
Accelerated Flow for Probability Distributions

Amirhossein Taghvaei¹ Prashant G. Mehta¹

Abstract

This paper presents a methodology and numerical algorithms for constructing accelerated gradient flows on the space of probability distributions. In particular, we extend the recent variational formulation of accelerated methods in (Wibisono et al., 2016) from vector valued variables to probability distributions. The variational problem is modeled as a mean-field optimal control problem. A quantitative estimate on the asymptotic convergence rate is provided based on a Lyapunov function construction, when the objective functional is displacement convex. An important special case is considered where the objective functional is the relative entropy. For this case, two numerical approximations are presented to implement the Hamilton's equations as a system of N interacting particles. The algorithm is numerically illustrated and compared with the MCMC and Hamiltonian MCMC algorithms.

1. Introduction

Optimization on the space of probability distributions is important to a number of machine learning models including variational inference (Blei et al., 2017), generative models (Goodfellow et al., 2014; Arjovsky et al., 2017), and policy optimization in reinforcement learning (Sutton et al., 2000). A number of recent studies have considered solution approaches to these problems based upon a construction of gradient flow on the space of probability distributions (Liu & Wang, 2016; Richemond & Maginnis, 2017; Zhang et al., 2018; Frogner & Poggio, 2018; Chizat & Bach, 2018; Chen et al., 2018; Liu et al., 2018). Such constructions are useful for convergence analysis as well as development of numerical algorithms.

¹Department of Mechanical Science and Engineering, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL, USA. Correspondence to: Amirhossein Taghvaei <taghvae2@illinois.edu>.

In this paper, we propose a methodology and numerical algorithms that achieve accelerated gradient flows on the space of probability distributions. The proposed numerical algorithms are related to yet distinct from the accelerated stochastic gradient descent (Jain et al., 2017) and Hamiltonian Markov chain Monte-Carlo (MCMC) algorithms (Neal et al., 2011; Cheng et al., 2017). The proposed methodology extends the variational formulation of (Wibisono et al., 2016) from vector valued variables to probability distributions. The original formulation of (Wibisono et al., 2016) was used to derive and analyze the convergence properties of a large class of accelerated optimization algorithms, most significant of which is the continuous-time limit of the Nesterov's algorithm (Su et al., 2014). In this paper, the limit is referred to as the Nesterov's ordinary differential equation.

The extension proposed in our work is based upon a generalization of the formula for the Lagrangian in (Wibisono et al., 2016): (i) the kinetic energy term is replaced with the expected value of kinetic energy; (ii) the potential energy term is replaced with a suitably defined functional on the space of probability distributions. The variational problem is to obtain a trajectory in the space of probability distributions that minimizes the action integral of the Lagrangian.

The variational problem is modeled as a mean-field optimal problem (Bensoussan et al., 2013; Carmona & Delarue, 2017). The maximum principle of the optimal control theory is used to derive the Hamilton's equations which represent the first order optimality conditions. The Hamilton's equations provide a generalization of the Nesterov's ODE to the space of probability distributions. A Lyapunov function is proposed for the convergence analysis of the solution of the Hamilton's equations. In this way, quantitative estimates on convergence rate are obtained for the case when the objective functional is displacement convex (McCann, 1997). Table 1 provides a summary of the relationship between the original variational formulation in (Wibisono et al., 2016) and the extension proposed in this paper.

We also consider the important special case when the objective functional is the relative entropy functional $D(\rho|\rho_\infty)$ defined with respect to a target probability distribution ρ_∞ . In this case, the accelerated gradient flow is shown to be related to the continuous limit of the Hamiltonian Monte-Carlo algorithm (Cheng et al., 2017) (Remark 2). The Hamil-

ton's equations are finite-dimensional for the special case when the initial and the target probability distributions are both Gaussian. In this case, the mean evolves according to the Nesterov's ODE. For the general case, the Lyapunov function-based convergence analysis applies when the target distribution is log-concave.

As a final contribution, the proposed methodology is used to obtain a numerical algorithm. The algorithm is an interacting particle system that empirically approximates the distribution with a finite but large number of N particles. The difficult part of this construction is the approximation of the interaction term between particles. For this purpose, two types of approximations are described: (i) Gaussian approximation which is asymptotically (as $N \rightarrow \infty$) exact in Gaussian settings; and (ii) Diffusion map approximation which is computationally more demanding but asymptotically exact for a general class of distributions.

The outline of the remainder of this paper is as follows: Sec. 2 provides a brief review of the variational formulation in (Wibisono et al., 2016). The proposed extension to the space of probability distribution appears in Sec. 3 where the main result is also described. The numerical algorithm along with the results of numerical experiments appear in Sec. 4. Comparisons with MCMC and Hamiltonian MCMC are also described. The conclusions appear in Sec. 5. The proof of the results appear in the supplementary note.

1.1. Related work

Construction of accelerated flows for probability distribution was proposed in (Liu et al., 2018) based on the generalization of the Nesterov's method to Riemannian manifolds (Liu et al., 2017). The procedure involves approximating the exponential map and parallel transport map for probability distributions in the Wasserstein space. Our construction of accelerated flow is different from (Liu et al., 2018) in several respects: i) we describe a variational formulation and make connection to mean-field control theory; ii) our variational construction yields a continuous-time algorithm providing a straightforward comparison to HMCMC; iii) we carry out convergence analysis based upon a Lyapunov function method; iv) and analysis in Gaussian setting shows we recover the Nesterov ode.

Another class of related work are the interacting particle-based numerical algorithms designed to sample from a target distribution. An example is the Stein variational gradient descent (SVGD) algorithm (Liu & Wang, 2016; Liu, 2017) based on the Riemannian construction of the gradient flow. Another example is the particle optimization method (Chen et al., 2018), whose update is obtained from a solution to an optimization problem based on the variational formulation of the Langevin dynamics. Interacting particle systems have also been shown to be useful for numerically solving the

nonlinear filtering problem (Del Moral et al., 1998; Reich, 2011; Yang et al., 2016; Zhang et al., 2019).

