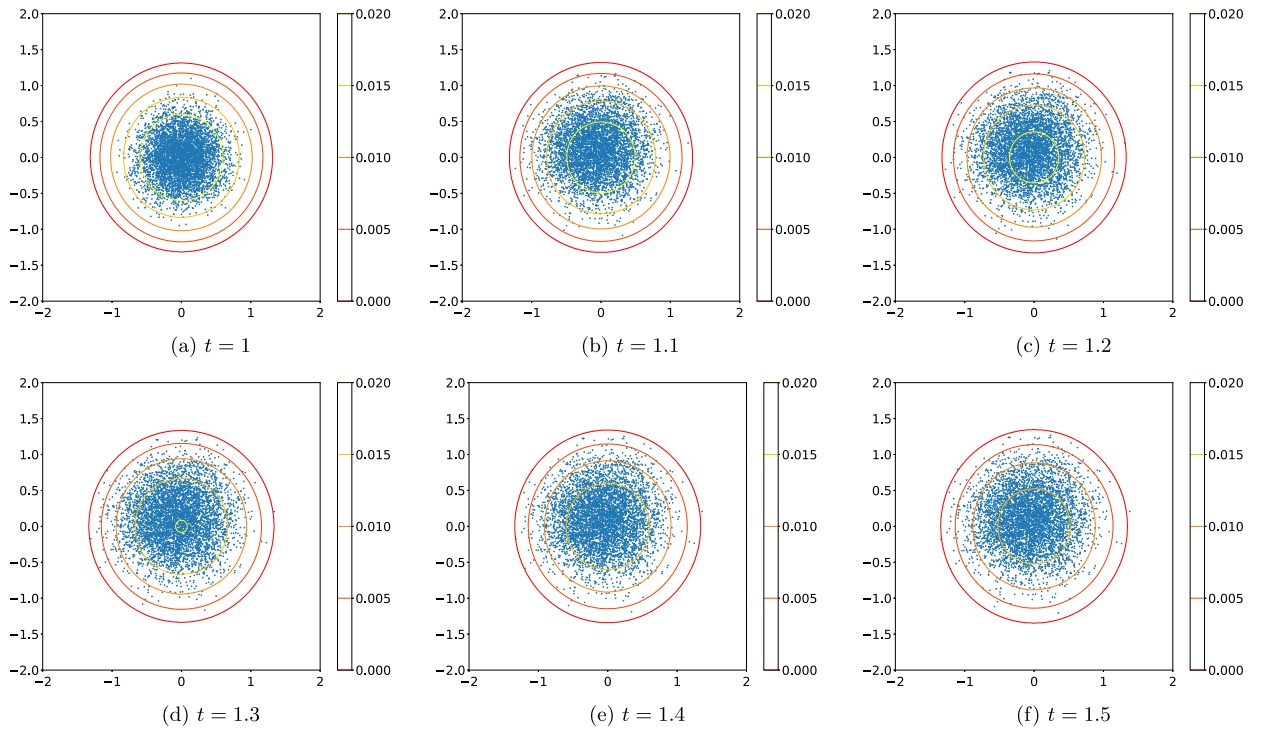


Fig. 5. Sample plots of computed  $\rho_\theta$  at different time  $t$  for porous medium equation with Dirac Delta function as the initial condition for  $d = 15$ . The figures are plotted with 5000 samples. In the level curves, darker colors correspond to smaller values to emphasize the support. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)




Fig. 6. Sample plots of computed  $\rho_\theta$  at different time  $t$  for porous medium equation with Gaussian mixture distribution as the initial condition. The figures are plotted with 5000 samples.




Fig. 7. Sample plots of computed  $\rho_\theta$  at different time  $t$  for porous medium equation with the mixture of Gaussian distribution and uniform distribution as the initial condition. The graphs are plotted with 5000 samples.

