RICCI CURVATURE FOR PARAMETRIC STATISTICS VIA OPTIMAL TRANSPORT

WUCHEN LI AND GUIDO MONTÚFAR

ABSTRACT. We elaborate the notion of a Ricci curvature lower bound for parametrized statistical models. Following the seminal ideas of Lott-Strum-Villani, we define this notion based on the geodesic convexity of the Kullback-Leibler divergence in a Wasserstein statistical manifold, that is, a manifold of probability distributions endowed with a Wasserstein metric tensor structure. Within these definitions, the Ricci curvature is related to both, information geometry and Wasserstein geometry. These definitions allow us to formulate bounds on the convergence rate of Wasserstein gradient flows and information functional inequalities in parameter space. We discuss examples of Ricci curvature lower bounds and convergence rates in exponential family models.

1. INTRODUCTION

The Ricci curvature lower bound on sample space plays a crucial role in various fields, including heat semi-groups [3] and differential geometry (Brunn-Minkowski inequality) [24]. In particular, it provides sharp bounds for convergence rates of diffusion processes [3] and functional inequalities [22]. In recent years, optimal transport contributes a viewpoint that connects Ricci curvature and information functionals. In this study, optimal transport, in particular the L^2 -Wasserstein metric, introduces a Riemannian structure in probability density space, named density manifold [14]. The Ricci curvature lower bound in sample space is equivalent to the geodesic convexity of the Kullback-Leibler divergence in the density manifold¹. Following this angle, Lott-Strum-Villani [19, 23] define the Ricci curvature on non-smooth metric sample spaces, and Erbar-Maas [9] introduce it on discrete sample spaces.

In statistics and machine learning, we often are interested in constructing, or selecting, a density that models the behavior of some observed data, according to some quality criterion. Very often we restrict the search to a subset of densities, as this allows us to handle large state spaces and also to incorporate prior knowledge into our search. Parametrized statistical models are a ubiquitous and powerful approach. In this paper, we develop the theory of Ricci curvature lower bounds for this situation. The Ricci curvature lower bound governs the dissipation rates of the cross entropy. In the context of learning, this corresponds to the rates of convergence of gradient descent methods for minimizing the Kullback-Leibler (KL) divergence and computing information projections.

Key words and phrases. Ricci curvature; information projection; Wasserstein statistical manifold; Fokker-Planck equation on parameter space; machine learning.

¹Geodesic convexity is a synthetic definition. If a function f on manifold (M, g) is second differentiable, then f is λ -geodesic convex whenever $\operatorname{Hess}_M f \succeq \lambda g$.

The Wasserstein metric tensor of a statistical manifold (a parametrized set of probability densities) has been defined in [16]. A statistical manifold endowed with a Wasserstein metric tensor structure is called Wasserstein statistical manifold. We define the Ricci curvature lower bound via geodesic convexity of the KL divergence on a Wasserstein statistical manifold. We obtain a definition of the Ricci-curvature that connects Wasserstein geometry [24] and information geometry [1, 2], much in the spirit of [15, 16], and take a natural further step towards connecting the two fields, in particular, relating notions from learning applications and the geometry of the statistical models. We focus on discrete sample spaces, which allows us to present a clear picture of the relations deriving from this theory, and leave the details of continuous settings for future work.

We consider a discrete statistical model described by a tuple (Θ, I, \mathbf{p}) consisting of a parameter space Θ , a discrete sample space (or state space) $I = \{1, \dots, n\}$, and a parametrization $\mathbf{p} \colon \Theta \to \mathcal{P}(I)$. Here $\mathcal{P}(I)$ denotes the set of all probability distributions on I. We say that (Θ, I, \mathbf{p}) has Ricci curvature lower bound $\kappa \in \mathbb{R}$ with respect to a given reference measure q, if and only if, for any $\theta \in \Theta$, it holds that

$$G_F(\theta) + \sum_{a \in I} \left(d_{\theta\theta} p_a(\theta) \log \frac{p_a(\theta)}{q_a} - \Gamma^{W,a}(\theta) \frac{d}{d\theta_a} \mathcal{D}_{\mathrm{KL}}(p(\theta) \| q) \right) \succeq \kappa G_W(\theta).$$

Here G_F is the Fisher-Rao metric tensor, G_W is the L^2 -Wasserstein metric tensor, $d_{\theta\theta}p$ is the second differential of the parameterization, $\Gamma^W(\theta)$ are the Christoffel symbols of the Wasserstein statistical manifold, and $D_{KL}(p(\theta)||q) = \sum_{i=1}^{n} p_i(\theta) \log \frac{p_i(\theta)}{q_i}$ is the KL divergence. This definition depends on the reference measure q. In statistics and learning applications, the reference measure will play the role of a target or empirical data distribution. A schematic illustration of the spaces and relations that we consider is provided in Figure 1.

The Ricci curvature on discrete state spaces has been studied by many groups. (i) Ollivier [20] introduces a discrete Ricci curvature via L_1 -Wasserstein metric. Many inequalities on graphs are shown under this setting; see, e.g., [12, 13, 21]. (ii) Lin-Yau et al. [17, 18] also define a Ricci curvature lower bound by heat semi-groups and Bakery-Emery Γ_2 operators. (iii) Erbar-Maas introduce the Ricci curvature lower bound in [9] by by means of equivalence relations with Lott-Strum-Villani in the Wasserstein probability manifold, under which several information functional inequalities are established. This notion has been studied extensively in [6, 7, 8, 10, 11]. However, the notion of a Ricci curvature lower bound on the parameter space of a statistical manifold has not been studied so far. Parametrized Wasserstein probability sub-manifolds were not introduced until recently in [4, 16]. Our definition of the Ricci curvature lower bound for parametrized statistical models is close in spirit to the definitions by Lott-Strum-Villani and Erbar-Maas.

This paper is organized as follows. In Section 2, we briefly review the connections between Ricci curvature, optimal transport, and KL divergence. We further demonstrate these connections in the context of information projections. In Section 3, we introduce Wasserstein statistical manifolds in parameter space. This is intended as a short review of the definitions from [16]. We derive the Fokker-Planck equation on parameter space, which is the Wasserstein gradient flow of the KL divergence. The main technical contributions

 $\mathbf{2}$



FIGURE 1. Our discussion involves a state space I, a parameter space Θ , and a parametrized set $\mathbf{p}(\Theta)$ in the space $\mathcal{P}(I)$ of probability distributions on I. For a reference measure $q \in \mathcal{P}(I)$, a positive Ricci curvature lower bound implies that the Wasserstein geodesic connecting two distributions, $p(\theta_0)$ and $p(\theta_1)$, 'bends' towards q. The figure depicts the geodesic as a thick curve, together with the level sets of $D_{\mathrm{KL}}(p(\cdot)||q)$, in Θ and $\mathbf{p}(\Theta)$. In terms of the state space I, when q is uniform, a decrease of the KL divergence with respect to q corresponds to an increase of the entropy, meaning that along the geodesic, the 'volume' of states under the distributions 'bulges'. This corresponds to the synthetic notion of positive curvature in sample space. Note how the geodesics are constrained to lie within the model $\mathbf{p}(\Theta)$, which in general does not contain q. See Definition 7, Theorem 8, Proposition 9, and Figures 2 and 3 for more details.

of this paper are contained in Section 4. We describe the convergence rate of the Fokker-Planck equation in terms of a Ricci curvature lower bound. Further, we use the notion of Ricci curvature lower bound to establish information functional inequalities. We also discuss methods to estimate the Ricci curvature lower bound in practice. In Section 5, we

present experiments on small examples of exponential families. These allow us to illustrate the notions introduced in the paper, and gain more intuition about their meaning.