Notation: The gradient and divergence operators are denoted as ∇ and $\nabla \cdot$ respectively. With multiple variables, ∇_z denotes the gradient with respect to the variable z . Therefore, the divergence of the vector field U is $\nabla \cdot U(x) = \sum_{n=1}^d \nabla_{x_n} U_n(x)$. The space of absolutely continuous probability measures on \mathbb{R}^d with finite second moments is denoted by $\mathcal{P}_{ac,2}(\mathbb{R}^d)$. For a measure $\mu \in \mathcal{P}_{ac,2}(\mathbb{R}^d)$ and a measurable map $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$, the push-forward of μ by T is denoted by $T\#\mu$. The second-order Wasserstein distance between any two measures $\mu, \nu \in \mathcal{P}_{ac,2}(\mathbb{R}^d)$ is denoted as $W_2(\mu, \nu)$. The Wasserstein gradient and the Gâteaux derivative of a functional F is denoted as $\nabla_\rho F(\rho)$ and $\frac{\partial F}{\partial \rho}(\rho)$ respectively (see supplementary note for definition). The probability distribution of a random variable Z is denoted as $\text{Law}(Z)$.

2. Review of the variational formulation of (Wibisono et al., 2016)

The basic problem is to minimize a C^1 smooth convex function f on \mathbb{R}^d . The standard form of the gradient descent algorithm for this problem is an ODE:

$$\frac{dX_t}{dt} = -\nabla f(X_t), \quad t \geq 0 \quad (1)$$

Accelerated forms of this algorithm are obtained based on a variational formulation due to (Wibisono et al., 2016). The formulation is briefly reviewed here using an optimal control formalism. The Lagrangian $L : \mathbb{R}^+ \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is defined as

$$L(t, x, u) := e^{\alpha_t + \gamma_t} \left(\underbrace{\frac{1}{2} |e^{-\alpha_t} u|^2}_{\text{kinetic energy}} - \underbrace{e^{\beta_t} f(x)}_{\text{potential energy}} \right) \quad (2)$$

where $t \geq 0$ is the time, $x \in \mathbb{R}^d$ is the state, $u \in \mathbb{R}^d$ is the velocity or control input, and the time-varying parameters $\alpha_t, \beta_t, \gamma_t$ satisfy the following scaling conditions: $\alpha_t = \log p - \log t$, $\beta_t = p \log t + \log C$, and $\gamma_t = p \log t$ where $p \geq 2$ and $C > 0$ are constants.

The variational problem is

$$\begin{aligned} \text{Minimize}_u \quad & J(u) = \int_0^\infty L(t, X_t, u_t) dt \\ \text{Subject to} \quad & \frac{dX_t}{dt} = u_t, \quad X_0 = x_0 \end{aligned} \quad (3)$$

where u_t is the control input chosen to minimize the objective function $J(u)$. over all control laws $\{u_t\}_{t>0}$ in \mathbb{R}^d .

The Hamiltonian function is

$$H(t, x, y, u) = y \cdot u - L(t, x, u) \quad (4)$$

	Vector	Probability distribution
State-space	\mathbb{R}^d	$\mathcal{P}_2(\mathbb{R}^d)$
Objective function	$f(x)$	$F(\rho) := D(\rho \rho_\infty)$
Lagrangian	$e^{\alpha_t+\gamma_t} \left(\frac{1}{2} e^{-\alpha_t} u ^2 - e^{\beta_t} f(x) \right)$	$e^{\alpha_t+\gamma_t} \mathbb{E} \left[\frac{1}{2} e^{-\alpha_t} U ^2 - e^{\beta_t} \log\left(\frac{\rho(X)}{\rho_\infty(X)}\right) \right]$
Lyapunov funct.	$\frac{1}{2} x + e^{-\gamma_t} y - \bar{x} ^2$ $+ e^{\beta_t} (f(x) - f(\bar{x}))$	$\frac{1}{2} \mathbb{E}[X_t + e^{-\gamma_t} Y_t - T_{\rho_t}^{\rho_\infty}(X_t) ^2]$ $+ e^{\beta_t} (F(\rho_t) - F(\rho_\infty))$
Convergence rate	$f(x_t) - f(\bar{x}) \leq O(e^{-\beta t})$	$F(\rho_t) - F(\rho_\infty) \leq O(e^{-\beta t})$

Table 1. Summary of the variational formulations for vectors and probability distributions.

where $y \in \mathbb{R}^d$ is dual variable and $y \cdot u$ denotes the dot product between vectors y and u .

According to the Pontryagin's Maximum Principle, the optimal control $u_t^* = \arg \max_v H(t, X_t, Y_t, v) = e^{\alpha_t - \gamma_t} Y_t$. The resulting Hamilton's equations are

$$\frac{dX_t}{dt} = +\nabla_y H(t, X_t, Y_t, u_t^*) = e^{\alpha_t - \gamma_t} Y_t, \quad X_0 = x_0 \quad (5a)$$

$$\frac{dY_t}{dt} = -\nabla_x H(t, X_t, Y_t, u_t^*) = -e^{\alpha_t + \beta_t + \gamma_t} \nabla f(X_t), \quad (5b)$$

The system (5) is an example of accelerated gradient descent algorithm. Specifically, if the parameters $\alpha_t, \beta_t, \gamma_t$ are defined using $p = 2$, one obtains the continuous-time limit of the Nesterov's accelerated algorithm. It is referred to as the Nesterov's ODE in this paper.

For this system, a Lyapunov function is as follows:

$$V(t, x, y) = \frac{1}{2} |x + e^{-\gamma_t} y - \bar{x}|^2 + e^{\beta_t} (f(x) - f(\bar{x})) \quad (6)$$

where $\bar{x} \in \arg \min_x f(x)$. It is shown in (Wibisono et al., 2016) that upon differentiating along the solution trajectory, $\frac{d}{dt} V(t, X_t, Y_t) \leq 0$. This yields the convergence rate:

$$f(X_t) - f(\bar{x}) \leq O(e^{-\beta t}), \quad \forall t \geq 0 \quad (7)$$

3. Variational formulation for probability distributions

3.1. Motivation and background

Let $F : \mathcal{P}_{ac,2}(\mathbb{R}^d) \rightarrow \mathbb{R}$ be a functional on the space of probability distributions. Consider the problem of minimizing $F(\rho)$. The (Wasserstein) gradient flow with respect to $F(\rho)$ is a curve ρ_t such that

$$\frac{\partial \rho_t}{\partial t}(x) = \nabla \cdot (\rho_t(x) \nabla_\rho F(\rho_t)(x)) \quad (8)$$

where (the vector field) $\nabla_\rho F(\rho) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the Wasserstein gradient of F .

