

WASSERSTEIN NATURAL GRADIENT IN STATISTICAL MANIFOLDS WITH CONTINUOUS SAMPLE SPACE

YIFAN CHEN AND WUCHEN LI

ABSTRACT. We study the Wasserstein natural gradient in parametric statistical models with continuous sample spaces. Our approach is to pull back the L^2 -Wasserstein metric tensor in the probability density space to a parameter space, equipping the latter with a positive definite metric tensor, under which it becomes a Riemannian manifold, named the Wasserstein statistical manifold. In general, it is not a totally geodesic sub-manifold of the density space, and therefore its geodesics will differ from the Wasserstein geodesics, except for the well-known Gaussian distribution case, a fact which can also be validated under our framework. We use the sub-manifold geometry to derive a gradient flow and natural gradient descent method in the parameter space. When parametrized densities lie in \mathbb{R} , the induced metric tensor establishes an explicit formula. In optimization problems, we observe that the natural gradient descent outperforms the standard gradient descent when the Wasserstein distance is the objective function. In such a case, we prove that the resulting algorithm behaves similarly to the Newton method in the asymptotic regime. The proof calculates the exact Hessian formula for the Wasserstein distance, which further motivates another preconditioner for the optimization process. To the end, we present examples to illustrate the effectiveness of the natural gradient in several parametric statistical models, including the Gaussian measure, Gaussian mixture, Gamma distribution, and Laplace distribution.

1. INTRODUCTION

The statistical distance between probability measures plays an important role in many fields such as data analysis and machine learning, which usually consist in minimizing a loss function as

$$\text{minimize } d(\rho, \rho_e) \quad \text{s.t. } \rho \in \mathcal{P}_\theta.$$

Here \mathcal{P}_θ is a parameterized subset of the probability density space, and ρ_e is the target density, which is often an empirical realization of a ground-truth distribution. The function d quantifies the difference between densities ρ and ρ_e .

An important example for d is the Kullback-Leibler (KL) divergence, also known as the relative entropy, which closely relates to the maximum likelihood estimate in statistics and the field of information geometry [2][7]. The Hessian operator of KL embeds \mathcal{P}_θ as a statistical manifold, in which the Riemannian metric is the Fisher-Rao metric. Due to Chentsov [15], the Fisher-Rao metric is the only one, up to scaling, that is invariant under sufficient statistics. Using Fisher-Rao metric, a natural gradient descent method, realized by a Forward-Euler discretization of the gradient flow in the manifold, has been introduced. It has found many successful applications in a variety of problems such as blind source separation [3], machine learning [1][30], filtering [36][37], statistics [27][28], optimization [29][38][39] and many others.

Key words and phrases. Optimal Transport; Information Geometry; Wasserstein Statistical Manifold; Wasserstein Natural Gradient.

Recently, the Wasserstein distance, introduced through the field of optimal transport (OT), has been attracting increasing attention in computation and learning [41]. One promising property of the Wasserstein distance is its ability to reflect the metric on sample space, rendering it very useful in machine learning [6][20][35], statistical models [12][14] and geophysics [18][19][13][32][33]. Further, OT theory provides the L^2 -Wasserstein metric tensor, which gives the probability density space (with smooth positive densities in a compact domain) a formal infinite-dimensional Riemannian differential structure [21][24]. This structure can be made mathematically rigorous and general to work for probability measures with finite second moments by using tools in metric geometry [5]. Under such a setting, the gradient flow with respect to the L^2 -Wasserstein metric tensor, known as the Wasserstein gradient flow, is well-defined and has been seen deep connections to fluid dynamics [11][40], differential geometry [25] and mean field games [16][17].

Nevertheless, compared to the Fisher-Rao metric, the Riemannian structure of the Wasserstein metric is mostly investigated in the whole probability space rather than the parameterized subset \mathcal{P}_θ . Therefore, there remains a gap in developing the natural gradient concept in a parametric model within the Wasserstein geometry context. Here we are primarily interested in the question whether there exist the Wasserstein metric tensor and the associated Wasserstein natural gradient in a general parameterized subset and whether we can gain computational benefits by considering these structures. We believe the answer to it will serve as a window to bring synergies between information geometry and optimal transport communities. We note that our motivation is very similar to [23], which introduces the natural gradient concept into parametric statistical models with discrete sample spaces, and discusses its connection and difference to the Fisher-Rao natural gradient in information geometry.

In this paper, we embed the Wasserstein geometry to parametric probability models with continuous sample spaces. Our treatment relies on the same “pull-back” idea as in [23]. Precisely, we pull back the L^2 -Wasserstein metric tensor into the parameter space, making it become a finite-dimensional Riemannian manifold, given the resulted tensor is positive definite. In [23], where the finite discrete sample space and positive probability simplex are considered, the pull-back relation implicitly defines the metric tensor in the parameter space, and its explicit form involves a weighted Laplacian on the graph. In our continuous world, we need solutions to an elliptic equation to define the associated metric tensor. We rely on Assumption 1 to guarantee the invertibility of this operator, which makes sure the metric tensor is well-defined. This assumption holds for the compact sample space and positive smooth densities, and also holds for Gaussian distributions in the unbounded \mathbb{R}^d , as we will see in section 4.1.

We remark that in general, the parametric model will not be a totally geodesic submanifold of the density manifold, except for the well-known Gaussian Wasserstein case [26][43][10], see also section 4.1. Thus generally, we will obtain a new geodesic distance in the parametric model that is different from the Wasserstein metric. However, this will not be a problem when it is used for certain computational purposes, as we will demonstrate both theoretically (in Theorem 1 and Proposition 6) and numerically (in section 4), the natural gradient concept derived from the submanifold geometry will still approximate the second order information of the geodesic distance in the whole density manifold. In the meantime, we will benefit from the fact that the new geometry is on the finite dimensional parameter space. This makes it easier to deal with mathematically compared to the infinite dimensional density manifold.

The Riemannian geometry in the parameter space allows us to derive the constrained Wasserstein gradient flow in it. The discretized version of the flow leads to the Wasserstein natural gradient descent method, in which the induced metric tensor acts as a preconditioning term in the standard gradient descent iteration. This is a standard approach for introducing the natural gradient concept and has been used in Fisher-Rao natural gradient [1] and Wasserstein natural gradient (discrete sample space) [23]. In our Wasserstein natural gradient (continuous sample space), when the dimension of densities is one, we obtain an explicit formula of the metric tensor. Precisely, given $\rho(x, \theta)$ as a parameterized density, $x \in \mathbb{R}^1$ and $\theta \in \Theta \subset \mathbb{R}^d$, the L^2 -Wasserstein metric tensor on Θ will be

$$G_W(\theta) = \int \frac{1}{\rho(x, \theta)} (\nabla_{\theta} F(x, \theta))^T \nabla_{\theta} F(x, \theta) dx,$$

where $F(y, \theta) = \int_{-\infty}^y \rho(y, \theta) dy$ is the cumulative distribution function of $\rho(x, \theta)$. We apply the natural gradient descent induced by $G_W(\theta)$ to the Wasserstein metric modeled problems. It is seen that the Wasserstein gradient descent outperforms the Euclidean and Fisher-Rao natural gradient descent in the iterations. We give theoretical justifications of this phenomenon by showing that the Wasserstein gradient descent behaves asymptotically as the Newton method in such case. A detailed description of the Hessian matrix is also presented by leveraging techniques in one-dimensional OT. Interestingly, this formula also provides us with a new preconditioner for the Wasserstein metric modeled problems, and results in a new algorithm which we call the modified Wasserstein gradient descent. We will compare its performance with the Wasserstein gradient descent in the experiments.

In the literature, there are pioneers toward the constrained Wasserstein gradient flow. [11] studies the density space with a fixed mean and variance. Compared to them, we focus on a density set parameterized by a finite dimensional parameter space. Also, there have been many works linking information geometry and optimal transport [4][45]. In particular, the Wasserstein metric tensor for Gaussian distributions exhibits an explicit form [26][43][10], which leads to extensive studies between the Wasserstein and Fisher-Rao metric for this model. In contrast to their works, our Wasserstein metric tensor can work for general parametric models. For consistency, in section 4.1 we will show that our defined Wasserstein statistical manifold gives the same metric tensor as in the literature when the parametric model is Gaussian. Thus, it can be seen as a direct extension to the Gaussian case. Under such an extension, we are able to discuss the natural gradient for a lot of parametric models in a systematic way.

This paper is organized as follows. In section 2, we briefly review the theory of optimal transport, with a concentration on its Riemannian differential structure. In section 3, we introduce the Wasserstein statistical manifolds by defining the metric tensor in the parameter space directly through the pull-back relation. The Wasserstein gradient flow and natural gradient descent method are then derived. We give a concise study of the metric tensor for one-dimensional densities, showing its connection to the Fisher information matrix. In this case, we theoretically analyze the effect of this natural gradient in the Wasserstein metric modeled problems. In section 4, examples are presented to justify the previous discussions. We also provide a detailed comparison between the performance of the Fisher-Rao and Wasserstein natural gradient in different inference tasks in section 5, with regard to different choices of loss functions, metric tensors and whether or not the ground truth density lies in the parametric family. Finally, we conclude the paper with several discussions in section 6.

2. REVIEW OF OPTIMAL TRANSPORT THEORY

In this section, we briefly review the theory of OT. We note that there are several equivalent definitions of OT, ranging from static to dynamic formulations. In this paper, we focus on the dynamic formulation and its induced Riemannian metric tensor in the density space.

The optimal transport problem is firstly proposed by Monge in 1781: given two probability densities ρ^0, ρ^1 on $\Omega \subset \mathbb{R}^n$, the goal is to find a transport plan $T : \Omega \rightarrow \Omega$ pushing ρ^0 to ρ^1 that minimizes the whole transportation cost, i.e.

$$(1) \quad \inf_T \int_{\Omega} d(x, T(x)) \rho^0(x) dx \quad \text{s.t.} \quad \int_A \rho^1(x) dx = \int_{T^{-1}(A)} \rho^0(x) dx,$$

for any Borel subset $A \subset \Omega$. Here the function $d : \Omega \times \Omega \rightarrow \mathbb{R}$ is the ground distance that measures the difference between x and $T(x)$. In the whole discussions we set $d(x, y) = \|x - y\|^2$ as the square of Euclidean distance. We assume all the densities belong to $\mathcal{P}_2(\Omega)$, which is defined as the collection of probability density functions on $\Omega \subset \mathbb{R}^n$ with finite second moments.