2. RICCI CURVATURE AND INFORMATION PROJECTIONS

In this section, we review the connection of optimal transport and information theory put forward in Villani's book [24], and we further connect with the notion of information projections described by Csiszár-Shields [5]. In later sections we will develop these connections for the case of parametric statistical models.

2.1. Wasserstein geometry. Consider a continuous measure space (Ω, g^{Ω}, q) . Here Ω is a finite dimensional compact smooth Riemannian manifold without boundary, g^{Ω} is its metric tensor, dx is the volume from of Ω , and $q \in C^{\infty}(\Omega)$ is the measure volume form with $\int_{\Omega} q(x) = 1$, q(x) > 0. The Ricci curvature tensor on (Ω, g^{Ω}, q) refers to

$$\operatorname{Ric} = \operatorname{Ric}_{\Omega} - \operatorname{Hess}_{\Omega} \log q, \tag{1}$$

where $\operatorname{Ric}_{\Omega}$ denotes the Ricci curvature on Ω and $\operatorname{Hess}_{\Omega}$ is the Hessian operator on Ω . Note that this notion of curvature depends on the reference measure q. Later in our discussion, the reference measure will play the role of a target or empirical data distribution.

On the one hand, optimal transport, in particular the L^2 -Wasserstein metric, introduces an infinite-dimensional Riemannian structure in density space. In the context of our discussion, consider the set of smooth and strictly positive densities

$$\mathcal{P}_{+}(\Omega) = \Big\{ \rho \in C^{\infty}(\Omega) \colon \rho(x) > 0, \ \int_{\Omega} \rho(x) dx = 1 \Big\}.$$

The tangent space of $\mathcal{P}_+(\Omega)$ at $\rho \in \mathcal{P}_+(\Omega)$ is given by

$$T_{\rho}\mathcal{P}_{+}(\Omega) = \Big\{ \sigma \in C^{\infty}(\Omega) \colon \int_{\Omega} \sigma(x) dx = 0 \Big\}.$$

Definition 1 (L²-Wasserstein metric tensor). Define the inner product $g_{\rho}: T_{\rho}\mathcal{P}_{+}(\Omega) \times T_{\rho}\mathcal{P}_{+}(\Omega) \to \mathbb{R}$ by

$$g_{\rho}(\sigma_1, \sigma_2) = \int_{\Omega} \sigma_1(x) (-\Delta_{\rho})^{\dagger} \sigma_2(x) dx,$$

where $\Delta_{\rho}^{\dagger} \colon T_{\rho}\mathcal{P}_{+}(\Omega) \to T_{\rho}\mathcal{P}_{+}(\Omega)$ is the inverse of elliptical operator $\Delta_{\rho} = \nabla \cdot (\rho \nabla)$. Here ∇ and $\nabla \cdot$ are the gradient and divergence operators in Ω , respectively.

Following [14], we call $(\mathcal{P}_+(\Omega), g)$ a Wasserstein density manifold or a Wasserstein manifold for short. The metric tensor introduces a variational formulation of a metric function. More precisely, the square of the L^2 -Wasserstein metric function is equal to the geometric energy (action) of geodesics in the Wasserstein manifold. For any $\rho_0, \rho_1 \in \mathcal{P}_+(\Omega)$, the L^2 -Wasserstein metric function is defined as

$$W(\rho_0,\rho_1)^2 = \inf \left\{ \int_0^1 g_{\rho_t}(\partial_t \rho_t, \partial_t \rho_t) dt \colon \rho_t \in \mathcal{P}_+(\Omega), t \in [0,1] \right\}.$$

One can extend the definitions from $\mathcal{P}_+(\Omega)$ to the set $\mathcal{P}_2(\Omega)$ of Borel probability measures with finite second moments. It is well known that the L^2 -Wasserstein metric defines a metric function on $\mathcal{P}_2(\Omega)$, and hence $(\mathcal{P}_2(\Omega), W)$ forms a length space. See related analytical treatments in [24].

2.2. Wasserstein gradient flow of the KL divergence. On the other hand, information theory considers a particular functional on density space, namely the KL divergence. Given a smooth reference measure $q \in \mathcal{P}_+(\Omega)$, the KL divergence of a given ρ with respect to q is defined by

$$D_{\mathrm{KL}}(\rho \| q) = \int_{\Omega} \rho(x) \log \frac{\rho(x)}{q(x)} dx.$$

Notice that the KL divergence is precisely the free energy. Indeed, if we write $q(x) = \frac{1}{K}e^{-V(x)}$ with $K = \int_{\Omega} e^{-V(x)} dx$, we see that

$$D_{\mathrm{KL}}(\rho \| q) = \int_{\Omega} \rho(x) \log \rho(x) dx + \int_{\Omega} V(x) \rho(x) dx + \log K$$
$$= -H(\rho) + \mathbb{E}_{\rho}[V(X)] + \log K ,$$

where $H(\rho) = -\int_{\Omega} \rho(x) \log \rho(x) dx$ is the Boltzmann-Shannon entropy, X is a random variable satisfying the law of density ρ and \mathbb{E} is the expectation operator.

The Ricci curvature on sample space is related both to the KL divergence and the L^2 -Wasserstein metric tensor. This interaction starts with the gradient flow of the KL divergence in the Wasserstein manifold $(\mathcal{P}_+(\Omega), g)$, which describes the time evolution of the density following the negative Wasserstein gradient of the KL divergence:

$$\frac{\partial \rho_t}{\partial t} = -\operatorname{grad}_W \operatorname{D}_{\operatorname{KL}}(\rho \| q)
= \nabla \cdot \left(\rho_t \nabla (\log \frac{\rho_t}{q} + 1)\right)
= \nabla \cdot \left(\rho_t \nabla V\right) + \Delta \rho_t.$$
(2)

The second line is by Definition 1 of the Wasserstein metric tensor. The last equality holds since $q(x) = \frac{1}{K}e^{-V(x)}$ and $\nabla \cdot (\rho \nabla \log \rho) = \nabla \cdot (\nabla \rho) = \Delta \rho$.

It is worth noting that there are several perspectives based on (2). Firstly, the flow (2) is a well-known dynamics called *Fokker-Planck equation* (FPE). It describes the probability transition equation of drift diffusion process

$$\dot{X}_t = -\nabla V(X_t) + \sqrt{2}\dot{B}_t,$$

where B_t is the canonical Brownian motion in sample space. Secondly, along the flow (2), the KL divergence converges to zero. I.e. ρ_t converges to the minimizer of the KL divergence (free energy), known as the Gibbs measure, $q(x) = \frac{1}{K}e^{-V(x)}$. This reminds of iterative methods for computing information projections [5] in statistics and machine learning. In this context, one seeks to reproduce the behavior of a teacher system in terms of a model. To this end, the learning rule proceeds by adjusting the model parameters so as to maximize the likelihood of the observations, which is equivalent to minimizing the divergence, for instance using Wasserstein gradient descent. The flow is the continuous limit of the gradient descent learning rule. We shall go to this connection shortly, in Section 2.5.