An important example is the relative entropy functional where $F(\rho) = D(\rho|\rho_\infty) := \int_{\mathbb{R}^d} \log\left(\frac{\rho(x)}{\rho_\infty(x)}\right) \rho(x) dx$ where $\rho_\infty \in \mathcal{P}_{ac,2}(\mathbb{R}^d)$ is referred to as the target distribution. The gradient of relative entropy is given by $\nabla_\rho F(\rho)(x) = \nabla \log\left(\frac{\rho(x)}{\rho_\infty(x)}\right)$ (Ch. 8.3 in (Villani, 2003)). The gradient flow

$$\frac{\partial \rho_t}{\partial t}(x) = -\nabla \cdot (\rho_t(x) \nabla \log(\rho_\infty(x))) + \Delta \rho_t(x) \quad (9)$$

is the Fokker-Planck equation (Jordan et al., 1998). The gradient flow achieves the density transport from an initial probability distribution ρ_0 to the target (here, also equilibrium) probability distribution ρ_∞ ; and underlies the construction and the analysis of Markov chain Monte-Carlo (MCMC) algorithms. The simplest MCMC algorithm is the Langevin stochastic differential equation (SDE):

$$dX_t = -\nabla f(X_t) dt + \sqrt{2} dB_t, \quad X_0 \sim \rho_0$$

where B_t is the standard Brownian motion in \mathbb{R}^d .

The main problem of this paper is to construct an accelerated form of the gradient flow (8). The proposed solution is based upon a variational formulation. As tabulated in Table 1, the solution represents a generalization of (Wibisono et al., 2016) from its original deterministic finite-dimensional to now probabilistic infinite-dimensional settings.

The variational problem can be expressed in two equivalent forms: (i) The probabilistic form and (ii) The partial differential equation (PDE)

3.2. Probabilistic form of the variational problem

Consider the stochastic process $\{X_t\}_{t \geq 0}$ that takes values in \mathbb{R}^d and evolves according to:

$$\frac{dX_t}{dt} = U_t, \quad X_0 \sim \rho_0$$

where the control input $\{U_t\}_{t \geq 0}$ also takes values in \mathbb{R}^d , and $\rho_0 \in \mathcal{P}_{ac,2}(\mathbb{R}^d)$ is the probability distribution of the initial condition X_0 . It is noted that the randomness here comes only from the random initial condition.

Suppose the objective functional is of the form $F(\rho) = \int \tilde{F}(\rho, x) \rho(x) dx$. The Lagrangian $L : \mathbb{R}^+ \times \mathbb{R}^d \times \mathcal{P}_{ac,2}(\mathbb{R}^d) \times \mathbb{R}^d \rightarrow \mathbb{R}$ is defined as

$$L(t, x, \rho, u) := e^{\alpha_t + \gamma_t} \left(\underbrace{\frac{1}{2} |e^{-\alpha_t} u|^2}_{\text{kinetic energy}} - \underbrace{e^{\beta_t} \tilde{F}(\rho, x)}_{\text{potential energy}} \right) \quad (10)$$

This formula is a natural generalization of the Lagrangian (2) and the parameters $\alpha_t, \beta_t, \gamma_t$ are defined exactly the same as in the finite-dimensional case. The stochastic optimal control problem is:

$$\begin{aligned} \text{Minimize } J(u) &= \mathbb{E} \left[\int_0^\infty L(t, X_t, \rho_t, U_t) dt \right] \\ \text{Subject to } \frac{dX_t}{dt} &= U_t, \quad X_0 \sim \rho_0 \end{aligned} \quad (11)$$

where $\rho_t = \text{Law}(X_t) \in \mathcal{P}_{ac,2}(\mathbb{R}^d)$ is the probability density function of the random variable X_t .

The Hamiltonian function $H : \mathbb{R}^+ \times \mathbb{R}^d \times \mathcal{P}_{ac,2}(\mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ for this problem is given by (see Sec. 6.2.3 in (Carmona & Delarue, 2017)):

$$H(t, x, \rho, y, u) := u \cdot y - L(t, x, \rho, u) \quad (12)$$

where $y \in \mathbb{R}^d$ is the dual variable.

Remark 1. The variational problem (11) is an example of a mean-field (McKean-Vlasov) optimal control problem. This is because the Lagrangian depends also upon the law of the stochastic process; cf., Ch. 6 in (Carmona & Delarue, 2017).

3.3. PDE formulation of the variational problem

An equivalent pde formulation is obtained by considering the stochastic optimal control problem (11) as a deterministic optimal control problem on the space of the probability distributions. Specifically, the process $\{\rho_t\}_{t \geq 0}$ is a deterministic process that takes values in $\mathcal{P}_{ac,2}(\mathbb{R}^d)$ and evolves according to the continuity equation

$$\frac{\partial \rho_t}{\partial t} = -\nabla \cdot (\rho_t u_t)$$

where $u_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is now a time-varying vector field. The Lagrangian $\mathcal{L} : \mathbb{R}^+ \times \mathcal{P}_{ac,2}(\mathbb{R}^d) \times L^2(\mathbb{R}^d; \mathbb{R}^d) \rightarrow \mathbb{R}$ is defined as:

$$\mathcal{L}(t, \rho, u) := e^{\alpha_t + \gamma_t} \left[\int_{\mathbb{R}^d} \frac{1}{2} |e^{-\alpha_t} u(x)|^2 \rho(x) dx - e^{\beta_t} F(\rho) \right] \quad (13)$$

The optimal control problem is:

$$\begin{aligned} \text{Minimize } & \int_0^\infty \mathcal{L}(t, \rho_t, u_t) dt \\ \text{Subject to } & \frac{\partial \rho_t}{\partial t} + \nabla \cdot (\rho_t u_t) = 0 \end{aligned} \quad (14)$$

The Hamiltonian function $\mathcal{H} : \mathbb{R}^+ \times \mathcal{P}_{ac,2}(\mathbb{R}^d) \times \mathcal{C}(\mathbb{R}^d; \mathbb{R}) \times L^2(\mathbb{R}^d; \mathbb{R}^d) \rightarrow \mathbb{R}$ is

$$\mathcal{H}(t, \rho, \phi, u) := \langle \nabla \phi, u \rangle_{L^2(\rho)} - \mathcal{L}(t, \rho, u) \quad (15)$$

where $\phi \in \mathcal{C}(\mathbb{R}^d; \mathbb{R})$ is the dual variable and the inner-product $\langle \nabla \phi, u \rangle_{L^2(\rho)} := \int_{\mathbb{R}^d} \nabla \phi(x) \cdot u(x) \rho(x) dx$

3.4. Main result

Theorem 1. Consider the variational problem (11)-(14).