In 1942, Kantorovich relaxed the problem to a linear programming:

$$(2) \quad \min_{\pi \in \Pi(\rho^0, \rho^1)} \int_{\Omega \times \Omega} \|x - y\|^2 \pi(x, y) dx dy,$$

where the infimum is taken over the set Π of joint probability measures on $\Omega \times \Omega$ that have marginals ρ^0, ρ^1 . This formulation finds a wide array of applications in computations [41].

In recent years, OT connects to a variational problem in density space, known as the Benamou-Brenier formula [8]:

$$(3a) \quad \inf_{\Phi_t} \int_0^1 \int_{\Omega} \|\nabla \Phi(t, x)\|^2 \rho(t, x) dx dt,$$

where the infimum is taken over the set of Borel potential functions $\Phi : [0, 1] \times \Omega \rightarrow \mathbb{R}$. Each gradient vector field of potential $\Phi_t = \Phi(t, x)$ on sample space determines a corresponding density path $\rho_t = \rho(t, x)$ as the solution of the continuity equation:

$$(3b) \quad \frac{\partial \rho(t, x)}{\partial t} + \nabla \cdot (\rho(t, x) \nabla \Phi(t, x)) = 0, \quad \rho(0, x) = \rho^0(x), \quad \rho(1, x) = \rho^1(x).$$

Here $\nabla \cdot$ and ∇ are the divergence and gradient operators in \mathbb{R}^n . If Ω is a compact set, the zero flux condition (Neumann condition) is proposed on the boundary of Ω . This is to ensure that $\int_{\Omega} \frac{\partial \rho(t, x)}{\partial t} dx = 0$, so that the total mass is conserved.

Under mild regularity assumptions (for example, if ρ_0, ρ_1 is absolutely continuous with respect to the Lebesgue measure), the above three formulations (1) (2) (3) are equivalent, see details in [44]. Their optimal quantity is denoted by $(W_2(\rho^0, \rho^1))^2$, which is called the square of the L^2 -Wasserstein distance between ρ^0 and ρ^1 . Here the subscript “2” in W_2 indicates that the L^2 ground distance is used. We note that formulation (1) (2) is static, in the sense that only the initial and final states of the transportation are considered. By taking the transportation path into consideration, OT enjoys a dynamical formulation (3). This will be our main interest in the following discussion.

The variational formulation (3) introduces an infinite-dimensional Riemannian structure in the density space. We note that the infinite-dimensional Riemannian geometry is well-defined for smooth positive densities in a compact domain [24], and can be extended to more

general situations where the measure is only with finite second moments through tools in metric geometry [5]. Here, for better illustration, assume Ω is compact and consider the set of smooth and strictly positive densities

$$\mathcal{P}_+(\Omega) = \left\{ \rho \in C^\infty(\Omega) : \rho(x) > 0, \int_{\Omega} \rho(x) dx = 1 \right\} \subset \mathcal{P}_2(\Omega).$$

Denote by $\mathcal{F}(\Omega) := C^\infty(\Omega)$ the set of smooth real valued functions on Ω . The tangent space of $\mathcal{P}_+(\Omega)$ is given by

$$T_\rho \mathcal{P}_+(\Omega) = \left\{ \sigma \in \mathcal{F}(\Omega) : \int_{\Omega} \sigma(x) dx = 0 \right\}.$$

Given $\Phi \in \mathcal{F}(\Omega)$ and $\rho \in \mathcal{P}_+(\Omega)$, define

$$V_\Phi(x) := -\nabla \cdot (\rho(x) \nabla \Phi(x)) \in T_\rho \mathcal{P}_+(\Omega).$$

Since ρ is positive in a compact region Ω , the elliptic operator identifies the function Φ on Ω modulo additive constants with the tangent vector V_Φ in $\mathcal{P}_+(\Omega)$. This gives an isomorphism

$$\mathcal{F}(\Omega)/\mathbb{R} \rightarrow T_\rho \mathcal{P}_+(\Omega), \quad \Phi \mapsto V_\Phi.$$

Here we treat $T_\rho^* \mathcal{P}_+(\Omega) = \mathcal{F}(\Omega)/\mathbb{R}$ as the smooth cotangent space of $\mathcal{P}_+(\Omega)$. The above facts introduce the L^2 -Wasserstein metric tensor on the density space [24][22]:

Definition 1 (L^2 -Wasserstein metric tensor). *Define the inner product on the tangent space of positive densities $g_\rho: T_\rho \mathcal{P}_+(\Omega) \times T_\rho \mathcal{P}_+(\Omega) \rightarrow \mathbb{R}$ by*

$$g_\rho(\sigma_1, \sigma_2) = \int_{\Omega} \nabla \Phi_1(x) \cdot \nabla \Phi_2(x) \rho(x) dx,$$

where $\sigma_1 = V_{\Phi_1}$, $\sigma_2 = V_{\Phi_2}$ with $\Phi_1(x), \Phi_2(x) \in \mathcal{F}(\Omega)/\mathbb{R}$.

With the inner product specified above, the variational problem (3) becomes a geometric action energy in $(\mathcal{P}_+(\Omega), g_\rho)$. As in Riemannian geometry, the square of distance equals the energy of geodesics, i.e.

$$(W_2(\rho^0, \rho^1))^2 = \inf_{\Phi_t} \left\{ \int_0^1 g_{\rho_t}(V_{\Phi_t}, V_{\Phi_t}) dt : \partial_t \rho_t = V_{\Phi_t}, \rho(0, x) = \rho^0, \rho(1, x) = \rho^1 \right\}.$$

This is exactly the form in (3). In this sense, it explains that the dynamical formulation of OT exhibits the Riemannian structure in the density space. In [21], $(\mathcal{P}_+(\Omega), g_\rho)$ is named density manifold. More geometric studies are provided in [24][22].

3. WASSERSTEIN NATURAL GRADIENT

In this section, we study parametric statistical models, which are parameterized subsets of the probability space $\mathcal{P}_2(\Omega)$. We pull back the L^2 -Wasserstein metric tensor into the parameter space, turning it to be a Riemannian manifold. This consideration allows us to introduce the Riemannian (natural) gradient flow on parameter spaces, which further leads to a natural gradient descent method in the field of optimization. When densities lie in \mathbb{R} , we show that the metric tensor establishes an explicit formula. It acts as a positive definite and asymptotically-Hessian preconditioner for the Wasserstein metric related minimizations.

3.1. Wasserstein statistical manifold. We adopt the definition of a statistical model from [7] (Chapter 3.2). It is represented by a triple (Ω, Θ, ρ) , where $\Omega \subset \mathbb{R}^n$ is the continuous sample space, $\Theta \subset \mathbb{R}^d$ is the statistical parameter space, and ρ is the probability density on Ω parameterized by θ such that $\rho: \Theta \rightarrow \mathcal{P}_2(\Omega)$ and $\rho = \rho(\cdot, \theta)$. In Chapter 3.2 of the book [7] some differentiability conditions are posed on the map ρ between Banach manifolds Θ and $\mathcal{P}_2(\Omega)$. Here for simplicity we assume Ω is either compact or $\Omega = \mathbb{R}^n$, and each $\rho(x, \theta)$ is positive and smooth. The parameter space is a finite dimensional manifold with metric tensor denoted by $\langle \cdot, \cdot \rangle_\theta$. The Riemannian gradient of a function $\rho(\theta)$ on Θ is denoted by $\nabla_\theta \rho(\theta)$. We also use $\langle \cdot, \cdot \rangle$ to represent the Euclidean inner product in \mathbb{R}^d .

The Riemannian metric g_θ on Θ will be the formal pull-back of $g_{\rho(\cdot, \theta)}$ on $\mathcal{P}_2(\Omega)$. That is, for $\xi, \eta \in T_\theta \Theta$, we have

$$(4) \quad g_\theta(\xi, \eta) := g_{\rho(\cdot, \theta)}(d_\xi \rho(\theta), d_\eta \rho(\theta)),$$

where by definition of the Riemannian gradient, $d_\xi \rho(\theta) = \langle \nabla_\theta \rho(\cdot, \theta), \xi \rangle_\theta$, $d_\eta \rho(\theta) = \langle \nabla_\theta \rho(\cdot, \theta), \eta \rangle_\theta$, in which we use the notation that $d_\xi \rho(\theta)$ is the derivative of $\theta \rightarrow \rho(\theta)$ in the direction of ξ computed at θ . The relation (4) implicitly defines the tensor $g_{\rho(\cdot, \theta)}$. It involves the solution of elliptic equations and in order to make the formula explicit we make the following assumptions on the statistical model (Ω, Θ, ρ) :

Assumption 1. *For the statistical model (Ω, Θ, ρ) , one of the following two conditions are satisfied:*

- (1) *The sample space Ω is compact, and for each $\xi \in T_\theta(\Theta)$, the elliptic equation*

$$\begin{cases} -\nabla \cdot (\rho(x, \theta) \nabla \Phi(x)) = \langle \nabla_\theta \rho(x, \theta), \xi \rangle_\theta \\ \frac{\partial \Phi}{\partial n} |_{\partial \Omega} = 0 \end{cases}$$

has a smooth solution Φ satisfying

$$(5) \quad \int_{\Omega} \rho(x, \theta) \|\nabla \Phi(x)\|^2 dx < +\infty.$$

- (2) *The sample space $\Omega = \mathbb{R}^n$, and for each $\xi \in T_\theta(\Theta)$, the elliptic equation*

$$-\nabla \cdot (\rho(x, \theta) \nabla \Phi(x)) = \langle \nabla_\theta \rho(x, \theta), \xi \rangle_\theta$$

has a smooth solution Φ satisfying

$$(6) \quad \int_{\Omega} \rho(x, \theta) \|\nabla \Phi(x)\|^2 dx < +\infty \quad \text{and} \quad \int_{\Omega} \rho(x, \theta) |\Phi(x)|^2 dx < +\infty.$$

Note that Φ will be dependent on ξ , the tangent vector at play. For simplicity of notation we do not write this dependence explicitly, but this will be assumed wherever it appears throughout this paper. Assumption 1 guarantees the action of ‘‘pull-back’’ we describe above is well-defined. The condition (5) is for the boundness of the energy of the solution. The additional assumption in condition (6) is used to guarantee the vanishing of boundary terms when integrating by parts in the non-bounded domain, i.e., one has $-\int_{\mathbb{R}^n} \Phi_1 \nabla \cdot \rho \nabla \Phi_2 = \int_{\mathbb{R}^n} \rho \nabla \Phi_1 \cdot \nabla \Phi_2$ for any Φ_1, Φ_2 satisfying condition (6). This will be used in the proof of the uniqueness of the solution to the equation. We will use the uniqueness result in section 4.1 for Gaussian measures, to show the submanifold geometry is totally geodesic.