2.3. Dissipation rates and the Ricci curvature lower bound. As it turns out, the Ricci curvature lower bound governs the exponential dissipation rate of (2) towards the Gibbs measure q. In the setting of learning, this corresponds precisely to the exponential rate of convergence of the learning dynamics. To see this, the following calculations in dynamical system are used. One can find the convergence rate of (2) by comparing the ratio between the first and second time derivatives along the flow. By some computations, the first time derivative of the KL divergence along the flow is found to be equal to

$$-\frac{d}{dt} \operatorname{D}_{\mathrm{KL}}(\rho_t \| q) = g_{\rho_t}(\partial_t \rho_t, \partial_t \rho_t)$$
$$= \int_{\Omega} \Gamma(\log \frac{\rho_t}{q}, \log \frac{\rho_t}{q}) \rho_t \, dx,$$

while the second time derivative is given by

10

$$\frac{d^2}{dt^2} \operatorname{D}_{\mathrm{KL}}(\rho_t \| q) = \operatorname{Hess}_W \operatorname{D}_{\mathrm{KL}}(\rho_t \| q) (\partial_t \rho_t, \partial_t \rho_t) = \int_{\Omega} \Gamma_2(\log \frac{\rho_t}{q}, \log \frac{\rho_t}{q}) \rho_t \, dx \; .$$

Here Hess_W is the Hessian operator with respect to the Wasserstein metric tensor, and Γ and Γ_2 are the Bakery-Emery operators defined by

$$\Gamma(f,f) = \int_{\Omega} g^{\Omega}(\nabla f, \nabla f) \ dx,$$

and

$$\Gamma_2(f, f) = (\operatorname{Ric}_{\Omega} - \operatorname{Hess}_{\Omega} \log q)(\nabla f, \nabla f) + \operatorname{tr}(\operatorname{Hess}_{\Omega} f, \operatorname{Hess}_{\Omega} f),$$

where $\operatorname{Ric}_{\Omega}$ is the Ricci curvature tensor on Ω , $\operatorname{Hess}_{\Omega}$ is the Hessian operator on Ω , and tr is the trace operator. By the above formulas, the ratio between $\frac{d}{dt} \operatorname{D}_{\mathrm{KL}}(\rho_t || q)$ and $\frac{d^2}{dt^2} \operatorname{D}_{\mathrm{KL}}(\rho_t || q)$ relates to the integral version of Γ , Γ_2 , i.e. the expectation values of the operators Γ , Γ_2 . Notice that $\operatorname{tr}(\operatorname{Hess}_{\Omega} f, \operatorname{Hess}_{\Omega} f) \geq 0$. Classical results [24] show that the lower bound of Ricci curvature governs the smallest ratio between $\frac{d}{dt} \operatorname{D}_{\mathrm{KL}}(\rho_t || q)$ and $\frac{d^2}{dt^2} \operatorname{D}_{\mathrm{KL}}(\rho_t || q)$, which further gives the exponential convergence rate of (2). In addition, the above computation demonstrate that the lower bound of the Ricci curvature, informally speaking, is equivalent to the smallest eigenvalue of the Hessian operator of the KL divergence in the Wasserstein manifold.

Theorem 2. Given $\kappa \in \mathbb{R}$ and $q(x) \in \mathcal{P}_+(\Omega)$, the following statements are equivalent.

(i) κ is a Ricci curvature lower bound of (Ω, g^{Ω}, q) . I.e. κ is the largest number for which, uniformly over Ω ,

$$\operatorname{Ric} = \operatorname{Ric}_{\Omega} - \operatorname{Hess}_{\Omega} \log q \succeq \kappa g^{\Omega};$$

- (ii) $\Gamma_2(f, f) \ge \kappa \Gamma(f, f)$, for any $f \in C^{\infty}(\Omega)$;
- (iii) For any constant speed geodesic ρ_t , $t \in [0, 1]$, connecting ρ_0 and ρ_1 in $(\mathcal{P}_2(\Omega), W)$,

$$D_{\mathrm{KL}}(\rho_t \| q) \le (1-t) D_{\mathrm{KL}}(\rho_0 \| q) + t D_{\mathrm{KL}}(\rho_1 \| q) - \frac{\kappa}{2} t (1-t) W(\rho_0, \rho_1)^2$$

Theorem 2 opens the door to define a notion of Ricci curvature lower bound on sample space via its equivalent statements. In the literature, Bakery-Emery [3] define the Ricci curvature lower bound by applying (ii) for smooth Riemannian sample spaces, while Lott-Strum-Villani [19, 23] define it using (iii) for non-smooth metric sample spaces, and Erbar-Maas [9] define it by (iii) in a discrete sample space. In this paper, we shall define the notion of Ricci curvature lower bound for parametric statistics taking an approach based on (iii), known as the geodesic convexity property of the KL divergence.

2.4. Learning in a parametrized model. In statistics and machine learning applications one often considers a parametrized set $\{p(\cdot; \theta) : \theta \in \mathbb{R}^d\}$ of candidate probability distributions from which one wishes to choose one to model the distribution of some given data.

One motivation for using parametrized models is reducing the dimensionality associated with large state spaces. For instance, we may be considering a state space consisting of images presented as arrays of pixel intensities, corresponding to $\{0,1\}^n$ with n easily in the order of thousands. In this case, storing a probability distribution as a vector $p \in \mathbb{R}^{2^n}$ of individual probabilities $p(x), x \in \{0,1\}^n$, is an impossibility. With a parametric model, instead of storing the probability vector, we store only a parameter vector $\theta \in \mathbb{R}^d$, with a more manageable d, and fix a mapping that allows us to recover individual values $p(x; \theta)$ of the probability distribution for a given x, or, in other cases, which allows us to generate samples from $p(\cdot; \theta)$ that we can also use to estimate any expectation values of interest. Reducing the dimensionality is useful not only in terms of storage, but also in a statistical sense, in relation to overfitting. Without going into details, the richer the class of hypotheses, with more free parameters, the more prone we are to fitting statistical nuisances of the data, instead of capturing the true general behavior of the data. By working with a parametrized model, we can incorporate priors into the learning system and limit its vulnerability to overfitting.

When working with a parametrized model, obtaining the best possible hypothesis, e.g. the maximizer of the data likelihood, is usually a non trivial problem and one has to resort to iterative methods. A relevant question then is the computational effort needed for this. In particular, one is interested in the number of iterations needed until reaching a solution that is within ϵ of the best possible. The Ricci curvature can be regarded as a way to obtain bounds on the convergence rate of gradient optimization of the KL divergence for a given a target distribution, uniformly over the start distribution. We elaborate on this in the next paragraph. The situation is illustrated in Figures 1, 2, and 3.

2.5. **I-projections.** In the context of information theory and statistics, Csiszár-Shields [5] define the *I-projection* of a distribution Q onto a non-empty closed convex set \mathcal{N} of distributions as the point $P^* \in \mathcal{N}$ such that

$$D(P^* \| Q) = \min_{P \in \mathcal{N}} D(P \| Q).$$

The notion of I-projection considers the minimization of the KL divergence with respect to the first argument, but it is also relevant in the context of maximum likelihood estimation, where the minimization is with respect to the second argument. Given an empirical data



FIGURE 2. For a distribution Q in an exponential family \mathcal{E} and a distribution P in an orthogonal linear family \mathcal{N} , the Pythagorean relation holds: $D(P||Q) = D(P||P^*) + D(P^*||Q)$, where P^* is the unique intersection point of \mathcal{E} and \mathcal{N} .

distribution P, a maximum likelihood estimator over a set \mathcal{E} is a point $P^* \in \overline{\mathcal{E}}$ (the closure of \mathcal{E}), with

$$D(P||P^*) = \inf_{Q \in \mathcal{E}} D(P||Q).$$

If we consider an exponential family model $\mathcal{E} = \{p \propto Q \exp(\theta^{\mathsf{T}} F) : \theta \in \mathbb{R}^d\}$ on a finite state space I, with sufficient statistics $F : I \to \mathbb{R}^d$ and reference measure $Q \in \mathcal{P}_+(I)$, then the maximum likelihood estimator P^* of the target distribution P can be obtained as the I-projection of Q onto the orthogonal linear family defined by $\mathcal{N} = \{p : \sum_x F(x)p(x) = \sum_x F(x)P(x)\}$. We have namely that

$$P^* = \operatorname{argmin}_{Q \in \mathcal{E}} D(P \| Q) = \operatorname{argmin}_{P \in \mathcal{N}} D(P \| Q).$$

This is a consequence of the well known Pythagorean relation [5, 1] illustrated in Figure 2.