(i) For the probabilistic form (11) of the variational problem, the optimal control $U_t^* = e^{\alpha_t - \gamma_t} Y_t$, where the optimal trajectory $\{(X_t, Y_t)\}_{t \geq 0}$ evolves according to the Hamilton's odes:

$$\frac{dX_t}{dt} = U_t^* = e^{\alpha_t - \gamma_t} Y_t, \quad X_0 \sim \rho_0 \quad (16a)$$

$$\frac{dY_t}{dt} = -e^{\alpha_t + \beta_t + \gamma_t} \nabla_\rho F(\rho_t)(X_t), \quad Y_0 = \nabla \phi_0(X_0) \quad (16b)$$

where ϕ_0 is a convex function, and $\rho_t = \text{Law}(X_t)$.

(ii) For the pde form (14) of the variational problem, the optimal control is $u_t^* = e^{\alpha_t - \gamma_t} \nabla \phi_t(x)$, where the optimal trajectory $\{(\rho_t, \phi_t)\}_{t \geq 0}$ evolves according to the Hamilton's pdes:

$$\frac{\partial \rho_t}{\partial t} = -\nabla \cdot (\rho_t \underbrace{e^{\alpha_t - \gamma_t} \nabla \phi_t}_{u_t^*}), \quad (17a)$$

$$\frac{\partial \phi_t}{\partial t} = -e^{\alpha_t - \gamma_t} \frac{|\nabla \phi_t|^2}{2} - e^{\alpha_t + \gamma_t + \beta_t} \nabla_\rho F(\rho) \quad (17b)$$

(iii) The solutions of the two forms are equivalent in the following sense:

$$\text{Law}(X_t) = \rho_t, \quad U_t = u_t(X_t), \quad Y_t = \nabla \phi_t(X_t)$$

(iv) Suppose additionally that the functional F is displacement convex and ρ_∞ is its minimizer. Define

$$\begin{aligned} V(t) &= \frac{1}{2} \mathbb{E}[|X_t + e^{-\gamma_t} Y_t - T_{\rho_t}^{\rho_\infty}(X_t)|^2] \\ &+ e^{\beta_t} (F(\rho) - F(\rho_\infty)) \end{aligned} \quad (18)$$

where the map $T_{\rho_t}^{\rho_\infty} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the optimal transport map from ρ_t to ρ_∞ . Assume the dimension $d = 1$. Consequently, the following rate of convergence is obtained along the optimal trajectory

$$F(\rho_t) - F(\rho_\infty) \leq O(e^{-\beta t}), \quad \forall t \geq 0$$

Proof sketch. The Hamilton's equations are derived using the standard mean-field optimal control theory (Carmona & Delarue, 2017). The Lyapunov function argument is based upon the variational inequality characterization of a displacement convex function (see Eq. 10.1.7 in (Ambrosio et al., 2008)). The detailed proof appears in the supplementary note. We expect that the assumption that $d = 1$ is not necessary. This is the subject of the continuing work. \square

3.5. Relative entropy as the functional

In the remainder of this paper, we assume that the functional $F(\rho) = D(\rho|\rho_\infty)$ is the relative entropy where $\rho_\infty \in \mathcal{P}_{ac,2}(\mathbb{R}^d)$ is a given target probability distribution. In this case the Hamilton's equations are given by

$$\frac{dX_t}{dt} = e^{\alpha_t - \gamma_t} Y_t, \quad X_0 \sim \rho_0 \quad (19a)$$

$$\frac{dY_t}{dt} = -e^{\alpha_t + \beta_t + \gamma_t} (\nabla f(X_t) + \nabla \log(\rho_t(X_t))), \quad (19b)$$

with $Y_0 = \nabla \phi_0(X_0)$, where $\rho_t = \text{Law}(X_t)$ and $f = -\log(\rho_\infty)$. Moreover, if f is convex (or equivalently ρ_∞ is log-concave), then F is displacement convex with the unique minimizer at ρ_∞ and the convergence estimate is given by $D(\rho_t|\rho_\infty) \leq O(e^{-\beta t})$.

Remark 2. The Hamilton's equations (19) with the relative entropy functional is related to the under-damped Langevin equation (Cheng et al., 2017). A basic form of the under-damped (or second order) Langevin equation is

$$\begin{aligned} dX_t &= v_t dt \\ dv_t &= -\gamma v_t dt - \nabla f(X_t) dt + \sqrt{2} dB_t \end{aligned} \quad (20)$$

where $\{B_t\}_{t \geq 0}$ is the standard Brownian motion.

Consider next, the the accelerated flow (19). Denote $v_t := e^{\alpha_t - \gamma_t} Y_t$. Then, with an appropriate choice of scaling parameters (e.g. $\alpha_t = 0$, $\beta_t = 0$ and $\gamma_t = -\gamma t$):

$$\begin{aligned} dX_t &= v_t dt \\ dv_t &= -\gamma v_t dt - \nabla f(X_t) dt - \nabla_x \log(\rho_t(X_t)) \end{aligned} \quad (21)$$

The scaling parameters are chosen here for the sake of comparison and do not satisfy the ideal scaling conditions of (Wibisono et al., 2016).

The sdes (20) and (21) are similar except that the stochastic term $\sqrt{2} dB_t$ in (20) is replaced with a deterministic term $-\nabla_x \log(\rho_t(X_t))$ in (21). Because of this difference, the resulting distributions are different. See the supplementary note for more details.