Proposition 1. *Under assumption 1, the solution Φ is unique modulo the addition of a spatially-constant function.*

Proof. It suffices to show the equation

$$(7) \quad \nabla \cdot (\rho(x, \theta) \nabla \Phi(x)) = 0$$

only has the trivial solution $\nabla \Phi = 0$ in the space described in assumption 1.

For case (1), we multiple Φ to (7) and integrate it in Ω . Integration by parts result in

$$\int_{\Omega} \rho(x, \theta) \|\nabla \Phi(x)\|^2 = 0$$

due to the zero flux condition. Hence $\nabla \Phi = 0$.

For case (2), we denote by $B_R(0)$ the ball in \mathbb{R}^n with center 0 and radius R . Multiply Φ to the equation and integrate in $B_R(0)$:

$$\int_{B_R(0)} \rho(x, \theta) \|\nabla \Phi(x)\|^2 = \int_{\partial B_R(0)} \rho(x, \theta) \Phi(x) (\nabla \Phi(x) \cdot n) dx.$$

By Cauchy-Schwarz inequality we can control the right hand side by

$$\left| \int_{\partial B_R(0)} \rho(x, \theta) \Phi(x) (\nabla \Phi(x) \cdot n) dx \right|^2 \leq \int_{\partial B_R(0)} \rho(x, \theta) \Phi(x)^2 dx \cdot \int_{\partial B_R(0)} \rho(x, \theta) \|\nabla \Phi(x)\|^2 dx.$$

However, due to (6), there exists a sequence $R_k, k \geq 1$, such that $R_{k+1} > R_k$, $\lim_{k \rightarrow +\infty} R_k = \infty$ and

$$\lim_{k \rightarrow +\infty} \int_{\partial B_{R_k}(0)} \rho(x, \theta) \Phi(x)^2 dx = \lim_{k \rightarrow +\infty} \int_{\partial B_{R_k}(0)} \rho(x, \theta) \|\nabla \Phi(x)\|^2 dx = 0.$$

Hence

$$\lim_{k \rightarrow +\infty} \left| \int_{\partial B_{R_k}(0)} \rho(x, \theta) \Phi(x) (\nabla \Phi(x) \cdot n) dx \right|^2 = 0,$$

which leads to

$$\int_{\mathbb{R}^n} \rho(x, \theta) \|\nabla \Phi(x)\|^2 = 0.$$

Thus $\nabla \Phi = 0$, which is the trivial solution. \square

Since we deal with positive ρ , the existence of solutions to case (1) in Assumption 1 is ensured by the theory of elliptic equations. For case (2), i.e. $\Omega = \mathbb{R}^n$, we show when ρ is Gaussian distribution in \mathbb{R}^d , the existence of solution Φ is guaranteed and exhibits explicit formulation in section 4.1. Although we only deal with compact Ω or the whole \mathbb{R}^n , the treatment to some other Ω , such as the half space of \mathbb{R}^n , is similar and omitted. Given these preparations, now we can explicitly write down the Wasserstein metric tensor in the parameter space. See the following Definition 2 and Proposition 2.

Definition 2 (L^2 -Wasserstein metric tensor in parameter space). *Under assumption 1, the inner product g_{θ} on $T_{\theta}(\Theta)$ is defined as*

$$g_{\theta}(\xi, \eta) = \int_{\Omega} \rho(x, \theta) \nabla \Phi_{\xi}(x) \cdot \nabla \Phi_{\eta}(x) dx,$$

where ξ, η are tangent vectors in $T_{\theta}(\Theta)$, Φ_{ξ} and Φ_{η} satisfy $\langle \nabla_{\theta} \rho(x, \theta), \xi \rangle_{\theta} = -\nabla \cdot (\rho \nabla \Phi_{\xi}(x))$ and $\langle \nabla_{\theta} \rho(x, \theta), \eta \rangle_{\theta} = -\nabla \cdot (\rho \nabla \Phi_{\eta}(x))$.

Generally, (Θ, g_{θ}) will be a Pseudo-Riemannian manifold. However, if the statistical model is non-degenerate, i.e., g_{θ} is positive definite on the tangent space $T_{\theta}(\Theta)$, then (Θ, g_{θ}) forms a Riemannian manifold. We call (Θ, g_{θ}) the Wasserstein statistical manifold.

Proposition 2. *The metric tensor can be written as*

$$(8) \quad g_\theta(\xi, \eta) = \xi^T G_W(\theta) \eta,$$

where $G_W(\theta) \in \mathbb{R}^{d \times d}$ is a positive definite matrix and can be represented by

$$G_W(\theta) = G_\theta^T A(\theta) G_\theta,$$

in which $A_{ij}(\theta) = \int_\Omega \partial_{\theta_i} \rho(x, \theta) (-\Delta_\theta)^{-1} \partial_{\theta_j} \rho(x, \theta) dx$ and $-\Delta_\theta = -\nabla \cdot (\rho(x, \theta) \nabla)$. The matrix G_θ associates with the original metric tensor in Θ such that $\langle \dot{\theta}_1, \dot{\theta}_2 \rangle_\theta = \dot{\theta}_1^T G_\theta \dot{\theta}_2$ for any $\dot{\theta}_1, \dot{\theta}_2 \in T_\theta(\Theta)$. If Θ is Euclidean space then $G_W(\theta) = A(\theta)$.

Proof. Write down the metric tensor

$$\begin{aligned} g_\theta(\xi, \eta) &= \int_\Omega \rho(x, \theta) \nabla \Phi_\xi(x) \cdot \nabla \Phi_\eta(x) dx \\ &\stackrel{a)}{=} \int_\Omega \langle \nabla_\theta \rho(x, \theta), \xi \rangle_\theta \cdot \Phi_\eta(x) dx \\ &= \int_\Omega \langle \nabla_\theta \rho(x, \theta), \xi \rangle_\theta (-\Delta_\theta)^{-1} \langle \nabla_\theta \rho(x, \theta), \eta \rangle_\theta dx \end{aligned}$$

where $a)$ is due to integration by parts. Comparing the above equation with (8) finishes the proof. \square

Given this $G_W(\theta)$, we derive the geodesic in this manifold and illustrate its connection to the geodesic in $\mathcal{P}_2(\Omega)$ as follows.

Proposition 3. *The geodesics in (Θ, g_θ) satisfies*

$$(9) \quad \begin{cases} \dot{\theta} - G_W(\theta)^{-1} S = 0 \\ \dot{S} + \frac{1}{2} S^T \frac{\partial}{\partial \theta} G_W(\theta)^{-1} S = 0 \end{cases}$$

Proof. In geometry, the square of geodesic distance d_W between $\rho(\cdot, \theta^0)$ and $\rho(\cdot, \theta^1)$ equals the energy functional:

$$(10) \quad d_W^2(\rho^0(\cdot, \theta), \rho^1(\cdot, \theta)) = \inf_{\theta(t) \in C^1(0,1)} \left\{ \int_0^1 \dot{\theta}(t)^T G_W(\theta) \dot{\theta}(t) dt : \theta(0) = \theta^0, \theta(1) = \theta^1 \right\}.$$

The minimizer of (10) satisfies the geodesic equation. Let us write down the Lagrangian $L(\dot{\theta}, \theta) = \frac{1}{2} \dot{\theta}^T G_W(\theta) \dot{\theta}$. The geodesic satisfies the Euler-Lagrange equation

$$\frac{d}{dt} \nabla_{\dot{\theta}} L(\dot{\theta}, \theta) = \nabla_\theta L(\dot{\theta}, \theta).$$

By the Legendre transformation,

$$H(S, \theta) = \sup_{\dot{\theta} \in T_\theta(\Theta)} S^T \dot{\theta} - L(\dot{\theta}, \theta).$$

Then $S = G_W(\theta) \dot{\theta}$ and $H(S, \theta) = \frac{1}{2} S^T G_W(\theta)^{-1} S$. Thus we derive the Hamilton's equations

$$\dot{\theta} = \partial_S H(\theta, S), \quad \dot{S} = -\partial_\theta H(\theta, S).$$

This recovers (9). \square

Remark 1. We recall the Wasserstein geodesic equation in $(\mathcal{P}_2(\Omega), W_2)$:

$$\begin{cases} \frac{\partial \rho(t, x)}{\partial t} + \nabla \cdot (\rho(t, x) \nabla \Phi(t, x)) = 0 \\ \frac{\partial \Phi(t, x)}{\partial t} + \frac{1}{2} (\nabla \Phi(t, x))^2 = 0 \end{cases}$$

The above PDE pair contains both continuity equation and Hamilton-Jacobi equation. Our equation (9) can be viewed as the continuity equation and Hamilton-Jacobi equation on the parameter space. The difference is that when restricted to a statistical model, the continuity equation and the associated Hamilton-Jacobi equation can only flow in the probability densities constrained in $\rho(\Theta)$.

Remark 2. If the optimal flow $\rho_t, 0 \leq t \leq 1$ in the continuity equation (3b) totally lies in the probability subspace parameterized by θ , then the two geodesic distances coincide:

$$d_W(\theta^0, \theta^1) = W_2(\rho^0(\cdot, \theta), \rho^1(\cdot, \theta)).$$

It is well known that the optimal transportation path between two Gaussian distributions will also be Gaussian distributions. Hence when $\rho(\cdot, \theta)$ are Gaussian measures, the above condition is satisfied. This means Gaussian is a totally geodesic submanifold. In general, d_W will be different from the L^2 Wasserstein metric. We will demonstrate this fact in our numerical examples.

3.2. Wasserstein natural gradient. Based on the Riemannian structure established in the previous section, we are able to introduce the gradient flow on the parameter space (Θ, g_θ) . Given an objective function $R(\theta)$, the associated gradient flow will be:

$$\frac{d\theta}{dt} = -\nabla_g R(\theta).$$

Here ∇_g is the Riemannian gradient operator satisfying

$$g_\theta(\nabla_g R(\theta), \xi) = \nabla_\theta R(\theta) \cdot \xi$$

for any tangent vector $\xi \in T_\theta \Theta$, where ∇_θ represents the Euclidean gradient operator.