Csiszár and Shields [5] consider iterative methods for computing I-projections, and obtain upper bounds on the divergence along the resulting parameter trajectories, which describe the convergence to the optimum value. For two sets of distributions, \mathcal{P} and \mathcal{Q} , together with two functions $D(\cdot, \cdot): \mathcal{P} \times \mathcal{Q} \to \mathbb{R}$ and $\delta(\cdot, \cdot): \mathcal{P} \times \mathcal{P} \to \mathbb{R}$, satisfying certain conditions, they describe an iterative algorithm (alternating divergence minimization) which iterates $p_n \in \mathcal{P}$ and $q_n \in \mathcal{Q}$, and give an upper bound of the form

$$D(p_{n+1}, q_n) - D_{\min} \le \delta(p_{\infty}, p_n) - \delta(p_{\infty}, p_{n+1}).$$
(3)

In this paper we are in the special setting where $\mathcal{Q} = \{q\}$ and \mathcal{P} is the set of all densities. There is a natural connection between (3) and the Fokker-Planck-equation (2). Indeed, setting D as the KL divergence and $q_n = q = p_{\infty}$, $p_n = \rho_t$, $p_{n+1} = \rho_{t+\Delta t}$, where Δt is the step size, we demonstrate in Proposition 9 that we can substitute

$$\delta(q, p) = \frac{1}{2\kappa\Delta t} \operatorname{D}_{\mathrm{KL}}(p\|q),$$



FIGURE 3. Illustration of the Ricci curvature lower bound κ in connection to the geodesic convexity of the KL divergence and the rate of convergence of the information projection flow. Here $p(\theta_t)$ is a Wasserstein geodesic connecting $p(\theta_0)$ and $p(\theta_1)$. When $q = p(\theta_1)$, the KL divergence $D_{KL}(p(\theta_t)||q)$ is monotonically decreasing.

where κ is the Ricci curvature lower bound that we will define later on. Strictly speaking, for this correspondence, we need to assume that $\kappa > 0$, which is a natural requirement similar to requiring that the KL divergence is geodesic convex in set \mathcal{P} . Each step dissipation in (3) then gives

$$\begin{aligned} \mathbf{D}_{\mathrm{KL}}(\rho_{t+\Delta t} \| q) - \mathbf{D}_{\mathrm{min}} &\leq \delta(p_{\infty}, p_n) - \delta(p_{\infty}, p_{n+1}) \\ &= \frac{1}{2\kappa\Delta t} \Big\{ \mathbf{D}_{\mathrm{KL}}(\rho_t \| q) - \mathbf{D}_{\mathrm{KL}}(\rho_{t+\Delta t} \| q) \Big\} \\ &= -\frac{1}{2\kappa} \frac{d}{dt} \mathbf{D}_{\mathrm{KL}}(\rho_t \| q) + o(\Delta t). \end{aligned}$$

In other words, the Fokker-Planck equation is a monotone information projection flow, in which the dissipation quantity is governed by the difference of relative entropy divided by twice the Ricci curvature lower bound. In the limit where Δt goes to zero,

$$\mathbf{D}_{\mathrm{KL}}(\rho_t \| q) - \mathbf{D}_{\min} \le -\frac{1}{2\kappa} \frac{d}{dt} \mathbf{D}_{\mathrm{KL}}(\rho_t \| q).$$

Grönwall's inequality then implies that this I-projection flow (2) converges to the minimizer at the rate of $e^{-2\kappa t}$, i.e.

$$\mathbf{D}_{\mathrm{KL}}(\rho_t \| q) - \mathbf{D}_{\mathrm{min}} \le e^{-2\kappa t} \Big(\mathbf{D}_{\mathrm{KL}}(\rho_0 \| q) - \mathbf{D}_{\mathrm{min}} \Big).$$

The above shows that the learning rate for the Fokker-Planck equation is linear, whose lower bound is governed by κ . Following these connections, we will pursue the definition of the Ricci curvature lower bound on parameter space. The convergence rate, in relation to the Ricci curvature lower bound and the geodesic convexity of the KL divergence, is illustrated schematically in Figure 3. More details will be provided in Proposition 9.

3. WASSERSTEIN STATISTICAL MANIFOLDS

In preparation for the definitions and results on the Ricci curvature that we will present in the next section, we briefly review the definition of a Wasserstein statistical manifold with discrete sample space from [16], and present the Fokker-Planck equation on parameter space.

3.1. Wasserstein geometry on the probability simplex. We recall the definition of discrete probability simplex with L^2 -Wasserstein Riemannian metric. Consider the discrete sample space $I = \{1, \dots, n\}$. The probability simplex on I is the set

$$\mathcal{P}(I) = \left\{ (p_1, \cdots, p_n) \in \mathbb{R}^n \colon \sum_{i=i}^n p_i = 1, \quad p_i \ge 0 \right\}.$$

Here $p = (p_1, \ldots, p_n)$ is a probability vector with coordinates p_i corresponding to the probabilities assigned to each node $i \in I$. The probability simplex $\mathcal{P}(I)$ is a manifold with boundary. We denote the interior by $\mathcal{P}_+(I)$. This consists of the strictly positive probability distributions, with $p_i > 0$ for all $i \in I$. To simplify the discussion, we will focus on the interior $\mathcal{P}_+(I)$. For the studies related to the boundary $\partial \mathcal{P}(I)$, we refer the reader to [15].

Next we define the L^2 -Wasserstein metric tensor on $\mathcal{P}_+(I)$, which also encodes the metric tensor of discrete states I. We need to give a ground metric notion on sample space. We do this in terms of a undirected graph with weighted edges, $G = (I, E, \omega)$, where I is the vertex set, $E \subseteq {I \choose 2}$ is the edge set, and $\omega = (\omega_{ij})_{i,j\in I} \in \mathbb{R}^{n \times n}$ is a matrix of edge weights satisfying

$$\omega_{ij} = \begin{cases} \omega_{ji} > 0, & \text{if } (i,j) \in E \\ 0, & \text{otherwise} \end{cases}.$$

The set of neighbors (adjacent vertices) of *i* is denoted by $N(i) = \{j \in V : (i, j) \in E\}$. The normalized volume form on node $i \in I$ is given by $d_i = \frac{\sum_{j \in N(i)} \omega_{ij}}{\sum_{i=1}^n \sum_{i' \in N(i)} \omega_{ii'}}$.

The graph structure $G = (I, E, \omega)$ induces a graph Laplacian matrix function.