3.6. Quadratic Gaussian case

Suppose the initial distribution ρ_0 and the target distribution ρ_∞ are both Gaussian, denoted as $\mathcal{N}(m_0, \Sigma_0)$ and

$\mathcal{N}(\bar{x}, Q)$, respectively. This is equivalent to the objective function $f(x)$ being quadratic of the form $f(x) = \frac{1}{2}(x - \bar{x})^\top Q^{-1}(x - \bar{x})$. Therefore, this problem is referred to as the *quadratic Gaussian case*. The following Proposition shows that the mean of the stochastic process (X_t, Y_t) evolves according to the Nesterov ODE (5):

Proposition 1. (*Quadratic Gaussian case*) Consider the variational problem (11) for the quadratic Gaussian case. Then

- (i) The stochastic process (X_t, Y_t) is a Gaussian process. The Hamilton's equations are given by:

$$\frac{dX_t}{dt} = e^{\alpha_t - \gamma_t} Y_t,$$

$$\frac{dY_t}{dt} = -e^{\alpha_t + \beta_t + \gamma_t} (Q^{-1}(X_t - \bar{x}) - \Sigma_t^{-1}(X_t - m_t))$$

where m_t and Σ_t are the mean and the covariance of X_t .

- (ii) Upon taking the expectation of both sides, and denoting $n_t := \mathbb{E}[Y_t]$

$$\frac{dm_t}{dt} = e^{\alpha_t - \gamma_t} n_t, \quad \frac{dn_t}{dt} = -e^{\alpha_t + \beta_t + \gamma_t} \underbrace{Q^{-1}(m_t - \bar{x})}_{\nabla f(m_t)}$$

which is identical to Nesterov ODE (5).

Proof sketch. Fix ρ_t . Consider the resulting pair (X_t, Y_t) from (19) and let $\tilde{\rho}_t = \text{Law}(X_t)$. The proof follows from showing that a Gaussian ρ_t is a fixed-point of the map $\rho_t \mapsto \tilde{\rho}_t$. \square

4. Numerical algorithm

The proposed numerical algorithm is based upon an interacting particle implementation of the Hamilton's equation (19). Consider a system of N particles $\{(X_t^i, Y_t^i)\}_{i=1}^N$ that evolve according to:

$$\begin{aligned} \frac{dX_t^i}{dt} &= e^{\alpha_t - \gamma_t} Y_t^i, \quad X_0^i \sim \rho_0 \\ \frac{dY_t^i}{dt} &= -e^{\alpha_t + \beta_t + \gamma_t} (\nabla f(X_t^i) + \underbrace{I_t^{(N)}(X_t^i)}_{\text{interaction term}}), \end{aligned}$$

with $Y_0^i = \nabla \phi_0(X_0^i)$. The interaction term $I_t^{(N)}$ is an empirical approximation of the $\nabla \log(\rho_t)$ term in (19). We propose two types of empirical approximations as follows:

1. Gaussian approximation: Suppose the density is approximated as a Gaussian $\mathcal{N}(m_t, \Sigma_t)$. In this case, $\nabla \log(\rho_t(x)) = -\Sigma_t^{-1}(x - m_t)$. This motivates the following empirical approximation of the interaction term:

$$I_t^{(N)}(x) = -\Sigma_t^{(N)-1}(x - m_t^{(N)}) \quad (23)$$

where $m_t^{(N)} := N^{-1} \sum_{i=1}^N X_t^i$ is the empirical mean and $\Sigma_t^{(N)} := \frac{1}{N-1} \sum_{i=1}^N (X_t^i - m_t^{(N)})(X_t^i - m_t^{(N)})^\top$ is the empirical covariance.

Even though the approximation is asymptotically (as $N \rightarrow \infty$) exact only under the Gaussian assumption, it may be used in a more general settings, particularly when the density ρ_t is unimodal. The situation is analogous to the (Bayesian) filtering problem, where an ensemble Kalman filter is used as an approximate solution for non-Gaussian distributions (Evensen, 2003).

2. Diffusion map approximation: This is based upon the diffusion map approximation of the weighted Laplacian operator (Coifman & Lafon, 2006; Hein et al., 2007). For a C^2 function f , the weighted Laplacian is defined as $\Delta_\rho f := \frac{1}{\rho} \nabla \cdot (\rho \nabla f)$. Denote $e(x) = x$ as the coordinate function on \mathbb{R}^d . It is a straightforward calculation to show that $\nabla \log(\rho) = \Delta_\rho e$. This allows one to use the diffusion map approximation of the weighted Laplacian to approximate the interaction term as follows:

$$(DM) \quad I_t^{(N)}(X_t^i) = \frac{1}{\epsilon} \frac{\sum_{j=1}^N k_\epsilon(X_t^i, X_t^j)(X_t^j - X_t^i)}{\sum_{j=1}^N k_\epsilon(X_t^i, X_t^j)} \quad (24)$$

where the kernel $k_\epsilon(x, y) = \frac{g_\epsilon(x, y)}{\sqrt{\sum_{i=1}^N g_\epsilon(y, X_t^i)}}$ is constructed empirically in terms of the Gaussian kernel $g_\epsilon(x, y) = \exp(-|x - y|^2 / (4\epsilon))$. The parameter ϵ is referred to as the kernel bandwidth. The approximation is asymptotically exact as $\epsilon \downarrow 0$ and $N \uparrow \infty$. The approximation error is of order $O(\epsilon) + O(\frac{1}{\sqrt{N\epsilon^{d/4}}})$ where the first term is referred to as the bias error and the second term is referred to as the variance error (Hein et al., 2007). The variance error is the dominant term in the error for small values of ϵ , whereas the bias error is the dominant term for large values of ϵ (see Figure 2(d)).

The resulting interacting particle algorithm is tabulated in Table 1. The symplectic method proposed in (Betancourt et al., 2018) is used to carry out the numerical integration. The algorithm is applied to two examples as described in the following sections.

Remark 3. For the case where there is only one particle ($N = 1$), the interaction term is zero and the system (22) reduces to the Nesterov ODE (5).