Proposition 4. The gradient flow of function $R \in C^1(\Theta)$ in (Θ, g_θ) satisfies

$$(11) \quad \frac{d\theta}{dt} = -G_W(\theta)^{-1} \nabla_\theta R(\theta).$$

Proof. By the definition of gradient operator,

$$\nabla_g R(\theta)^T G_W(\theta) \xi = \nabla_\theta R(\theta) \cdot \xi,$$

for any ξ . Thus $\nabla_g R(\theta) = G_W(\theta)^{-1} \nabla_\theta R(\theta)$. \square

When $R(\theta) = R(\rho(\cdot, \theta))$, i.e. the function is implicitly determined by the density $\rho(\cdot, \theta)$, the Riemannian gradient can naturally reflect the change in the probability density domain. This will be expressed in our experiments by using Forward-Euler to solve the gradient flow numerically. The iteration writes

$$(12) \quad \theta^{n+1} = \theta^n - \tau G_W(\theta^n)^{-1} \nabla_\theta R(\rho(\cdot, \theta^n)).$$

This iteration of θ^{n+1} can also be understood as an approximate solution to the following problem:

$$\arg \min_{\theta} R(\rho(\cdot, \theta)) + \frac{d_W(\rho(\cdot, \theta^n), \rho(\cdot, \theta))^2}{2\tau}.$$

The approximation goes as follows. Note the Wasserstein metric tensor satisfies

$$d_W(\rho(\cdot, \theta + \Delta\theta), \rho(\cdot, \theta))^2 = \frac{1}{2}(\Delta\theta)^T G_W(\theta)(\Delta\theta) + o((\Delta\theta)^2) \quad \text{as } \Delta\theta \rightarrow 0,$$

and $R(\rho(\cdot, \theta + \Delta\theta)) = R(\rho(\cdot, \theta)) + \langle \nabla_{\theta} R(\rho(\cdot, \theta)), \Delta\theta \rangle + O((\Delta\theta)^2)$. Ignoring high-order items we obtain

$$\theta^{n+1} = \arg \min_{\theta} \langle \nabla_{\theta} R(\rho(\cdot, \theta^n)), \theta - \theta^n \rangle + \frac{(\theta - \theta^n)^T G_W(\theta^n)(\theta - \theta^n)}{2\tau}.$$

This recovers (12). It explains (12) is the steepest descent with respect to the change of probability distributions measured by W_2 .

In fact, (12) shares the same spirit in natural gradient [1] with respect to Fisher-Rao metric. To avoid ambiguity, we call it the Fisher-Rao natural gradient. It considers θ^{n+1} as an approximate solution of

$$\arg \min_{\theta} R(\rho(\cdot, \theta)) + \frac{D_{\text{KL}}(\rho(\cdot, \theta) \parallel \rho(\cdot, \theta^n))}{\tau},$$

where D_{KL} represents the Kullback-Leibler divergence, i.e. given two densities p, q on Ω , then

$$D_{\text{KL}}(p \parallel q) = \int_{\Omega} p(x) \log\left(\frac{p(x)}{q(x)}\right) dx.$$

In our case, we replace the KL divergence by the constrained Wasserstein metric. For this reason, we call (12) the Wasserstein natural gradient descent method.

3.3. 1D densities. In the following we concentrate on the one dimensional sample space, i.e. $\Omega = \mathbb{R}$. We show that g_{θ} exhibits an explicit formula. From it, we demonstrate that when the minimization is modeled by the Wasserstein distance, namely $R(\rho(\cdot, \theta))$ is related to W_2 , then $G_W(\theta)$ will approach the Hessian matrix of $R(\rho(\cdot, \theta))$ at the minimizer.

Proposition 5. *Suppose $\Omega = \mathbb{R}$, $\Theta = \mathbb{R}^d$ is the Euclidean space, and assumption 1 is satisfied, then the Riemannian inner product on the Wasserstein statistical manifold (Θ, g_{θ}) has explicit form*

$$(13) \quad G_W(\theta) = \int_{\mathbb{R}} \frac{1}{\rho(x, \theta)} (\nabla_{\theta} F(x, \theta))^T \nabla_{\theta} F(x, \theta) dx,$$

such that $g_{\theta}(\xi, \eta) = \langle \xi, G_W(\theta)\eta \rangle$.

Proof. When $\Omega = \mathbb{R}$, we have

$$g_{\theta}(\xi, \eta) = \int_{\mathbb{R}} \rho(x, \theta) \Phi'_{\xi}(x) \cdot \Phi'_{\eta}(x) dx,$$

where $\langle \nabla_{\theta} \rho(x, \theta), \xi \rangle = (\rho \Phi'_{\xi}(x))'$ and $\langle \nabla_{\theta} \rho(x, \theta), \eta \rangle = (\rho \Phi'_{\eta}(x))'$.

Integrating the two sides yields

$$\int_{-\infty}^y \langle \nabla_{\theta} \rho(x, \theta), \xi \rangle = \rho \Phi'_{\xi}(y).$$

Denote by $F(y) = \int_{-\infty}^y \rho(x)dx$ the cumulative distribution function of ρ , then

$$\Phi'_\xi(x) = \frac{1}{\rho(x, \theta)} \langle \nabla_\theta F(x, \theta), \xi \rangle,$$

and

$$g_\theta(\xi, \eta) = \int_{\mathbb{R}} \frac{1}{\rho(x, \theta)} \langle \nabla_\theta F(x, \theta), \xi \rangle \langle \nabla_\theta F(x, \theta), \eta \rangle dx.$$

This means $g_\theta(\xi, \eta) = \langle \xi, G_W(\theta)\eta \rangle$ and we obtain (13). Since assumption 1 is satisfied, this integral is well-defined. \square

Recall the Fisher-Rao metric tensor, also known as the Fisher information matrix:

$$\begin{aligned} G_F(\theta) &= \int_{\mathbb{R}} \rho(x, \theta) (\nabla_\theta \log \rho(x, \theta))^T \nabla_\theta \log \rho(x, \theta) dx \\ &= \int_{\mathbb{R}} \frac{1}{\rho(x, \theta)} (\nabla_\theta \rho(x, \theta))^T \nabla_\theta \rho(x, \theta) dx, \end{aligned}$$

where we use the fact $\nabla_\theta \log \rho(x, \theta) = \frac{1}{\rho(x, \theta)} \nabla_\theta \rho(x, \theta)$. Compared to the Fisher-Rao metric tensor, our Wasserstein metric tensor $G_W(\theta)$ only changes the density function in the integral to the corresponding cumulative distribution function. We note the condition that ρ is everywhere positive can be relaxed, for example, by assuming each component of $\nabla_\theta F(x, \theta)$, when viewed as a density in \mathbb{R} , is absolutely continuous with respect to $\rho(x, \theta)$. Then we can use the associated Radon-Nikodym derivative to define the integral. This treatment is similar to the one for Fisher-Rao metric tensor [7].

Now we turn to study the computational property of the natural gradient method with the Wasserstein metric. For standard Fisher-Rao natural gradient, it is known that when $R(\rho(\cdot, \theta)) = \text{KL}(\rho(\cdot, \theta), \rho(\cdot, \theta^*))$, then

$$\lim_{\theta \rightarrow \theta^*} G_F(\theta) = \nabla_{\theta}^2 R(\rho(\cdot, \theta^*)).$$

Hence, $G_F(\theta)$ will approach the Hessian of R at the minimizer. Regarding this, the Fisher-Rao natural gradient descent iteration

$$\theta^{n+1} = \theta^n - G_F(\theta^n)^{-1} \nabla_\theta R(\rho(\cdot, \theta))$$

will be asymptotically Newton method for KL divergence related minimization.

We would like to demonstrate a similar result for the Wasserstein natural gradient. In other words, we shall show the Wasserstein natural gradient will be asymptotically Newton method for the Wasserstein distance related minimization. To achieve this, we start by proving a detailed description of the Hessian matrix for the Wasserstein metric in Theorem 1. Throughout the following discussion, we use the notation $T'(x, \theta)$ to represent the derivative of T with respect to the x variable. We make the following assumption which is needed in the proof of Theorem 1 to interchange the differentiation and integration.

Assumption 2. For any $\theta_0 \in \Theta$, there exists a neighborhood $N(\theta_0) \subset \Theta$, such that

$$\begin{aligned} \int_{\Omega} \max_{\theta \in N(\theta_0)} \left| \frac{\partial^2 F(x, \theta)}{\partial \theta_i \partial \theta_j} \right| dx &< +\infty \\ \int_{\Omega} \max_{\theta \in N(\theta_0)} \left| \frac{\partial \rho(x, \theta)}{\partial \theta_i} \right| dx &< +\infty \\ \int_{\Omega} \max_{\theta \in N(\theta_0)} \frac{1}{\rho(x, \theta)} \left| \frac{\partial F(x, \theta)}{\partial \theta_i} \frac{\partial F(x, \theta)}{\partial \theta_j} \right| dx &< +\infty \end{aligned}$$

for each $1 \leq i, j \leq d$.

Theorem 1. Consider the statistical model (Ω, Θ, ρ) , in which $\rho(\cdot, \theta)$ is positive and Ω is a compact region in \mathbb{R} . Suppose Assumption 1 and 2 are satisfied and $T'(x, \theta)$ is uniformly bounded for all x when θ is fixed. If the objective function has the form

$$R(\rho(\cdot, \theta)) = \frac{1}{2} (W_2(\rho(\cdot, \theta), \rho^*))^2,$$

where ρ^* is the ground truth density, then

$$(14) \quad \nabla_{\theta}^2 R(\rho(\cdot, \theta)) = \int_{\Omega} (T(x, \theta) - x) \nabla_{\theta}^2 F(x, \theta) dx + \int_{\Omega} \frac{T'(x, \theta)}{\rho(x, \theta)} (\nabla_{\theta} F(x, \theta))^T \nabla_{\theta} F(x, \theta) dx,$$

in which $T(\cdot, \theta)$ is the optimal transport map between $\rho(\cdot, \theta)$ and ρ^* , the function $F(\cdot, \theta)$ is the cumulative distribution function of $\rho(\cdot, \theta)$.

Proof. We recall the three formulations of OT in section 2 and the following facts for 1D Wasserstein distance. They will be used in the proof.

(i) When $\Omega \subset \mathbb{R}$, the optimal map will have explicit formula, namely $T(x) = F_1^{-1}(F_0(x))$, where F_0, F_1 are cumulative distribution functions of ρ^0, ρ^1 respectively. Moreover, T satisfies

$$(15) \quad \rho^0(x) = \rho^1(T(x))T'(x).$$

(ii) The dual of linear programming (2) has the form

$$(16) \quad \max_{\phi} \int_{\Omega} \phi(x) \rho^0(x) dx + \int_{\Omega} \phi^c(x) \rho^1(x) dx,$$

in which ϕ and ϕ^c satisfy

$$\phi^c(y) = \inf_{x \in \Omega} \|x - y\|^2 - \phi(x).$$

(iii) We have the relation $\nabla \phi(x) = 2(x - T(x))$ for the optimal T and ϕ .