Definition 3 (Weighted Laplacian matrix). Given an undirected weighted graph $G = (I, E, \omega)$, with $I = \{1, \ldots, n\}$, the matrix function $L(\cdot) : \mathbb{R}^n \to \mathbb{R}^{n \times n}$ is defined by

$$L(p) = D^{\mathsf{T}} \Lambda(p) D, \quad p = (p_i)_{i=1}^n \in \mathbb{R}^n,$$

where

• $D \in \mathbb{R}^{|E| \times n}$ is the discrete gradient operator defined by

$$D_{(i,j)\in E,k\in V} = \begin{cases} \sqrt{\omega_{ij}}, & \text{if } i = k, \ i > j \\ -\sqrt{\omega_{ij}}, & \text{if } j = k, \ i > j \\ 0, & \text{otherwise} \end{cases}$$

• $-D^{\mathsf{T}} \in \mathbb{R}^{n \times |E|}$ is the oriented incidence matrix, and

• $\Lambda(p) \in \mathbb{R}^{|E| \times |E|}$ is a weight matrix depending on p,

$$\Lambda(p)_{(i,j)\in E,(k,l)\in E} = \begin{cases} \frac{1}{2}(\frac{1}{d_i}p_i + \frac{1}{d_j}p_j) & \text{if } (i,j) = (k,l)\in E\\ 0 & \text{otherwise} \end{cases}$$

The Laplacian matrix function L(p) is the discrete analog of the weighted Laplacian operator $-\nabla \cdot (\rho \nabla)$ from Definition 1.

We are now ready to present the L^2 -Wasserstein metric tensor. Consider the tangent space of $\mathcal{P}_+(I)$ at p,

$$T_p \mathcal{P}_+(I) = \Big\{ (\sigma_i)_{i=1}^n \in \mathbb{R}^n \colon \sum_{i=1}^n \sigma_i = 0 \Big\}.$$

Denote the space of *potential functions* on I by $\mathcal{F}(I) = \mathbb{R}^n$, and consider the quotient space

$$\mathcal{F}(I)/\mathbb{R} = \{ [\Phi] \mid (\Phi_i)_{i=1}^n \in \mathbb{R}^n \},\$$

where $[\Phi] = \{(\Phi_1 + c, \cdots, \Phi_n + c) : c \in \mathbb{R}\}$ are functions defined up to addition of constants.

We introduce an identification map via the weighted Laplacian matrix L(p):

$$\mathbf{V}: \mathcal{F}(I)/\mathbb{R} \to T_p \mathcal{P}_+(I), \qquad \mathbf{V}_\Phi = L(p)\Phi.$$

We know that L(p) has only one simple zero eigenvalue with eigenvector $c(1, 1, \dots, 1)$, for any $c \in \mathbb{R}$. This is true since for $(\Phi_i)_{i=1}^n \in \mathbb{R}^n$,

$$\Phi^{\mathsf{T}}L(p)\Phi = (D\Phi)^{\mathsf{T}}\Lambda(p)(D\Phi) = \sum_{(i,j)\in E} \omega_{ij}(\Phi_i - \Phi_j)^2(\frac{1}{2}(\frac{1}{d_i}p_i + \frac{1}{d_j}p_j)) = 0,$$

implies $\Phi_i = \Phi_j$, $(i, j) \in E$. It the graph is connected, as we assume, then $(\Phi_i)_{i=1}^n$ is a constant vector. Thus $V_{\Phi} \colon \mathcal{F}(I)/\mathbb{R} \to T_p \mathcal{P}_+(I)$ is a well defined map, linear, and one to one. I.e., $\mathcal{F}(I)/\mathbb{R} \cong T_p^* \mathcal{P}_+(I)$, where $T_p^* \mathcal{P}_+(I)$ is the cotangent space of $\mathcal{P}_+(I)$. This identification induces the following inner product on $T_p \mathcal{P}_+(I)$.

Definition 4 (L²-Wasserstein metric tensor). The inner product $g_p : T_p \mathcal{P}_+(I) \times T_p \mathcal{P}_+(I) \rightarrow \mathbb{R}$ takes any two tangent vectors $\sigma_1 = \mathbf{V}_{\Phi_1}$ and $\sigma_2 = \mathbf{V}_{\Phi_2} \in T_p \mathcal{P}_+(I)$ to

$$g_p(\sigma_1, \sigma_2) = \sigma_1^{\mathsf{T}} \Phi_2 = \sigma_2^{\mathsf{T}} \Phi_1 = \Phi_1^{\mathsf{T}} L(p) \Phi_2.$$
 (4)

In other words,

 $g_p(\sigma_1, \sigma_2) := \sigma_1^{\mathsf{T}} L(p)^{\dagger} \sigma_2, \quad \text{for any } \sigma_1, \sigma_2 \in T_p \mathcal{P}_+(I),$

where $L(p)^{\dagger}$ is the pseudo inverse of L(p).

Following the inner product (4), the Wasserstein metric (distance function) $W: \mathcal{P}_+(I) \times \mathcal{P}_+(I) \to \mathbb{R}$ is defined by

$$W(p^{0}, p^{1})^{2} := \inf_{p(t), \Phi(t)} \left\{ \int_{0}^{1} \Phi(t)^{\mathsf{T}} L(p(t)) \Phi(t) dt \right\} .$$
(5)

Here the infimum is taken over pairs $(p(t), \Phi(t))$ with $p \in H^1((0, 1), \mathbb{R}^n)$ and $\Phi: [0, 1] \to \mathbb{R}^n$ measurable, satisfying

$$\frac{d}{dt}p(t) - L(p(t))\Phi(t) = 0, \quad p(0) = p^0, \quad p(1) = p^1.$$

3.2. Wasserstein statistical manifold. We next consider a statistical model defined by a triplet (Θ, I, \mathbf{p}) . Here, $I = \{1, \dots, n\}$ is the sample space, Θ is the parameter space, which is an open subset of \mathbb{R}^d , $d \leq n-1$, and $\mathbf{p} \colon \Theta \to \mathcal{P}_+(I)$ is the parametrization function,

$$p(\theta) = (p_i(\theta))_{i=1}^n, \quad \theta \in \Theta.$$

We define a Riemannian metric g^W on Θ as the pull-back of metric g on $\mathcal{P}_+(I)$. In other words, we require that $\mathbf{p} \colon (\Theta, g^W) \to (\mathcal{P}_+(I), g)$ is an isometric embedding:

$$g_{\theta}^{W}(a,b) := g_{p(\theta)}^{W}(d_{\theta}p(\theta)(a), d_{\theta}p(\theta)(b))$$

= $(d_{\theta}p(\theta)(a))^{\mathsf{T}}L(p(\theta))^{\dagger}(d_{\theta}p(\theta)(b)), \text{ for all } a, b \in T_{\theta}(\Theta).$

Since $dp(\theta)(a) = \left(\sum_{j=1}^{d} \frac{\partial p_i(\theta)}{\partial \theta_j} a_j\right)_{i=1}^n = J_{\theta} p(\theta) a$, we arrive at the following definition.

Definition 5 (L^2 -Wasserstein metric tensor on parameter space). For any pair of tangent vectors $a, b \in T_{\theta}\Theta = \mathbb{R}^d$, define

$$G_W(\theta) := J_\theta p(\theta)^\mathsf{T} L(p(\theta))^\dagger J_\theta p(\theta), \tag{6}$$

and

$$g_{\theta}^{W}(a,b) := a^{\mathsf{T}} G_{W}(\theta) b,$$

where $J_{\theta}(p(\theta)) = (\frac{\partial p_i(\theta)}{\partial \theta_j})_{1 \le i \le n, 1 \le j \le d} \in \mathbb{R}^{n \times d}$ is the Jacobi matrix of the parametrization **p**.

This inner product is consistent with the restriction of the Wasserstein metric g^W to $p(\theta)$. We will assume that rank $(J_{\theta}p(\theta)) = d$, so that the parametrization \mathbf{p}_i is locally injective and the metric tensor g^W is positive definite. We call (Θ, I, \mathbf{p}) , together with the induced Riemannian metric g^W , Wasserstein statistical manifold (WSM).