Remark 4. (Comparison with density estimation) The diffusion map approximation algorithm is conceptually different from an explicit density estimation-based approach. A basic density estimation is to approximate $\rho(x) \approx \frac{1}{N} \sum_{i=1}^N g_\epsilon(x, X_t^i)$ where $g_\epsilon(x, y)$ is the Gaussian kernel. Using such an approximation, the interaction term is approximated as

$$(DE) \quad I_t^{(N)}(X_t^i) = \frac{1}{\epsilon} \frac{\sum_{j=1}^N g_\epsilon(X_t^i, X_t^j)(X_t^j - X_t^i)}{2 \sum_{j=1}^N g_\epsilon(X_t^i, X_t^j)} \quad (25)$$

Algorithm 1 Interacting particle implementation of the accelerated gradient flow

Input: $\rho_0, \phi_0, N, t_0, \Delta t, p, C, K$

Output: $\{X_k^i\}_{i=1, k=0}^{N, K}$

Initialize $\{X_0^i\}_{i=1}^N \stackrel{i.i.d.}{\sim} \rho_0, Y_0^i = \nabla \phi_0(X_0^i)$

Compute $I_0^{(N)}(X_0^i)$ with (23) or (24)

for $k = 0$ to $K - 1$ **do**

$t_{k+\frac{1}{2}} = t_k + \frac{1}{2} \Delta t$

$Y_{k+\frac{1}{2}}^i = Y_k^i - \frac{1}{2} C p t_{k+\frac{1}{2}}^{2p-1} (\nabla f(X_k^i) + I_k^{(N)}(X_k^i)) \Delta t$

$X_{k+1}^i = X_k^i + \frac{p}{t_{k+\frac{1}{2}}^{p+1}} Y_k^i \Delta t$

Compute $I_{k+1}^{(N)}(X_{k+1}^i)$ with (23) or (24)

$Y_{k+1}^i = Y_{k+\frac{1}{2}}^i - \frac{1}{2} C p t_{k+\frac{1}{2}}^{2p-1} (\nabla f(X_{k+1}^i) +$

$I_{k+1}^{(N)}(X_{k+1}^i)) \Delta t$

$t_{k+1} = t_{k+\frac{1}{2}} + \frac{1}{2} \Delta t$

end for

Despite the apparent similarity of the two formulae, (24) for diffusion map approximation and (25) for density estimation, the nature of the two approximations is different. The difference arises because the kernel $k_\epsilon(x, y)$ in (24) is data-dependent whereas the kernel in (25) is not. While both approximations are exact in the asymptotic limit as $N \uparrow \infty$ and $\epsilon \downarrow 0$, they exhibit different convergence rates. Numerical experiments presented in Figure 2(a)-(d) show that the diffusion map approximation has a much smaller variance for intermediate values of N . Theoretical understanding of the difference is the subject of continuing work.

4.1. Gaussian Example

Consider the Gaussian example as described in Sec. 3.6. The simulation results for the scalar ($d = 1$) case with initial distribution $\rho_0 = \mathcal{N}(2, 4)$ and target distribution $\mathcal{N}(\bar{x}, Q)$ where $\bar{x} = -5.0$ and $Q = 0.25$ is depicted in Figure 1-(a)-(b). For this simulation, the numerical parameters are as follows: $N = 100, \phi_0(x) = 0.5(x - 2), t_0 = 1, \Delta t = 0.1, p = 2, C = 0.625$, and $K = 400$. The result numerically verifies the $O(e^{-\beta t}) = O(\frac{1}{t^2})$ convergence rate derived in Theorem 1 for the case where the target distribution is Gaussian.

4.2. Non-Gaussian example

This example involves a non-Gaussian target distribution $\rho_\infty = \frac{1}{2} \mathcal{N}(-m, \sigma^2) + \frac{1}{2} \mathcal{N}(m, \sigma^2)$ which is a mixture of two one-dimensional Gaussians with $m = 2.0$ and $\sigma^2 = 0.8$. The simulation results are depicted in Figure 1-(c)-(d). The numerical parameters are same as in the Example 4.1. The interaction term is approximated using the diffusion map approximation with $\epsilon = 0.01$. The numerical result depicted in Figure 1-(c) show that the diffusion map

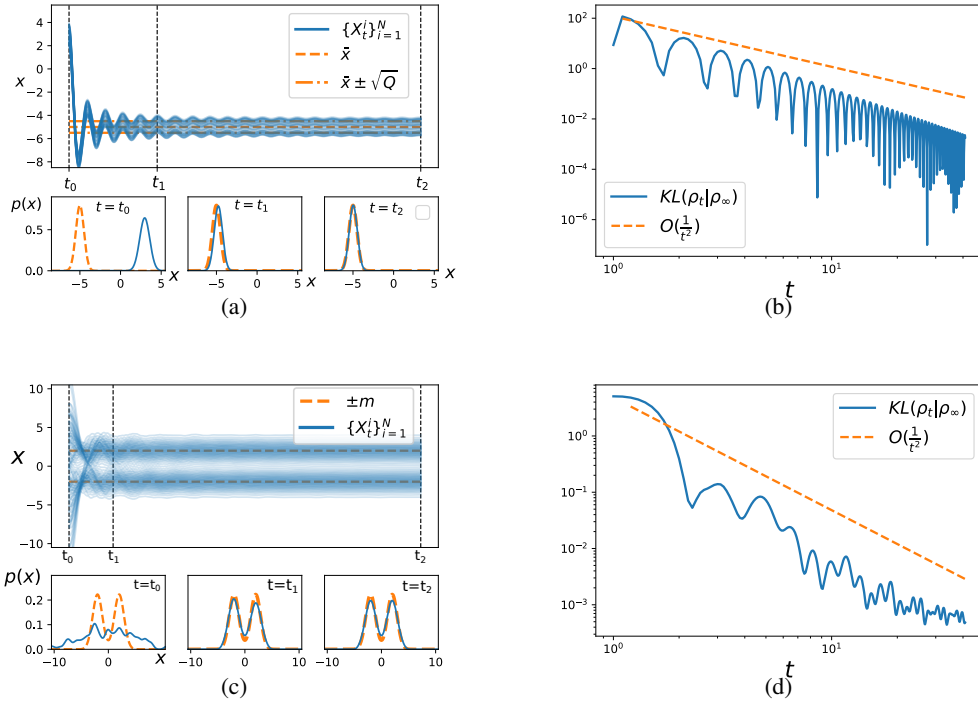


Figure 1. Simulation result for the Gaussian case (Example 4.1): (a) The time traces of the particles; (b) The KL-divergence as a function of time. Simulation result for the non-Gaussian case (Example 4.2): (c) The time traces of the particles; (d) The KL-divergence as a function of time.

algorithm converges to the mixture of Gaussian target distribution. The result depicted in Figure 1-(d) suggests that the convergence rate $O(e^{-\beta t})$ also appears to hold for this non-log-concave target distribution. Theoretical justification of this is subject of continuing work.