Using the above three facts, we have

$$R(\rho(\cdot, \theta)) = \frac{1}{2} (W_2(\rho(\cdot, \theta), \rho^*))^2 = \frac{1}{2} \int_{\Omega} |x - F_*^{-1}(F(x, \theta))|^2 \rho(x, \theta) dx,$$

where F_* is the cumulative distribution function of ρ^* . We first compute $\nabla_{\theta} R(\rho(\cdot, \theta))$. Fix θ and assume the dual maximum is achieved by ϕ^* and ϕ^{c*} :

$$R(\rho(\cdot, \theta)) = \frac{1}{2} \int_{\Omega} \phi(x)^* \rho(x, \theta) dx + \int_{\Omega} \phi^{c*}(x) \rho^*(x) dx.$$

Then, for any $\hat{\theta} \in \Theta$,

$$R(\rho(\cdot, \hat{\theta})) \leq \frac{1}{2} \int_{\Omega} \phi^*(x) \rho(x, \hat{\theta}) dx + \int_{\Omega} \phi^{c*}(x) \rho^*(x) dx,$$

and the equality holds when $\hat{\theta} = \theta$. Thus

$$\nabla_{\theta} R(\rho(\cdot, \theta)) = \nabla_{\theta} \frac{1}{2} \int_{\Omega} \phi^*(x) \rho(x, \theta) dx = \frac{1}{2} \int_{\Omega} \phi(x, \theta) \nabla_{\theta} \rho(x, \theta) dx,$$

in which integration and differentiation are interchangeable due to Lebesgue dominated convergence theorem and assumption 2. The function $\phi(x, \theta)$ is the Kantorovich potential associated with $\rho(x, \theta)$.

As is mentioned in (iii), $(\phi(x, \theta))' = 2(x - T(x, \theta))$, which leads to

$$\nabla_{\theta} R(\rho(\cdot, \theta)) = - \int_{\Omega} (x - T(x, \theta)) \nabla_{\theta} F(x, \theta) dx.$$

Differentiation with respect to θ and interchange integration and differentiation:

$$\nabla_{\theta}^2 R(\rho(\cdot, \theta)) = - \int_{\Omega} (x - T(x, \theta)) \nabla_{\theta}^2 F(x, \theta) dx + \int_{\Omega} (\nabla_{\theta} T(x, \theta))^T \nabla_{\theta} F(x, \theta) dx.$$

On the other hand, T satisfies the following equation as in (15):

$$(17) \quad \rho(x, \theta) = \rho^*(T(x, \theta)) T'(x, \theta).$$

Differentiating with respect to θ and noticing that the derivative of the right hand side has a compact form:

$$\nabla_{\theta} \rho(x, \theta) = (\rho^*(T(x, \theta)) \nabla_{\theta} T(x, \theta))',$$

and hence

$$\nabla_{\theta} T(x, \theta) = \frac{1}{\rho^*(T(x, \theta))} \int_{-\infty}^x \nabla_{\theta} \rho(y, \theta) dy = \frac{\nabla_{\theta} F(x, \theta)}{\rho^*(T(x, \theta))}.$$

Combining them together we obtain

$$(18) \quad \nabla_{\theta}^2 R(\rho(\cdot, \theta)) = \int_{\Omega} (T(x, \theta) - x) \nabla_{\theta}^2 F(x, \theta) dx + \int_{\Omega} \frac{1}{\rho^*(T(x, \theta))} (\nabla_{\theta} F(x, \theta))^T \nabla_{\theta} F(x, \theta) dx.$$

Substituting $\rho^*(T(x, \theta))$ by $\rho(x, \theta)$ and $T(x, \theta)$ based on (17), we obtain (14). Since assumption 1 and 2 are satisfied, and $T'(x, \theta)$ is uniformly bounded, the integral in (14) is well-defined. \square

Proposition 6. *Under the condition in Theorem 1, if the ground-truth density satisfies $\rho^* = \rho(\cdot, \theta^*)$ for some $\theta^* \in \Theta$, then*

$$\lim_{\theta \rightarrow \theta^*} \nabla_{\theta}^2 R(\rho(\cdot, \theta)) = \int_{\mathbb{R}} \frac{1}{\rho(x, \theta^*)} (\nabla_{\theta} F(x, \theta^*))^T \nabla_{\theta} F(x, \theta^*) dx = G_W(\theta^*).$$

Proof. When θ approaches θ^* , $T(x, \theta) - x$ will go to zero and T' will go to the identity. We finish the proof by the result in Theorem 1. \square

Proposition 6 explains that the Wasserstein natural gradient descent is asymptotically Newton method for the Wasserstein metric based minimization. The preconditioner $G_W(\theta)$ equals the Hessian matrix of R at the ground truth. Moreover, the formula (14) contains more information than proposition 6, and they can be used to find suitable Hessian-like preconditioners. For example, we can see the second term in (14) is different from $G_W(\theta)$ when $\theta \neq \theta^*$. It seems more accurate to use the term

$$(19) \quad \bar{G}_W(\theta) := \int \frac{T'(x, \theta)}{\rho(x, \theta)} (\nabla_{\theta} F(x, \theta))^T \nabla_{\theta} F(x, \theta) dx$$

to approximate the Hessian of $R(\rho(\cdot, \theta))$. When θ is near to θ^* , \bar{G}_W is likely to achieve slightly faster convergence to $\nabla_{\theta}^2 R(\rho(\cdot, \theta^*))$ than G_W . However, the use of $\bar{G}_W(\theta)$ could also have several difficulties. The presence of $T'(x, \theta)$ limits its application to general minimization problem in space Θ which does not involve a Wasserstein metric objective function. Also, the computation of $T'(x, \theta)$ might suffer from potential numerical instability, especially when T is not smooth, which is often the case when the ground-truth density is a sum of delta functions. This fact is also expressed in our numerical examples.

4. EXAMPLES

In this section, we consider several concrete statistical models. We compute the related metric tensor $G_W(\theta)$, either explicitly or numerically, and further calculate the geodesic in the Wasserstein statistical manifold. Moreover, we test the Wasserstein natural gradient descent method in the Wasserstein distance based inference and fitting problems [9]. We show that the preconditioner $G_W(\theta)$ exhibits promising performance, leading to stable and fast convergence of the iterations. We also compare our results with the Fisher-Rao natural gradient.

4.1. Gaussian measures. We consider the multivariate Gaussian densities $\mathcal{N}(\mu, \Sigma)$ in \mathbb{R}^n :

$$\rho(x, \theta) = \frac{1}{\sqrt{\det(2\pi\Sigma)}} \exp\left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu)\right),$$

where $\theta = (\mu, \Sigma) \in \Theta := \mathbb{R}^n \times \text{Sym}^+(n, \mathbb{R})$. Here $\text{Sym}^+(n, \mathbb{R})$ is the $n \times n$ positive symmetric matrix set, which is an open subset of the $n \times n$ symmetric matrix vector space $\text{Sym}(n, \mathbb{R})$. This implies that the tangent space at each Σ is $\text{Sym}(n, \mathbb{R})$ with metric tensor $\text{tr}(S_1 S_2)$ for tangent vectors S_1, S_2 , i.e., the tangent bundle is trivial. Due to this reason, we drop the θ -subscript in inner product define on θ , i.e, using $\langle \nabla_{\theta} \rho(x, \theta), \xi \rangle$ here instead of $\langle \nabla_{\theta} \rho(x, \theta), \xi \rangle_{\theta}$. We obtain an explicit formula for g_{θ} by using definition 2.

Proposition 7. *The Wasserstein metric tensor for the multivariate Gaussian model is*

$$g_{\theta}(\xi, \eta) = \langle \dot{\mu}_1, \dot{\mu}_2 \rangle + \text{tr}(S_1 \Sigma S_2),$$

for any $\xi, \eta \in T_{\theta}\Theta$. Here $\xi = (\dot{\mu}_1, \dot{\Sigma}_1)$ and $\eta = (\dot{\mu}_2, \dot{\Sigma}_2)$, in which $\dot{\mu}_1, \dot{\mu}_2 \in \mathbb{R}^n, \dot{\Sigma}_1, \dot{\Sigma}_2 \in \text{Sym}(n, \mathbb{R})$, and the symmetric matrix S_1, S_2 satisfy $\dot{\Sigma}_1 = S_1 \Sigma + \Sigma S_1, \dot{\Sigma}_2 = S_2 \Sigma + \Sigma S_2$.

Proof. First we examine the elliptic equation in definition 2 has the solution explained in Proposition 7. Write down the equation

$$\langle \nabla_{\theta} \rho(x, \theta), \xi \rangle = -\nabla \cdot (\rho(x, \theta) \nabla \Phi_{\xi}(x)).$$

By some computations we have

$$\begin{aligned} \langle \nabla_{\theta} \rho(x, \theta), \xi \rangle &= \langle \nabla_{\mu} \rho(x, \theta), \dot{\mu}_1 \rangle + \text{tr}(\nabla_{\Sigma} \rho(x, \theta) \dot{\Sigma}_1), \\ \langle \nabla_{\mu} \rho(x, \theta), \dot{\mu}_1 \rangle &= \dot{\mu}_1^T \Sigma^{-1}(x - \mu) \cdot \rho(x, \theta), \\ \text{tr}(\nabla_{\Sigma} \rho(x, \theta) \dot{\Sigma}_1) &= -\frac{1}{2} \left(\text{tr}(\Sigma^{-1} \dot{\Sigma}_1) - (x - \mu)^T \Sigma^{-1} \dot{\Sigma}_1 \Sigma^{-1}(x - \mu) \right) \cdot \rho(x, \theta), \\ -\nabla \cdot (\rho \nabla \Phi_{\xi}(x)) &= (\nabla \Phi_{\xi}(x)) \Sigma^{-1}(x - \mu) \cdot \rho(x, \theta) - \rho(x, \theta) \Delta \Phi_{\xi}(x). \end{aligned}$$

Observing these equations, we let $\nabla \Phi_{\xi}(x) = (S_1(x - \mu) + \dot{\mu}_1)^T$, and $\Delta \Phi_{\xi} = \text{tr}(S_1)$, where S_1 is a symmetric matrix to be determined. By comparison of the coefficients and the fact $\dot{\Sigma}_1$ is

symmetric, we obtain

$$\dot{\Sigma}_1 = S_1 \Sigma + \Sigma S_1.$$

Similarly $\nabla \Phi_\eta(x) = (S_2(x - \mu) + \dot{\mu}_2)^T$ and

$$\dot{\Sigma}_2 = S_2 \Sigma + \Sigma S_2.$$

Then

$$\begin{aligned} g_\theta(\xi, \eta) &= \int \rho(x, \theta) \nabla \Phi_\xi(x) \cdot \nabla \Phi_\eta(x) dx \\ &= \langle \dot{\mu}_1, \dot{\mu}_2 \rangle + \text{tr}(S_1 \Sigma S_2). \end{aligned}$$

It is easy to check Φ satisfies the condition (6) and the uniqueness is guaranteed. \square

For Gaussian distributions, the above derived metric tensor has already been revealed in [43][26][10]. Our calculation shows that it is a particular formulation of G_W . Hence, this demonstrates our defined submanifold geometry is totally geodesic in the Gaussian case. In the following, we turn to several one-dimensional non-Gaussian distributions and illustrate the metric tensor and the related geodesics, gradient flow numerically. We will see in general, the Wasserstein statistical manifold defined here will not be totally geodesic.