In this case, the constrained Wasserstein distance function $d_W : \Theta \times \Theta \to \mathbb{R}_+$ is given by the geometric action energy

$$d_W(\theta_0, \theta_1)^2 = \inf_{\theta(t) \in C^1([0,1],\Theta)} \Big\{ \int_0^1 \dot{\theta}(t)^\mathsf{T} G_W(\theta(t)) \dot{\theta}(t) dt \colon \theta(0) = \theta_0, \ \theta(1) = \theta_1 \Big\}.$$
(7)

When working on the full probability simplex, with $\theta = p$, the metric function d_W corresponds precisely to the metric function W given in (5).

3.3. Fokker-Planck equation on parameter space. We next derive the Fokker-Planck equation on parameter space by Wasserstein gradient flow of KL divergence.

Given a reference measure $q \in \mathcal{P}_+(I)$, consider the Kullback-Leibler divergence (relative entropy) on parameter space

$$D_{\mathrm{KL}}(p(\theta) \| q) = \sum_{i=1}^{n} p_i(\theta) \log \frac{p_i(\theta)}{q_i}.$$

Proposition 6 (Fokker-Planck equation on parameter space). The gradient flow for the negative Boltzmann-Shannon entropy in (Θ, g) is

$$\frac{d\theta}{dt} = -\left(J_{\theta}p(\theta)^{\mathsf{T}}L(p(\theta))^{\dagger}J_{\theta}p(\theta)\right)^{\dagger}J_{\theta}p(\theta)^{\mathsf{T}}\log\frac{p(\theta)}{q}.$$
(8)

Proof. The gradient flow of entropy on (Θ, g_W) satisfies

$$\begin{aligned} \frac{d\theta}{dt} &= -\operatorname{grad}_W \operatorname{D}_{\mathrm{KL}}(p(\theta) \| q) \\ &= -G_W(\theta)^{\dagger} \nabla_{\theta} \operatorname{D}_{\mathrm{KL}}(p(\theta) \| q) \\ &= -\left(J_{\theta} p(\theta)^{\mathsf{T}} L(p(\theta))^{\dagger} J_{\theta} p(\theta)\right)^{\dagger} J_{\theta} p(\theta)^{\mathsf{T}} (\log \frac{p(\theta)}{q} + \vec{1}), \end{aligned}$$

where ∇_{θ} represents the Euclidean gradient operator, $\log \frac{p(\theta)}{q} = (\log \frac{p_i(\theta)}{q_i})_{i=1}^n$ and $\vec{1} = (1, \dots, 1) \in \mathbb{R}^n$. Since $p(\theta)^{\mathsf{T}}\vec{1} = 1$, we have $J_{\theta}p(\theta)^{\mathsf{T}}\vec{1} = \vec{0}$. This completes the proof. \Box *Remark* 1. Consider the full probability set with continuous sample space Ω . Denote the probability $p_i(t)$ by a density $\rho(t, x) \in \mathcal{P}(\Omega)$, then (8) recovers the FPE (2).

We next study the convergence properties of the Fokker-Planck equation on parameter space. In other words, how fast does the solution of (8) converge to its equilibrium? As in the full probability space, we define the concept of Ricci curvature lower bound on parameter space to give the bound of the convergence rate for (8).

4. RICCI CURVATURE LOWER BOUND ON PARAMETER SPACE

This section contains the main contributions of this paper. We define the Ricci curvature lower bound on parameter space and prove equivalent conditions for this definition, which connect information geometry and Wasserstein geometry. In addition, we present several information functional inequalities on parameter space. Finally, we give a simple guide for computing these quantities in practice.

4.1. Ricci curvature lower bound on parameter space.

Definition 7. We say (Θ, I, \mathbf{p}) has the Ricci curvature lower bound $\kappa \in \mathbb{R}$ if for any constant speed geodesic θ_t , $t \in [0, 1]$, connecting θ_0 , θ_1 in (Θ, g^W) , it holds that

$$D_{\rm KL}(p(\theta_t) \| q) \le (1-t) D_{\rm KL}(p(\theta_0) \| q) + t D_{\rm KL}(p(\theta_1) \| q) - \frac{\kappa}{2} t (1-t) d_W(\theta_0, \theta_1)^2.$$

In this case we also write

$$\operatorname{Ric}(\Theta, I, \mathbf{p}) \geq \kappa.$$

If (Θ, g_W) forms a compact smooth Riemannian manifold and $p(\theta)$ is smooth. Then κ is the smallest eigenvalue of the Hessian of the KL divergence over the Wasserstein statistical manifold, i.e.

$$\operatorname{Hess}_W \operatorname{D}_{\operatorname{KL}}(p(\theta) \| q) \succeq \kappa G_W(\theta),$$

for any $\theta \in \Theta$.

Definition 7 is based on the definition of geodesic convexity in geometry. It is a more general definition than the one in terms of the Hessian operator begin bounded below by κ . The reason is as follows. On the one hand, the probability set is a manifold with boundary. Suitable regularity studies are needed to take care of the boundary when using the Hessian operator [15]. On the other hand, not all parameterizations $p(\theta)$ are twice differentiable.

Definition 7 shares the same spirit of Lott-Strum-Villani and Erbar-Maas. If $p(\theta) = \mathcal{P}(I)$ is the whole probability simplex, then $\operatorname{Ric}(\Theta, I, \mathbf{p})$ is the Ricci curvature bound on discrete sample space. Our definition extends this idea to a statistical manifold. In other words, $\operatorname{Ric}(\Theta, I, \mathbf{p})$ inherits properties from both probability submanifold and Ricci curvature bound on sample space. Note that $\operatorname{Ric}(\Theta, I, \mathbf{p})$ is different from the Ricci curvature on (Θ, g^W) , in which the former represents how the changes ratio KL divergence takes effect on the parameterized sample space, while the later reflects the curvature on the set of probability itself.

We next given an equivalent condition for Definition 7. It naturally connects Ricci curvature (R), Information geometry (I) and Wasserstein geometry (W). We call it Ricci-Information-Wasserstein (RIW) condition.

Theorem 8 (RIW condition). Assume Θ is a compact set. $\operatorname{Ric}(\Theta, I, \mathbf{p}) \geq \kappa$ holds if and only if for any $\theta \in \Theta$,

$$G_F(\theta) + \sum_{a \in I} \left(d_{\theta\theta} p_a(\theta) \log \frac{p_a(\theta)}{q} - \Gamma^{W,a}(\theta) \frac{d}{d_{\theta_a}} \mathcal{D}_{\mathrm{KL}}(p(\theta) \| q) \right) \succeq \kappa G_W(\theta), \tag{9}$$

where $G_F(\theta) = (G_F(\theta)_{ab})_{1 \le a, b \le d}$ is the Fisher-Rao metric tensor

$$G_F(\theta)_{ab} = \sum_{i \in I} \frac{d \log p_i(\theta)}{d\theta_a} \frac{d \log p_i(\theta)}{d\theta_b} p_i(\theta), \tag{10}$$

 $\Gamma^{W,k} = (\Gamma^{W,k}_{ij})_{1 \leq i,j \leq n}$ is the Wasserstein Christoffel symbol

$$\Gamma_{ij}^{W,k} = \frac{1}{2} \sum_{l=1}^{n} (G_{W,kl})^{-1} \Big(\nabla_{\theta_i} G_{W,jl} + \nabla_{\theta_j} G_{W,il} - \nabla_{\theta_l} G_{W,ij} \Big),$$

and $G_W(\theta)$ is the Wasserstein metric tensor defined in (6).