*An Example with Mixed Gaussian Distribution and Uniform Distribution as the Initial Condition:* In this example, we apply Algorithm 4.1 to the porous medium equation with  $d = 5$ ,  $m = 3$  and a mixture of Gaussian and uniform distributions as the initial condition. The equations are

$$\frac{\partial \rho}{\partial t} = \Delta(\rho^m) \tag{5.10}$$

$$\rho(0, \cdot) = 0.2 \cdot \mathcal{N}(\mu, \sigma) + 0.8 \cdot \mathcal{U}_5([-1, 1]^5) \tag{5.11}$$

where  $\mu = (2, 2, \dots, 2)^T \in \mathbb{R}^5$ ,  $\sigma = (0.1, 0.1, \dots, 0.1)^T \in \mathbb{R}^5$  and  $\mathcal{U}_5([-1, 1]^5)$  is a uniform distribution in  $[-1, 1]^5$ . Note that this equation doesn't have a closed-form solution and it's not tractable in high dimensions. The sample size to evaluate  $\hat{G}(\theta)$  and  $F(\theta)$  are both 20,000. The computed samples are shown in Fig. 7 plotted with 5000 samples.

### 5.3. Aggregation model as a Wasserstein gradient flow

In this example, we apply Algorithm 4.1 to the aggregation problem with  $d = 2$  and  $a = 4$ ,  $b = 2$ , i.e., the energy functional  $\mathcal{F}$  is set as in (3.10) with  $J(x)$  given by

$$J(x) = \frac{|x|^4}{4} - \frac{|x|^2}{2}. \tag{5.12}$$

We set the initial condition to be Gaussian distribution with mean  $\mu_0 = (1.25, 1.25)^T$  and variance  $\gamma I$  where  $\gamma = 0.6$ .

$$\rho_0(x) = \frac{1}{\sqrt{2\pi\gamma}} e^{-\frac{|x-\mu_0|^2}{2\gamma^2}} \tag{5.13}$$

We use a neural network with residual structure as the push-forward map in this example:

$$T_\theta = Id + R_\theta \tag{5.14}$$




Fig. 8. Sample plots of computed  $\rho_\theta$  at different time  $t$  for aggregation problem.

where  $R_\theta$  is multi-layer perceptron with 2 hidden layers and 50 neurons in each hidden layer. The sample size to evaluate  $\hat{G}(\theta)$  and  $F(\theta)$  are both 10,000.

The steady solution  $\rho_*$  is a Dirac distribution uniformly concentrated on the ring with radius 0.5 centered at  $\mu_0$ .  $\rho$  converges to  $\rho_*$  when  $t \rightarrow \infty$ . This evolution is captured as shown in Fig. 8, which shows the distribution  $\rho_{\theta(t)}$  of samples from at different times  $t$ . Note that, despite of the strong singularity of  $\rho_*$ , the samples generated by Algorithm 4.1 can approximate this distribution well. The code is run on an NVIDIA RTX-3070 GPU with CUDA enabled.

## 6. Discussion

We proposed a new numerical approach to solve Wasserstein gradient flows (WGFs), which is particularly scalable for high-dimensional cases. Our approach used general reduced-order models, like deep neural networks, to parameterize the push-forward maps such that they can push a simple reference density to the density solving the given WGF. Essentially, the gradient flow defined on the infinite-dimensional Wasserstein manifold is reduced to a finite-dimensional dynamical system, the PWGF, for the parameters used in the reduced-order model. At the core of our design is the pullback Wasserstein metric on the parameter space. It facilitates the derivation, simplifies the numerical computation, and plays a pivotal role in the error-bound estimates that offer theoretical assurance for PWGF to the original WGF.

Though the parametrization of the WGF provides a fast algorithm, the accuracy and speed of the algorithm largely rely on the linear system solver for the neural network parameter  $\theta$ . Since the metric matrix  $\hat{G}(\theta)$  is semi-positive definite and can be degenerate in some scenarios, the minimal residual method is utilized to find the least square solution.

Our method is sample-based and provides the value of the density function at sampled points, it has the potential to adapt to tasks such as developing a fast sampler of target distribution by evolving samples from the initial probability density following the gradient flow or image generation following certain dynamics of probability densities. In these scenarios, as the task is more complex, the deep neural network for parametrizing the push-forward map could be large or have sophisticated architecture. One may consider dimension reduction techniques to reduce the dimension of the parameter space. In addition, other improvements such as more efficient linear system solvers or preconditioning techniques may also be considered to boost the accuracy and speed. Meanwhile, better understanding on  $\delta_0$  in the error analysis needs to be analyzed regarding various neural network structures.