4.2. Mixture model. We consider a generalized version of the Gaussian distribution, namely the Gaussian mixture model. For simplicity we assume there are two components, i.e. $a\mathcal{N}(\mu_1, \sigma_1) + (1 - a)\mathcal{N}(\mu_2, \sigma_2)$ with density functions:

$$\rho(x, \theta) = \frac{a}{\sigma_1 \sqrt{2\pi}} e^{-\frac{(x-\mu_1)^2}{2\sigma_1^2}} + \frac{1-a}{\sigma_2 \sqrt{2\pi}} e^{-\frac{(x-\mu_2)^2}{2\sigma_2^2}},$$

where $\theta = (a, \mu_1, \sigma_1^2, \mu_2, \sigma_2^2)$ and $a \in [0, 1]$.

We first compute the geodesics of Wasserstein statistical manifold numerically. Set $\theta^0 = (0.3, -3, 0.5^2, -5, 0.4^2)$ and $\theta^1 = (0.6, 7, 0.4^2, 5, 0.3^2)$. Their density functions are shown in Figure 1.

To compute the geodesics in the Wasserstein statistical manifold, we solve the optimal control problem in (10) numerically via a direct method. Discretize the problem as

$$\min_{\theta_i, 1 \leq i \leq N-1} N \sum_{i=0}^{N-1} (\theta_{i+1} - \theta_i)^T G_W(\theta_i) (\theta_{i+1} - \theta_i),$$

where $\theta_0 = \theta^0, \theta_N = \theta^1$ and the discrete time step-size is $1/N$. We use coordinate descent method, i.e. applying gradient on each $\theta_i, 1 \leq i \leq N - 1$ alternatively till convergence.

The geodesics in the whole density space is obtained by first computing the optimal transportation map T , using the explicit formula in one dimension $Tx = F_1^{-1}(F_0(x))$. Here F_0, F_1 are the cumulative distribution functions of ρ^0 and ρ^1 . Then the geodesic probability densities satisfies $\rho(t, x) = (tT + (1-t)I)\# \rho_0(x)$ for $0 \leq t \leq 1$, where $\#$ is the push forward operator. The result is shown in Figure 2.

Figure 2 demonstrates that the geodesics in the whole density manifold does not lie in the sub-manifold formed by the mixture distribution, and thus the distance d_W differs from the L^2 Wasserstein metric. Hence the optimal transport in the whole density space destroys the geometric shape during its path, which is not a desired property when we perform transportation.




FIGURE 1. Densities of Gaussian mixture distribution




FIGURE 2. Geodesic of Gaussian mixtures; left: in the Wasserstein statistical manifold; right: in the whole density space

Next, we test the Wasserstein natural gradient method in optimization. Consider the Gaussian mixture fitting problem: given N data points $\{x_i\}_{i=1}^N$ obeying the distribution $\rho(x; \theta^1)$ (unknown), we want to infer θ^1 by using these data points, which leads to a minimization as:

$$\min_{\theta} d \left(\rho(\cdot; \theta), \frac{1}{N} \sum_{i=1}^N \delta_{x_i}(\cdot) \right),$$

where d are certain distance functions on probability space. If we set d to be KL divergence, then the problem will correspond to the maximum likelihood estimate. Since $\rho(x; \theta^1)$ has small compact support, using KL divergence is risky and needs very good initial guess and careful optimization. Here we use the L^2 Wasserstein metric instead and set $N = 10^3$. We truncate the distribution in $[-r, r]$ for numerical computation. We choose $r = 15$. The Wasserstein

metric is effectively computed by the explicit formula

$$\frac{1}{2} \left(W_2(\rho(\cdot; \theta), \frac{1}{n} \sum_{i=1}^N \delta_{x_i}) \right)^2 = \frac{1}{2} \int_{-r}^r |x - T(x)|^2 \rho(x; \theta) dx,$$

where $T(x) = F_{em}^{-1}(F(x))$ and F_{em} is the cumulative distribution function of the empirical distribution. The gradient with respect to θ can be computed through

$$\nabla_{\theta}(\frac{1}{2}W^2) = \int_{-r}^r \phi(x) \nabla_{\theta} \rho(x; \theta) dx,$$

where $\phi(x) = \int_{-r}^x (y - T(y)) dy$ is the Kantorovich potential. The derivative $\nabla_{\theta} \rho(x; \theta)$ is obtained by numerical differentiation. We perform the following five iterative algorithms to solve the optimization problem:

Gradient descent (GD) : $\theta_{n+1} = \theta_n - \tau \nabla_{\theta}(\frac{1}{2}W^2)|_{\theta_n}$

GD with diag-preconditioning : $\theta_{n+1} = \theta_n - \tau P^{-1} \nabla_{\theta}(\frac{1}{2}W^2)|_{\theta_n}$

(A) Wasserstein GD : $\theta_{n+1} = \theta_n - \tau G_W(\theta_n)^{-1} \nabla_{\theta}(\frac{1}{2}W^2)|_{\theta_n}$

Modified Wasserstein GD : $\theta_{n+1} = \theta_n - \tau (\tilde{G}_W(\theta_n))^{-1} \nabla_{\theta}(\frac{1}{2}W^2)|_{\theta_n}$

Fisher-Rao GD : $\theta_{n+1} = \theta_n - \tau G_F(\theta_n)^{-1} \nabla_{\theta}(\frac{1}{2}W^2)|_{\theta_n}$

We consider the diagonal preconditioning because the scale of parameter a is very different from $\mu_i, \sigma_i, 1 \leq i \leq 2$. The diagonal matrix P is set to be $\text{diag}(40, 1, 1, 1, 1)$. We choose the initial step-size $\tau = 1$ with line search such that the objective value is always decreasing. The initial guess $\theta = \theta^0$. Figure 3 shows the experimental results.




FIGURE 3. objective value

From the figure, it is seen that the Euclidean gradient descent fails to converge. We observed that during iterations the parameter a goes very fast to 1 and then stop updating anymore. This is due to the ill-conditioned nature of the problem, in the sense that the scale of parameter differs drastically. If we use the diagonal matrix P to perform preconditioning, then it converges after approximately 70 steps. If we use Wasserstein gradient descent, then the iterations converge very efficiently, taking less than 10 steps. This demonstrates that $G_W(\theta)$ is well suited for the Wasserstein metric minimization problems, exhibiting very stable behavior. It can automatically detect the parameter scale and the underlying geometry. As a comparison, Fisher-Rao gradient descent fails, which implies $G_F(\theta)$ is not suitable for this Wasserstein metric modeled minimization.

The modified Wasserstein gradient descent does not converge because of the numerical instability of T' and further $\bar{G}_W(\theta)$ in the computation. In the next example with lower dimensional parameter space, however, we will see that $\bar{G}_W(\theta)$ performs better than $G_W(\theta)$. This implies $\bar{G}_W(\theta)$, if computed accurately, might achieve smaller approximation error to the Hessian matrix. Nevertheless, the difference is very slight, and since the matrix $\bar{G}_W(\theta)$ can only be applied to the Wasserstein modeled problem, we tend to believe that $G_W(\theta)$ is a better preconditioner.

4.3. Gamma distribution. Consider gamma distribution $\Gamma(\alpha, \beta)$, which has the probability density function

$$\rho(x; \alpha, \beta) = \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)}.$$

Set $\theta = (\alpha, \beta)$ and $\theta^0 = (2, 3), \theta^1 = (20, 2)$. Their density functions are shown in Figure 4.




FIGURE 4. Gamma density functions

We compute the related geodesics, in the Wasserstein statistical manifold and the whole density space respectively. The results are presented in Figure 5. We can see that these two do not differ very much. This means the optimal transport in the whole space could nearly keep the gamma distribution shape along with transportation.




FIGURE 5. Geodesic of Gamma distribution; left: in the Wasserstein statistical manifold; right: in the whole density space

Then, we consider the gamma distribution fitting problem. The model is similar to the one in the mixture examples, except that the parameterized family changes. The minimization problem is:

$$\min_{\theta} \frac{1}{2} \left(W_2(\rho(\cdot; \theta), \frac{1}{N} \sum_{i=1}^N \delta_{x_i}) \right)^2,$$

where $x_i \sim \rho(\cdot, \theta^1)$ and we set $N = 10^3$. The initial guess is $\theta = \theta^0$. Convergence results are presented in Figure 6.




FIGURE 6. objective value

The figure shows that the Euclidean gradient descent method takes a very long time to reach convergence, while Wasserstein GD and its modified version needs less than 10 steps, with the Fisher-Rao GD taking around 50 steps. This comparison demonstrates the efficiency of Wasserstein natural gradient in this Wasserstein metric modeled optimization problems.

As is mentioned in the previous example, the difference between using G_W and \bar{G}_W is very small. Since \bar{G}_W fails in the mixture example, we conclude that G_W , the Wasserstein gradient descent, will be a more stable choice for preconditioning.