Proof. Let θ_t be a constant speed geodesic, i.e. $\ddot{\theta}_t + \Gamma^W(\dot{\theta}_t, \dot{\theta}_t) = 0$ with $\theta_0 = \theta \in \Theta$ and $\dot{\theta}_0 = a \in T_\theta \Theta$. Consider the Taylor expansion

$$D_{\rm KL}(p(\theta_t)||q) = D_{\rm KL}(p(\theta)||q) + \left. \frac{d}{dt} \right|_{t=0} D_{\rm KL}(p(\theta_t)||q)t + \frac{1}{2} \left. \frac{d^2}{dt^2} \right|_{t=0} D_{\rm KL}(p(\theta_t)||q)t^2 + o(t^2).$$

Then the Hessian operator on Riemannian manifold (Θ, g^W) forms

$$\begin{aligned} \operatorname{Hess}_{W} \mathrm{D}_{\mathrm{KL}}(p(\theta_{t}) \| q)(\dot{\theta}_{t}, \dot{\theta}_{t}) &= \frac{d^{2}}{dt^{2}} D_{\mathrm{KL}}(p(\theta_{t}) \| q) \\ &= \frac{d}{dt} (d_{\theta} D_{\mathrm{KL}}(p(\theta_{t}) \| q)^{\mathsf{T}} \dot{\theta}_{t}) \\ &= \dot{\theta}_{t}^{\mathsf{T}} d_{\theta\theta} \operatorname{D}_{\mathrm{KL}}(p(\theta_{t}) \| q) \dot{\theta}_{t} - d_{\theta} \operatorname{D}_{\mathrm{KL}}(p(\theta_{t}) \| q)^{\mathsf{T}} \Gamma^{W}(\dot{\theta}_{t}, \dot{\theta}_{t}) \\ &= \dot{\theta}_{t}^{\mathsf{T}} d_{\theta\theta} \operatorname{D}_{\mathrm{KL}}(p(\theta_{t}) \| q) \dot{\theta}_{t} - \dot{\theta}_{t}^{\mathsf{T}} (\sum_{k \in I} \frac{d}{d\theta_{k}} \operatorname{D}_{\mathrm{KL}}(p(\theta_{t}) \| q) \Gamma^{W,k}) \dot{\theta}_{t}. \end{aligned}$$

In addition,

$$d_{\theta} \operatorname{D}_{\mathrm{KL}}(p(\theta_{t}) \| q) = \sum_{i=1}^{n} \left(d_{\theta} p_{i}(\theta_{t}) \log p_{i}(\theta_{t}) + p_{i}(\theta_{t}) d_{\theta} \log p_{i}(\theta_{t}) - d_{\theta} p_{i}(\theta_{t}) \log q_{i} \right)$$
$$= \sum_{i=1}^{n} \left(d_{\theta} p_{i}(\theta_{t}) \log p_{i}(\theta_{t}) - d_{\theta} p_{i}(\theta_{t}) \log q_{i} \right),$$

where $\sum_{i=1}^{n} p_i(\theta_t) d_\theta \log p_i(\theta_t) = \sum_{i=1}^{n} p_i(\theta_t) \frac{1}{p_i(\theta_t)} J_\theta p_i(\theta_t) = 0$, since $\sum_{i=1}^{n} p_i(\theta) = 1$. Thus

$$d_{\theta\theta} \operatorname{D}_{\operatorname{KL}}(p(\theta_t) \| q) = \sum_{i=1}^n d_{\theta\theta} p_i(\theta_t) \log \frac{p_i(\theta_t)}{q} + \sum_{i=1}^n \frac{1}{p_i(\theta_t)} d_{\theta} p_i(\theta_t) d_{\theta} p_i(\theta_t)^{\mathsf{T}}$$
$$= \sum_{i=1}^n d_{\theta\theta} p_i(\theta_t) \log \frac{p_i(\theta_t)}{q} + G_F(\theta),$$

where G_F denotes the Fisher-Rao metric tensor: $G_F(\theta_t) = \sum_{i=1}^n \frac{1}{p_i(\theta_t)} d_\theta p_i(\theta_t) d_\theta p_i(\theta_t)^\mathsf{T} = \sum_{i=1}^n d_\theta \log p_i(\theta_t) d_\theta \log p_i(\theta_t)^\mathsf{T} p_i(\theta)$, with the fact $\frac{1}{p_i(\theta_t)} d_\theta p_i(\theta_t) = d_\theta \log p_i(\theta_t)$.

Thus $\operatorname{Hess}_W \operatorname{D}_{\operatorname{KL}}(p(\theta) \| q) \succeq \kappa G_W(\theta)$ is equivalent to (9). This concludes the proof. *Remark* 2. If we replace I by the continuous sample space (Ω, g^{Ω}) and consider the full probability simplex, then RIW condition (9) is equivalent to the integral version of Bakery-Emery condition. See details in [15, Proposition 19].

4.2. Entropy dissipation on parameter space. With the Ricci curvature lower bound in hand, we can prove the following convergence properties of Fokker-Planck equations on parameter space.

Proposition 9 (Bakery-Emery condition on parameter space). Assume Θ is a compact set. If $\operatorname{Ric}(\Theta, I, \mathbf{p}) \geq \kappa > 0$, then there exists a unique equilibrium $\theta^* \in \Theta$, with

$$\theta^* = \arg\min_{\theta \in \Theta} \mathcal{D}_{\mathrm{KL}}(p(\theta) \| q)$$

In addition, for any initial condition $\theta_0 \in \Theta$, the solution $\theta(t)$ of (2) converges to θ^* exponentially fast, with

$$D_{\mathrm{KL}}(p(\theta_t)\|q) - D_{\mathrm{KL}}(p(\theta^*)\|q) \le e^{-2\kappa t} \Big(D_{\mathrm{KL}}(p(\theta_0)\|q) - D_{\mathrm{KL}}(p(\theta^*)\|q) \Big), \quad \text{for all } t.$$
(11)

Remark 3. This result will apply for any geometry defined on Θ , whenever κ is the smallest eigenvalue of the corresponding Hessian operator of the divergence function.

Proof. The proof comes from the classical study of gradient flow in Riemannian manifold (Θ, g_W) . Since $\operatorname{Hess}_W \operatorname{D}_{\operatorname{KL}}(p(\theta) || q) \geq \kappa > 0$, the $\operatorname{D}_{\operatorname{KL}}(p(\theta) || q)$ is κ -geodesics convex in (Θ, g_W) . Thus $\theta(t)$ converges to the unique equilibrium θ^* , which is also the unique minimizer of KL divergence.

We next investigate how fast $\theta(t)$ converges to θ^* . The speed of convergence is obtained by comparing the first and second derivatives of the KL divergence w.r.t. time t along (2). We have

$$\frac{d}{dt} \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q) = -g_W(\operatorname{grad}_W \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q), \operatorname{grad}_W \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q)),$$

and

$$\frac{d^2}{dt^2} \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q) = 2 \operatorname{Hess}_W \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q) \big(\operatorname{grad}_W \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q), \operatorname{grad}_W \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q) \big).$$

From $\operatorname{Ric}(\Theta, I, \mathbf{p}) \ge \kappa > 0$, then $\operatorname{Hess}_W \operatorname{D}_{\operatorname{KL}}(p(\theta) || q) \ge \kappa > 0$, i.e.

$$\frac{d^2}{dt^2} \mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q) \ge -2\kappa \frac{d}{dt} \mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q), \quad \text{for all } t \ge 0.$$
(12)

Then by integrating the above formula over $[t, +\infty)$, one obtains

$$\frac{d}{dt}[\mathbf{D}_{\mathrm{KL}}(p(\theta^*)\|q) - \mathbf{D}_{\mathrm{KL}}(p(\theta_t)\|q)] \ge -2\kappa[\mathbf{D}_{\mathrm{KL}}(p(\theta^*)\|q) - \mathbf{D}_{\mathrm{KL}}(p(\theta_t)\|q)].$$

Proceed with the Grönwall's inequality, the result is proved.