4.3. Comparison with MCMC and HMC

This section contains numerical experiment comparing the performance of the accelerated algorithm 1 using the diffusion map (DM) approximation (24) and the density estimation (DE)-based approximation (25) with the Markov chain Monte-Carlo (MCMC) algorithm studied in (Durmus & Moulines, 2016) and the Hamiltonian MCMC algorithm studied in (Cheng et al., 2017).

We consider the problem setting of the mixture of Gaussians as in example 4.2. All algorithms are simulated with a fixed step-size of $\Delta t = 0.1$ for $K = 1000$ iterations. The performance is measured by computing the mean-squared error in estimating the expectation of the function $\psi(x) = x1_{x \geq 0}$ denoted as $\hat{\psi} := \int \psi(x)\rho_\infty(x) dx$. The mean-square error at the k -th iteration is computed by averaging the error over

$M = 100$ runs:

$$\text{m.s.e}_k = \frac{1}{M} \sum_{m=1}^M \left(\frac{1}{N} \sum_{i=1}^N \psi(X_{t_k}^{i,m}) - \hat{\psi} \right)^2 \quad (26)$$

The numerical results are depicted in Figure 2. Figure 2(a) depicts the m.s.e as a function of N . It is observed that the accelerated algorithm 1 with the diffusion map approximation admits an order of magnitude better m.s.e for the same number of particles. It is also observed that the m.s.e decreases rapidly for intermediate values of N before saturating for large values of N , where the bias term dominates (see discussion following Eq. 24).

Figure 2(b) depicts the m.s.e as a function of the number of iterations for a fixed number of particles $N = 100$. It is observed that the accelerated algorithm 1 displays the quickest convergence amongst the algorithms tested.

Figure 2(c) depicts the average computational time per iteration as a function of the number of samples N . The computational time of the diffusion map approximation scales as $O(N^2)$ because it involves computing a $N \times N$ matrix $[k_\epsilon(X^i, X^j)]_{i,j=1}^N$, while the computational cost of the MCMC and HMC algorithms scale as $O(N)$. The computational complexity may be improved by (i) exploiting the

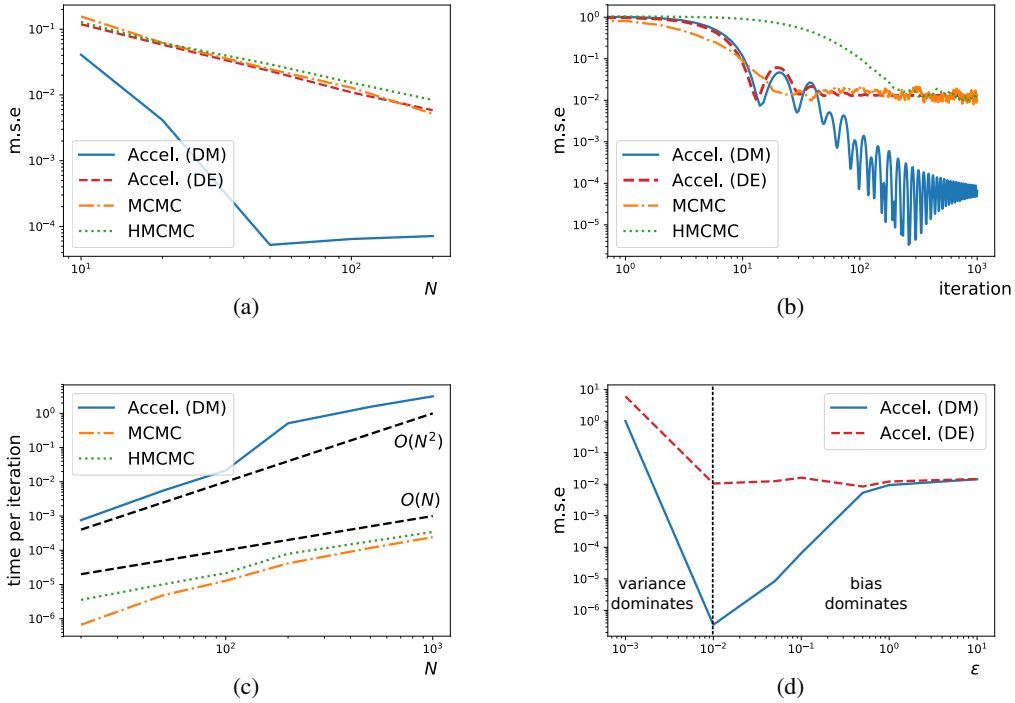


Figure 2. Simulation-based comparison of the performance of the accelerated algorithm 1 using the diffusion map (DM) approximation (24), the density estimation (DE)-based approximation (25) with the MCMC and HMCMC algorithms: (a) the mean-squared error (m.s.e) (26) as a function of the number of samples N ; (b) the m.s.e as a function of the number of iterations; (c) the average computational time per iteration as a function of the number of samples; (d) m.s.e comparison between the diffusion map and the density estimation-based approaches as a function of the kernel bandwidth ϵ .

sparsity structure of the $N \times N$ matrix; (ii) sub-sampling the particles in computing the empirical averages; (iii) adaptively updating the $N \times N$ matrix according to a certain error criteria.

Finally, we provide comparison between diffusion map approximation (25) and the density-based approximation (25): Figure 2(d) depicts the m.s.e for these two approximations as a function of the kernel-bandwidth ϵ for a fixed number of particles $N = 100$. For very large and for very small values of ϵ , where bias and variance dominates the error, respectively, the two algorithms have similar m.s.e. However, for intermediate values of ϵ , the diffusion map approximation has smaller variance, and thus lower m.s.e.

5. Conclusion and directions for future work

The main contribution of this paper is to extend the variational formulation of (Wibisono et al., 2016) to the space of probability distributions (see Theorem 1). In particular, we obtain theoretical results and numerical algorithms for accelerated gradient flow in such settings. Through this procedure, we provide an accelerated version of the Langevin equation (see equation (19)), and compare it with the under-

damped Langevin equation (see Remark 2). Moreover, we recover the ode limit of the Nesterov accelerated method in the special quadratic Gaussian case (see Proposition 1).