5. NUMERICAL COMPARISON BETWEEN FISHER-RAO AND WASSERSTEIN NATURAL GRADIENT

In the previous section, we provide several statistical estimation examples, in which the Wasserstein metric is served as the loss function, to investigate the performance of different kinds of gradient descent algorithms. Among the five iterative schemes in (A), we observe that the Wasserstein natural gradient outperforms the others. This demonstrates our theoretical arguments before. However, these discussions are all constrained to the Wasserstein metric loss function case and limited to the occasions that the ground truth density lies in the parametric family, i.e., the parametrization is well-specified, which may not be true in practice. Thus, it would be interesting to see the performance of these algorithms when different loss functions are used (say, KL divergence or Maximum Likelihood Estimation) and when the true density lies outside the parametric family. This is the main target of the present section.

Here, we consider three different iterative rules, namely Gradient Descent (GD), Wasserstein GD, and Fisher-Rao GD. In general, the loss is denoted by $d(\rho(\cdot, \theta), \rho^*(\cdot))$. Then, the iterative schemes write

$$\begin{aligned} \text{Gradient descent (GD)} : \quad & \theta_{n+1} = \theta_n - \tau \nabla_{\theta} d(\rho(\cdot, \theta), \rho^*(\cdot))|_{\theta_n} \\ \text{(B) Wasserstein GD} : \quad & \theta_{n+1} = \theta_n - \tau G_W(\theta_n)^{-1} \nabla_{\theta} d(\rho(\cdot, \theta), \rho^*(\cdot))|_{\theta_n} \\ \text{Fisher-Rao GD} : \quad & \theta_{n+1} = \theta_n - \tau G_F(\theta_n)^{-1} \nabla_{\theta} d(\rho(\cdot, \theta), \rho^*(\cdot))|_{\theta_n} \end{aligned}$$

We omit the GD with diag-preconditioning in (A) because it would be very difficult to come up with a general diagonal preconditioning rule for arbitrary problems. The modified Wasserstein GD is also omitted because in previous examples it behaves less stable compared to the unmodified version and more crucially, it is only applicable to the Wasserstein metric loss based problems.

The experiment is designed as follows: we set d to be the Wasserstein metric and KL divergence (or equivalently, Maximum Likelihood Estimate), respectively. The parametric density $\rho(\cdot, \theta)$ is the two-component Gaussian mixture

$$\rho(x, \theta) = \frac{1}{1 + \exp(a)} \frac{1}{\sigma_1 \sqrt{2\pi}} e^{-\frac{(x-\mu_1)^2}{2\sigma_1^2}} + \frac{1}{1 + \exp(-a)} \frac{1}{\sigma_2 \sqrt{2\pi}} e^{-\frac{(x-\mu_2)^2}{2\sigma_2^2}},$$

where $\theta = (a, \mu_1, \sigma_1^2, \mu_2, \sigma_2^2)$. The target $\rho^*(\cdot)$ is the empirical distribution of data samples generated from two-component Gaussian mixture distribution (well-specified case) or Laplace distribution (misspecified case, such that ρ^* is not in the parametric family) with density

$$\rho^*(x, \mu, b) = \frac{1}{2b} \exp\left(-\frac{|x - \mu|}{b}\right).$$

The number of data samples is determined later. We randomly pick some values of the parameters to generate $\rho(x, \theta^0)$ and ρ^* , where θ^0 is the initial parameter value. We draw a certain number of samples (determined later) to form the empirical data and construct the loss function, either the Wasserstein metric or Maximum Likelihood Estimate. Then, we

apply the three iterative schemes in (B) according to the following rule: the step size τ is chosen to monotonically decrease the objective value along with the iteration (by line search: $\tau \rightarrow \tau/2$ with initial stepsize 1). The iteration is stopped if

- $\|\nabla_{\theta} d(\rho(\cdot, \theta), \rho^*(\cdot))\| \leq \epsilon$ where we set $\epsilon = 10^{-1}$; or
- stepsize is smaller than δ but line search still fails, where we set $\delta = 10^{-4}$; or
- iteration step is larger than N ; where we set $N = 200$

In the stopping rule, the first case corresponds to the iteration converges. The second case may happen when the metric tensor is poorly conditioned, or the iteration jumps out of the computation region. The third case is for computational efficiency consideration, to prevent too many iterations. We collect the statistical behaviors of these iterations for many differently generated $\rho(x, \theta)$ and ρ^* . We classify the results based on their final objective values and calculate the mean value and standard variance. We also record the number of iterations taken in the process when different methods are used.

We present the experimental results in the following two subsection, from well-specified case to misspecified case.

5.1. Well-specified case. In the well-specified case, we set $\rho(\cdot, \theta_0)$ (initial point of the iteration) and ρ^* both from two-component mixture family, parametrized in the following way:

$$\rho(x, \theta) = \frac{1}{1 + \exp(a)} \frac{1}{\sigma_1 \sqrt{2\pi}} e^{-\frac{(x-\mu_1)^2}{2\sigma_1^2}} + \frac{1}{1 + \exp(-a)} \frac{1}{\sigma_2 \sqrt{2\pi}} e^{-\frac{(x-\mu_2)^2}{2\sigma_2^2}},$$

where $\theta = (a, \mu_1, \sigma_1^2, \mu_2, \sigma_2^2)$.

We choose a uniformly distributed in $[-2, 2]$, μ_1, μ_2 uniformly in $[-10, 10]$, σ_1^2, σ_2^2 uniformly in $[1, 11]$ to generate $\rho(\cdot, \theta^0)$ and ρ^* . We collect 200 samples from ρ^* , whose empirical distribution will be the target data we use in optimization. The loss function is chosen as the Wasserstein metric and Maximum Likelihood Estimate, respectively. Given the loss function, we perform the three iterative methods. 100 trials of different $\rho(\cdot, \theta^0)$ and ρ^* are tested. We output the mean / standard derivation (std) of the final objective values and number of total iteration steps in the 100 trials. We also plot the histogram of the final objective (obj) value for the three methods for better comparisons. Table 1 and Figure 8 are for the Wasserstein metric loss, while Table 2 and Figure 8 are for the Maximum Likelihood Estimation case.

	Wasserstein GD	Fisher-Rao GD	GD
obj mean	0.0378	0.0489	0.1543
obj std	0.0208	0.0304	0.4251
iterations mean	5.0842	9.6947	47.6105
iterations std	1.7052	5.5373	57.4925

TABLE 1. Wasserstein metric loss function: well-specified case

From the Tables 1, 2 and Figures 7, 8, we observe that Wasserstein GD behaves well when the Wasserstein metric is used as loss function, while it may perform worse than GD if it is applied to Maximum Likelihood Estimate case. Similarly, Fisher-Rao GD is very efficient in the Maximum Likelihood Estimate case but is less effective than Wasserstein GD when the Wasserstein metric loss function is used. This may be explained by the argument that the

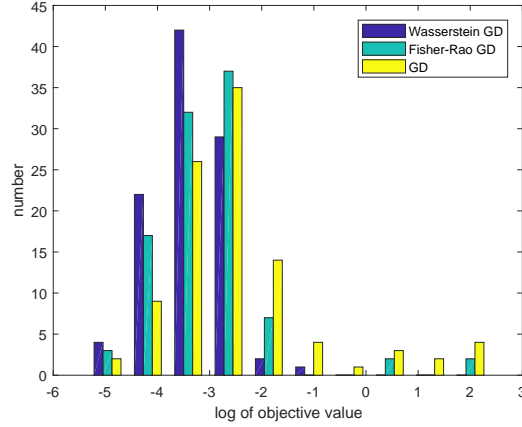


FIGURE 7. Histogram of obj (Wasserstein metric loss function: well-specified case)

	Wasserstein GD	Fisher-Rao GD	GD
obj mean	2.5613	2.4265	2.4119
obj std	0.4290	0.3548	0.2686
iterations mean	102.1684	3.9053	15.4632
iterations std	84.5325	0.9461	6.2481

TABLE 2. Maximum Likelihood Estimate: well-specified case




FIGURE 8. Histogram of obj (Maximum Likelihood Estimate: well-specified case)

natural gradient behaves as asymptotic Newton algorithm. Interestingly, in our experiments, the Fisher-Rao GD seems to be more stable and robust (better than GD) when used in the Wasserstein metric loss case, compared to the performance of Wasserstein GD applied to Maximum Likelihood Estimation.

5.2. **Misspecified case.** In the misspecified case, the target distribution is Laplace distribution

$$\rho^*(x, \mu, b) = \frac{1}{2b} \exp\left(-\frac{|x - \mu|}{b}\right),$$

while the parametric family is still Gaussian mixtures. We choose μ uniformly distributed from $[-10, 10]$, and b uniformly from $[1, 4]$ to generate the ground truth distributions and draw 200 samples from it, which will be the empirical data. We conduct 100 runs of experiments using different generated $\rho(x, \theta), \rho^*$. The setting of the experiment is the same as the well-specified case. Table 3 and Figure 10 are the results for the Wasserstein metric loss, while Table 4 and Figure 10 are for the Maximum Likelihood Estimate case.

	Wasserstein GD	Fisher-Rao GD	GD
obj mean	0.2490	0.3238	0.4986
obj std	0.2119	0.2493	0.3860
iterations mean	6.2947	10.0211	56.3579
iterations std	1.7127	4.4840	55.3561

TABLE 3. Wasserstein metric loss function: misspecified case




FIGURE 9. Histogram of obj (Wasserstein metric loss function: misspecified case)

	Wasserstein GD	Fisher-Rao GD	GD
obj mean	2.8450	2.6446	2.6426
obj std	0.5207	0.3366	0.3792
iterations mean	85.5895	4.2421	16.6211
iterations std	80.7177	1.1552	5.6251

TABLE 4. Maximum Likelihood Estimate: misspecified case

From these tables and figures, we observe that the results are very similar to the well-specified case. Wasserstein GD is efficient for the Wasserstein metric loss function, while Fisher-Rao GD performs very well for Maximum Likelihood Estimation. These results still hold for the situation that the ground truth distribution ρ^* is not in the parametric family. This implies the robustness of the natural gradient method. Though here we only conduct

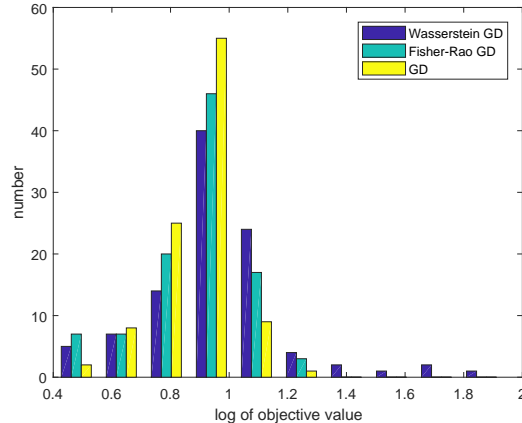


FIGURE 10. histogram of obj (Maximum Likelihood Estimate: misspecified case)

experiments on Gaussian mixtures and Laplace distribution, we believe our results reveal that when using the natural gradient concept for computations, one needs to choose a suitable geometry to design the natural gradient. This geometry should explore the structures of objective function and the optimization problem in general.