## CRedit authorship contribution statement

**Yijie Jin:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Shu Liu:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Hao Wu:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Xiaojing Ye:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Haomin Zhou:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

This research is partially supported by NSF grants DMS-1925263, DMS-2152960, DMS-2307465, DMS-2307466, and ONR grant N00014-21-1-2891. All the authors made equal contributions.

## Data availability

Data will be made available on request.



## References

- [1] W. Anderson, M. Farazmand, Evolution of nonlinear reduced-order solutions for PDEs with conserved quantities, *SIAM J. Sci. Comput.* 44 (2022) A176–A197.
- [2] M. Arjovsky, S. Chintala, L. Bottou, Wasserstein generative adversarial networks, in: *International Conference on Machine Learning*, PMLR, 2017, pp. 214–223.
- [3] J.-D. Benamou, G. Carlier, Q. Mérigot, E. Oudet, Discretization of functionals involving the Monge–Ampère operator, *Numer. Math.* 134 (2016) 611–636.
- [4] B.A. Berg, *Markov Chain Monte Carlo Simulations and Their Statistical Analysis: with Web-Based Fortran Code*, World Scientific Publishing Company, 2004.
- [5] A. Blanchet, P. Laurençot, The parabolic-parabolic Keller–Segel system with critical diffusion as a gradient flow in  $\mathbb{R}^d$ ,  $d \geq 3$ , *Comm. Partial Differential Equations* 38 (2013) 658–686.
- [6] S.C. Brenner, *The Mathematical Theory of Finite Element Methods*, Springer, 2008.
- [7] J. Bruna, B. Peherstorfer, E. Vanden-Eijnden, Neural Galerkin schemes with active learning for high-dimensional evolution equations, *J. Comput. Phys.* 496 (2024) 112588.
- [8] G. Carlier, V. Duval, G. Peyré, B. Schmitzer, Convergence of entropic schemes for optimal transport and gradient flows, *SIAM J. Math. Anal.* 49 (2017) 1385–1418.
- [9] J.A. Carrillo, K. Craig, L. Wang, C. Wei, Primal dual methods for Wasserstein gradient flows, *Found. Comput. Math.* (2022) 1–55.
- [10] R.T. Chen, Y. Rubanova, J. Bettencourt, D.K. Duvenaud, Neural ordinary differential equations, *Adv. Neural Inf. Process. Syst.* 31 (2018).
- [11] S. Chewi, T. Le Gouic, C. Lu, T. Maunu, P. Rigollet, Svdg as a kernelized Wasserstein gradient flow of the chi-squared divergence, *Adv. Neural Inf. Process. Syst.* 33 (2020) 2098–2109.
- [12] L. Dinh, J. Sohl-Dickstein, S. Bengio, Density estimation using real nvp, *arXiv preprint*, arXiv:1605.08803, 2016.
- [13] Y. Du, T.A. Zaki, Evolutional deep neural network, *Phys. Rev. E* 104 (2021) 045303.
- [14] J. Fan, Q. Zhang, A. Taghvaei, Y. Chen, Variational Wasserstein gradient flow, in: *International Conference on Machine Learning*, PMLR, 2022, pp. 6185–6215.
- [15] R.C. Fetecau, Y. Huang, T. Kolokolnikov, Swarm dynamics and equilibria for a nonlocal aggregation model, *Nonlinearity* 24 (2011) 2681.
- [16] N. Gaby, X. Ye, H. Zhou, Neural control of parametric solutions for high-dimensional evolution pdes, *SIAM J. Sci. Comput.* 46 (2024) C155–C185.
- [17] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [18] R. Holley, D.W. Stroock, Logarithmic Sobolev inequalities and stochastic Ising models, *J. Stat. Phys.* 46 (1987) 1159–1194, <https://doi.org/10.1007/BF01011611>.