Two algorithms based upon an interacting particle representation are presented and illustrated with numerical examples. The examples serve to verify the theoretical result, and demonstrate that our proposed approach admits faster convergence and better accuracy with the same number of samples, compared to the standard MCMC and HMCMC methods. The main bottleneck for the implementation of our approach is the $O(N^2)$ computational complexity, which is common to the related interacting particle based approaches in the literature. This is an important problem in its own right and subject of future work.

Some direction for future include: (i) removing the assumption $d = 1$ in the convergence result of the Theorem 1; (ii) using the variational formulation to obtain meaningful approximations of the interaction term for the finite- N algorithm (ii) analysis of the convergence under the weaker assumption that the target distribution satisfies the logarithmic-Sobolev inequality; and (iii) error analysis of the numerical algorithms in the finite- N and in the discrete-time setting.

Acknowledgements

Financial support from the NSF grant CMMI-1462773 and ARO grant W911NF1810334 is gratefully acknowledged.

References

- Ambrosio, L., Gigli, N., and Savaré, G. Gradient flows: in metric spaces and in the space of probability measures. Springer Science & Business Media, 2008.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein gan. arXiv preprint arXiv:1701.07875, 2017.
- Bensoussan, A., Frehse, J., Yam, P., et al. Mean field games and mean field type control theory, volume 101. Springer, 2013.
- Betancourt, M., Jordan, M. I., and Wilson, A. C. On symplectic optimization. arXiv preprint arXiv:1802.03653, 2018.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. Variational inference: A review for statisticians. Journal of the American Statistical Association, 112(518):859–877, 2017.
- Carmona, R. and Delarue, F. Probabilistic Theory of Mean Field Games with Applications I-II. Springer, 2017.
- Chen, C., Zhang, R., Wang, W., Li, B., and Chen, L. A unified particle-optimization framework for scalable bayesian sampling. arXiv preprint arXiv:1805.11659, 2018.
- Cheng, X., Chatterji, N. S., Bartlett, P. L., and Jordan, M. I. Underdamped langevin mcmc: A non-asymptotic analysis. arXiv preprint arXiv:1707.03663, 2017.
- Chizat, L. and Bach, F. On the global convergence of gradient descent for over-parameterized models using optimal transport. arXiv preprint arXiv:1805.09545, 2018.
- Coifman, R. R. and Lafon, S. Diffusion maps. Applied and computational harmonic analysis, 21(1):5–30, 2006.
- Del Moral, P. et al. Measure-valued processes and interacting particle systems. application to nonlinear filtering problems. The Annals of Applied Probability, 8(2):438–495, 1998.
- Durmus, A. and Moulines, E. High-dimensional bayesian inference via the unadjusted langevin algorithm. arXiv preprint arXiv:1605.01559, 2016.
- Evensen, G. The ensemble kalman filter: Theoretical formulation and practical implementation. Ocean dynamics, 53(4):343–367, 2003.
- Frogner, C. and Poggio, T. Approximate inference with wasserstein gradient flows. arXiv preprint arXiv:1806.04542, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In Advances in neural information processing systems, pp. 2672–2680, 2014.
- Hein, M., Audibert, J.-Y., and Luxburg, U. v. Graph laplacians and their convergence on random neighborhood graphs. Journal of Machine Learning Research, 8(Jun): 1325–1368, 2007.
- Jain, P., Kakade, S. M., Kidambi, R., Netrapalli, P., and Sidford, A. Accelerating stochastic gradient descent. arXiv preprint arXiv:1704.08227, 2017.
- Jordan, R., Kinderlehrer, D., and Otto, F. The variational formulation of the fokker–planck equation. SIAM journal on mathematical analysis, 29(1):1–17, 1998.
- Liu, C., Zhuo, J., Cheng, P., Zhang, R., Zhu, J., and Carin, L. Accelerated first-order methods on the wasserstein space for bayesian inference. arXiv preprint arXiv:1807.01750, 2018.
- Liu, Q. Stein variational gradient descent as gradient flow. In Advances in neural information processing systems, pp. 3115–3123, 2017.
- Liu, Q. and Wang, D. Stein variational gradient descent: A general purpose bayesian inference algorithm. In Advances In Neural Information Processing Systems, pp. 2378–2386, 2016.
- Liu, Y., Shang, F., Cheng, J., Cheng, H., and Jiao, L. Accelerated first-order methods for geodesically convex optimization on riemannian manifolds. In Advances in Neural Information Processing Systems, pp. 4868–4877, 2017.
- McCann, R. J. A convexity principle for interacting gases. Advances in mathematics, 128(1):153–179, 1997.
- Neal, R. M. et al. Mcmc using hamiltonian dynamics. Handbook of Markov Chain Monte Carlo, 2(11):2, 2011.
- Reich, S. A dynamical systems framework for intermittent data assimilation. BIT Numerical Mathematics, 51(1): 235–249, 2011.
- Richemond, P. H. and Maginnis, B. On wasserstein reinforcement learning and the fokker-planck equation. arXiv preprint arXiv:1712.07185, 2017.
- Su, W., Boyd, S., and Candes, E. A differential equation for modeling nesterov’s accelerated gradient method: Theory and insights. In Advances in Neural Information Processing Systems, pp. 2510–2518, 2014.

- Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In Advances in neural information processing systems, pp. 1057–1063, 2000.
- Villani, C. Topics in optimal transportation. Number 58. American Mathematical Soc., 2003.
- Wibisono, A., Wilson, A. C., and Jordan, M. I. A variational perspective on accelerated methods in optimization. Proceedings of the National Academy of Sciences, pp. 201614734, 2016.
- Yang, T., Laugesen, R. S., Mehta, P. G., and Meyn, S. P. Multivariable feedback particle filter. Automatica, 71: 10–23, 2016.
- Zhang, C., Taghvaei, A., and Mehta, P. G. A mean-field optimal control formulation for global optimization. IEEE Transactions on Automatic Control, 64(1): 279–286, 2019.
- Zhang, R., Chen, C., Li, C., and Carin, L. Policy optimization as wasserstein gradient flows. arXiv preprint arXiv:1808.03030, 2018.