6. DISCUSSION

To summarize, we introduce the Wasserstein statistical manifold for parametric models with continuous sample space. The metric tensor is derived by pulling back the L^2 -Wasserstein metric tensor in density space to parameter spaces. Given this Riemannian structure, the Wasserstein natural gradient is then proposed. In one-dimensional sample space, we obtain an explicit formula for this metric tensor, and from it, we show that the Wasserstein natural gradient descent method achieves asymptotically Newton method for Wasserstein metric modeled minimizations. Our numerical examples justify these arguments.

One potential future direction is using Theorem 1 to design various efficient algorithms for solving Wasserstein metric modeled problems. The Wasserstein gradient descent only takes the asymptotic behavior into consideration, and we think a careful investigation of the structure (14) will lead to better non-asymptotic results. Moreover, generalizing (14) to higher dimensions also remains a challenging and interesting issue. We are working on designing an efficient computational method for obtaining $G_W(\theta)$ and hope to report it in subsequent papers.

Analytically, the treatment of the Wasserstein statistical manifold could be generalized. This paper takes an initial step in introducing Wasserstein geometry to parametric models. More analysis on the solution of the elliptic equation and its regularity will be conducted.

Further, we believe ideas and studies from information geometry could lead to natural extensions in Wasserstein statistical manifold. The Wasserstein distance has shown its effectiveness in illustrating and measuring low dimensional supported densities in high dimensional space, which is often the target of many machine learning problems. We are interested in geometric properties of Wasserstein metric in these models, and we will continue to work on it.

Acknowledgments: This research is partially supported by AFOSR MURI proposal number 18RT0073. The research of Yifan Chen is partly supported by Tsinghua undergraduate Xuetao Mathematics Program and Caltech graduate Kortchak Scholarship. The authors thank Prof. Shui-Nee Chow for his farseeing viewpoints on the related topics, and we acknowledge many fruitful discussions with Prof. Wilfrid Gangbo and Prof. Wotao Yin. We gratefully thank Prof. Guido Montúfar for many valuable comments regarding the experimental design part about an earlier version of this manuscript.

REFERENCES

- [1] S. Amari. Natural Gradient Works Efficiently in Learning. *Neural Computation*, 10(2):251–276, 1998.
- [2] S. Amari. *Information Geometry and Its Applications*. Volume 194. Springer, 2016.
- [3] S. Amari and A. Cichocki. Adaptive blind signal processing-neural network approaches. *Proceedings of the IEEE*, 86(10):2026–2048, 1998.
- [4] S. Amari, R. Karakida, and M. Oizumi. Information Geometry Connecting Wasserstein Distance and Kullback-Leibler Divergence via the Entropy-Relaxed Transportation Problem. *Information Geometry*, 1(1),13–37, 2018.
- [5] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Birkhäuser Basel, Basel, 2005.
- [6] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *arXiv:1701.07875 [cs, stat]*, 2017.
- [7] N. Ay, J. Jost, H. V. Lê, and L. J. Schwachhöfer. *Information Geometry*. Ergebnisse der Mathematik und ihrer Grenzgebiete A series of modern surveys in mathematics. Folge, volume 64. Springer, Cham, 2017.
- [8] J. D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.
- [9] E. Bernton, P. E. Jacob, M. Gerber, and C. P. Robert. Inference in generative models using the wasserstein distance. *arXiv:1701.05146 [math, stat]*, 2017.
- [10] R. Bhatia, T. Jian, Y. Lim. On the Bures-Wasserstein distance between positive definite matrices. *Expositiones Mathematicae*, 2018.
- [11] E. A. Carlen and W. Gangbo. Constrained Steepest Descent in the 2-Wasserstein Metric. *Annals of Mathematics*, 157(3):807–846, 2003.
- [12] F. P. Carli, L. Ning, and T. T. Georgiou. Convex Clustering via Optimal Mass Transport. *arXiv:1307.5459 [cs]*, 2013.
- [13] J. Chen, Y. Chen, H. Wu, and D. Yang. The quadratic Wasserstein metric for earthquake location. *Journal of Computational Physics*, 373:188–209, 2018.
- [14] Y. Chen, T. T. Georgiou, and A. Tannenbaum. Optimal transport for Gaussian mixture models. *IEEE Access*, 7:6269–6278, 2019.
- [15] N. N. Chentsov. *Statistical decision rules and optimal inference*. American Mathematical Society, Providence, R.I., 1982.
- [16] S. N. Chow, W. Li, J. Lu, and H. Zhou. Population games and Discrete optimal transport. *Journal of Nonlinear Science*, 29(3):871–896, 2019.
- [17] P. Degond, J. G. Liu, and C. Ringhofer. Large-Scale Dynamics of Mean-Field Games Driven by Local Nash Equilibria. *Journal of Nonlinear Science*, 24(1):93–115, 2014.
- [18] B. Engquist and B. D. Froese. Application of the Wasserstein metric to seismic signals. *Communications in Mathematical Sciences*, 12(5):979–988, 2014.
- [19] B. Engquist, B. D. Froese, and Y. Yang. Optimal transport for seismic full waveform inversion. *Communications in Mathematical Sciences*, 14(8):2309–2330, 2016.
- [20] C. Frogner, C. Zhang, H. Mobahi, M. Araya-Polo, and T. Poggio. Learning with a Wasserstein Loss. In *Advances in Neural Information Processing Systems*, pages 2053–2061, 2015.
- [21] J. D. Lafferty. The density manifold and configuration space quantization. *Transactions of the American Mathematical Society*, 305(2):699–741, 1988.
- [22] W. Li. Geometry of probability simplex via optimal transport. *arXiv:1803.06360 [math]*, 2018.
- [23] W. Li and G. Montufar. Natural gradient via optimal transport. *Information Geometry*, 1(2): 181-214, 2018.
- [24] J. Lott. Some Geometric Calculations on Wasserstein Space. *Communications in Mathematical Physics*, 277(2):423–437, 2007.

- [25] J. Lott and C. Villani. Ricci curvature for metric-measure spaces via optimal transport. *Annals of Mathematics*, 169(3):903–991, 2009.
- [26] L. Malagò, L. Montrucchio, and G. Pistone. Wasserstein Riemannian Geometry of Positive Definite Matrices. *arXiv:1801.09269 [math, stat]*, 2018.
- [27] L. Malagò and G. Pistone. Natural Gradient Flow in the Mixture Geometry of a Discrete Exponential Family. *Entropy*, 17(6):4215–4254, 2015.
- [28] L. Malagò, M. Matteucci, and G. Pistone. Natural gradient, fitness modelling and model selection: A unifying perspective. *2013 IEEE Congress on Evolutionary Computation*, Cancun, pages 486–493, 2013.
- [29] L. Malagò and M. Matteucci. Robust Estimation of Natural Gradient in Optimization by Regularized Linear Regression. *Geometric Science of Information*, Springer Berlin Heidelberg, pages 861–867, 2013.
- [30] J. Martens. New insights and perspectives on the natural gradient method. *arXiv:1412.1193 [cs, stat]*, 2014.
- [31] G. Marti, S. Andler, F. Nielsen, and P. Donnat. Optimal transport vs. Fisher-Rao distance between copulas for clustering multivariate time series. *2016 IEEE Statistical Signal Processing Workshop*, pages 1–5, 2016.
- [32] L. Métivier, R. Brossier, Q. Mérigot, E. Oudet and J. Virieux. Measuring the misfit between seismograms using an optimal transport distance: application to full waveform inversion *Geophysical Supplements to the Monthly Notices of the Royal Astronomical Society*, 205(1): 345–377, 2016.
- [33] L. Métivier, R. Brossier, Q. Mérigot, E. Oudet and J. Virieux. An optimal transport approach for seismic tomography: application to 3D full waveform inversion. *Inverse Problems*, 32(11): 115008, 2016.
- [34] K. Modin. Geometry of Matrix Decompositions Seen Through Optimal Transport and Information Geometry. *Journal of Geometric Mechanics*, 9(3):335–390, 2017.
- [35] G. Montavon, K. R. Müller, and M. Cuturi. Wasserstein Training of Restricted Boltzmann Machines. In *Advances in Neural Information Processing Systems 29*, pages 3718–3726, 2016.
- [36] Y. Ollivier. Online natural gradient as a Kalman filter. *Electronic Journal of Statistics*, 12(2): 2930–2961, 2018.
- [37] Y. Ollivier. The Extended Kalman Filter is a Natural Gradient Descent in Trajectory Space. *arXiv:1901.00696*, 2019.
- [38] Y. Ollivier. True Asymptotic Natural Gradient Optimization. *arXiv:1712.08449*, 2017.
- [39] Y. Ollivier, L. Arnold, A. Auger and N. Hansen. Information-Geometric Optimization Algorithms: A Unifying Picture via Invariance Principles. *Journal of Machine Learning Research*, 18(18):1–65, 2017.
- [40] F. Otto. The geometry of dissipative evolution equations the porous medium equation. *Communications in Partial Differential Equations*, 26(1-2):101–174, 2001.
- [41] G. Peyré and M. Cuturi. Computational Optimal Transport. *arXiv:1803.00567 [stat]*, 2018.
- [42] A. Sanctis and S. Gattone. A Comparison between Wasserstein Distance and a Distance Induced by Fisher-Rao Metric in Complex Shapes Clustering. *Multidisciplinary Digital Publishing Institute Proceedings*, 2(4):163, 2017.
- [43] A. Takatsu. Wasserstein geometry of Gaussian measures. *Osaka Journal of Mathematics*, 48(4):1005–1026, 2011.
- [44] C. Villani. *Optimal Transport: Old and New*. Volume 338. Springer Science & Business Media, 2008.
- [45] T. L. Wong. Logarithmic divergences from optimal transport and Rényi geometry. *Information Geometry*, 1(1): 39–78, 2018.

DEPARTMENT OF COMPUTING + MATHEMATICAL SCIENCES, CALIFORNIA INSTITUTE OF TECHNOLOGY, PASADENA, CALIFORNIA 91106, EMAIL: YIFANC@CALTECH.EDU

DEPARTMENT OF MATHEMATICS, UCLA, LOS ANGELES, CA 90095 USA, EMAIL: WCLI@MATH.UCLA.EDU