- [19] Z. Hu, C. Liu, Y. Wang, Z. Xu, Energetic variational neural network discretizations to gradient flows, *arXiv preprint*, arXiv:2206.07303, 2022.
- [20] R. Jordan, D. Kinderlehrer, F. Otto, The variational formulation of the Fokker–Planck equation, *SIAM J. Math. Anal.* 29 (1998) 1–17.
- [21] I. Kobyzev, S.J. Prince, M.A. Brubaker, Normalizing flows: an introduction and review of current methods, *IEEE Trans. Pattern Anal. Mach. Intell.* 43 (2020) 3964–3979.
- [22] A.J. Ladd, P. Szymczak, Reactive flows in porous media: challenges in theoretical and numerical methods, *Annu. Rev. Chem. Biomol. Eng.* 12 (2021) 543–571.
- [23] J.D. Lafferty, The density manifold and configuration space quantization, *Trans. Amer. Math. Soc.* 305 (1988) 699–741.
- [24] W. Lee, L. Wang, W. Li, Deep JKO: time-implicit particle methods for general nonlinear gradient flows, *arXiv preprint*, arXiv:2311.06700, 2023.
- [25] R.J. LeVeque, *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems*, SIAM, 2007.
- [26] W. Li, S. Liu, H. Zha, H. Zhou, *Parametric Fokker–Planck Equation*, Springer International Publishing, 2019, pp. 715–724.
- [27] W. Li, J. Lu, L. Wang, Fisher information regularization schemes for Wasserstein gradient flows, *J. Comput. Phys.* 416 (2020) 109449.
- [28] W. Li, J. Zhao, Wasserstein information matrix, *Inf. Geom.* (2023) 1–53.
- [29] S. Liu, W. Li, H. Zha, H. Zhou, Neural parametric Fokker–Planck equation, *SIAM J. Numer. Anal.* 60 (2022) 1385–1449.
- [30] P. Mokrov, A. Korotin, L. Li, A. Genevay, J.M. Solomon, E. Burnaev, Large-scale Wasserstein gradient flows, *Adv. Neural Inf. Process. Syst.* 34 (2021) 15243–15256.
- [31] L. Nurbekyan, W. Lei, Y. Yang, Efficient natural gradient descent methods for large-scale pde-based optimization problems, *SIAM J. Sci. Comput.* 45 (2023) A1621–A1655.
- [32] F. Otto, The geometry of dissipative evolution equations: the porous medium equation, *Comm. Partial Differential Equations* 26 (2001) 101–174.
- [33] G. Peyré, Entropic approximation of Wasserstein gradient flows, *SIAM J. Imaging Sci.* 8 (2015) 2323–2351.
- [34] D. Rezende, S. Mohamed, Variational inference with normalizing flows, in: *International Conference on Machine Learning*, PMLR, 2015, pp. 1530–1538.
- [35] Y. Saad, *Iterative Methods for Sparse Linear Systems*, SIAM, 2003.
- [36] J.L. Vázquez, *The Porous Medium Equation: Mathematical Theory*, Oxford University Press, 2007.
- [37] M. Wang, J. Lu, Neural network-based variational methods for solving quadratic porous medium equations in high dimensions, *Commun. Math. Stat.* 11 (2023) 21–57.
- [38] H. Wu, *Theory and Computation of Wasserstein Geometric Flows with Application to Time-Dependent Schrödinger Equation*, Publisher: Georgia Institute of Technology, 2023.
- [39] H. Wu, S. Liu, X. Ye, H. Zhou, Parameterized Wasserstein Hamiltonian flow, *arXiv preprint*, arXiv:2306.00191, 2023.
- [40] M. Yi, S. Liu, Bridging the gap between variational inference and Wasserstein gradient flows, <https://arxiv.org/abs/2310.20090>, 2023.
- [41] X. Zuo, J. Zhao, S. Liu, S. Osher, W. Li, Numerical analysis on neural network projected schemes for approximating one dimensional Wasserstein gradient flows, <https://arxiv.org/abs/2402.16821>, 2024.