4.3. Functional inequalities on parameter space. In literature [22], the convergence rate of FPE is used to prove several functional inequalities, including Log-Sobolev, Talagrand and HWI inequalities. The HWI inequality is a relation between the relative entropy (H), Wasserstein metric (W), relative Fisher information functional (\mathcal{I}). We shall derive the counterparts of these inequalities on parameter space.

Here the Log-Sobolev inequality describes a relationship between relative entropy and relative Fisher information functional on parameter space. Here the relative Fisher information functional is defined by

$$\mathcal{I}(p(\theta)\|q) := g_W(\operatorname{grad}_W \operatorname{D}_{\operatorname{KL}}(p(\theta_t)\|q), \operatorname{grad}_W \operatorname{D}_{\operatorname{KL}}(p(\theta_t)\|q)).$$
(13)

In particular, we formulate (13) as follows:

$$\begin{aligned} \mathcal{I}(p(\theta) \| q) &= \log \frac{p(\theta)}{q}^{\mathsf{T}} J_{\theta}(p(\theta)) \left(J_{\theta} p(\theta)^{\mathsf{T}} L(p(\theta))^{\dagger} J_{\theta} p(\theta) \right)^{\dagger} J_{\theta} p(\theta)^{\mathsf{T}} \log \frac{p(\theta)}{q} \\ &= \left(\operatorname{Proj}_{\theta} \log \frac{p(\theta)}{q} \right)^{\mathsf{T}} L(p(\theta)) \left(\operatorname{Proj}_{\theta} \log \frac{p(\theta)}{q} \right) \\ &= \sum_{i=1}^{n} \sum_{j \in N(i)} \frac{1}{2d_{i}} \omega_{ij} \left((\operatorname{Proj}_{\theta} \log \frac{p(\theta)}{q})_{i} - (\operatorname{Proj}_{\theta} \log \frac{p(\theta)}{q})_{j} \right)^{2} p_{i}(\theta), \end{aligned}$$

where $\operatorname{Proj}_{\theta} = (J_{\theta}p(\theta)^{\mathsf{T}})^{\dagger}J_{\theta}p(\theta)^{\mathsf{T}} \in \mathbb{R}^{n \times n}$ is the projection matrix, which projects the differential operator in full probability into the one in parameter space. We compare (13) with the one in continuous sample space and full probability space:

$$\mathcal{I}(\rho \| q) = \int_{\Omega} g^{\Omega}(\nabla \log \frac{\rho}{q}, \nabla \log \frac{\rho}{q}) \rho \ dx.$$

We note that the functional (13) is different from the commonly known Fisher information matrix (10) in parameter space. It contains the ground metric structure in the sample space, which is inherited from the L^2 -Wasserstein metric tensor $L(p)^{\dagger}$. In other words, when applying the Fisher information in full probability set into parameter space, the following two angles arrive. Here (13) keeps the differential structure of sample space and project the differential of KL divergence into the parameter space, while Fisher information matrix (10) replaces the differential structures of sample space to the ones in parameters.

In the following, we derive inequalities based on (13).

16

Proposition 10 (Functional inequalities on parameter space). Consider a statistical manifold (Θ, I, \mathbf{p}) . The following inequalities hold.

(i) If $\operatorname{Ric}(\Theta, I, \mathbf{p}) \geq \kappa > 0$, then the Logarithmic Sobolev inequality on parameter space

$$D_{\mathrm{KL}}(p(\theta)\|q) - D_{\mathrm{KL}}(p(\theta^*)\|q) \le \frac{1}{2\kappa} \mathcal{I}(p(\theta)\|q),$$
(14)

holds for any $\theta \in \Theta$.

(ii) If $\operatorname{Ric}(\Theta, I, \mathbf{p}) \ge \kappa > 0$, then the Talagrand inequality on parameter space $\frac{\kappa}{2} d_W(\theta, \theta^*)^2 \le \operatorname{D}_{\operatorname{KL}}(p(\theta) \| q) - \operatorname{D}_{\operatorname{KL}}(p(\theta^*) \| q),$

holds for any $\theta \in \Theta$.

(iii) If $\operatorname{Ric}(\Theta, I, \mathbf{p}) \geq \kappa \in \mathbb{R}$ (κ not necessarily positive), then the HWI inequality on parameter space

$$D_{\mathrm{KL}}(p(\theta)\|q) - D_{\mathrm{KL}}(p(\theta^*)\|q) \le \sqrt{\mathcal{I}(p(\theta)\|q)} d_W(\theta,\theta^*) - \frac{\kappa}{2} d_W(\theta,\theta^*)^2,$$

holds for any $\theta \in \Theta$.

Proof. Here we mainly follow the heuristic arguments in [22]. In finite dimensional parameter space, these approaches are rigorous. We demonstrate the proofs for the completeness of paper.

(i) The proof follows Proposition 9. Consider the Fokker-Planck equation (2) with initial condition $\theta(0) = \theta$. The dissipation along gradient flow of entropy gives

$$\mathcal{I}(p(\theta_t) \| q) = -\frac{d}{dt} D_{\mathrm{KL}}(p(\theta_t) \| q)$$

$$= g_W \big(\operatorname{grad}_W D_{\mathrm{KL}}(p(\theta_t) \| q), \operatorname{grad}_W D_{\mathrm{KL}}(p(\theta_t) \| q) \big).$$
(15)

Since (12) holds, by integrating over time $t \in [0, \infty)$, we have

$$-\frac{d}{dt} \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q) |_{t=0}^{\infty} \ge -2\kappa [\operatorname{D}_{\mathrm{KL}}(p(\theta^*) \| q) - \operatorname{D}_{\mathrm{KL}}(p(\theta_0) \| q)]$$

From (15) and $\theta(0) = \theta$, we have

$$\mathcal{I}(p(\theta)\|q) - \mathcal{I}(p(\theta^*)\|q) \le 2\kappa [\mathrm{D}_{\mathrm{KL}}(p(\theta)\|q) - \mathrm{D}_{\mathrm{KL}}(p(\theta^*)\|q)],$$

where we use the fact $\operatorname{grad}_W \operatorname{D}_{\operatorname{KL}}(p(\theta) || q) = 0$, so that $\mathcal{I}(p(\theta^*) || q) = 0$. It proves the result.

(ii) Consider $\theta(t)$ satisfy the FPE (2) on parameter space with $\theta(0) = \theta$. Since $\operatorname{Ric}(\Theta, I, \mathbf{p}) \geq \kappa > 0$, then $\lim_{t \to \infty} \theta(t) = \theta^*$. Define

$$\Psi(t) = d_W(\theta, \theta(t)) + \sqrt{\frac{2}{\kappa}} \sqrt{\mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q) - \mathcal{D}_{\mathrm{KL}}(p(\theta^*) \| q)}.$$

Thus $\Psi(0) = \sqrt{\frac{2}{\kappa}} \sqrt{D_{\text{KL}}(p(\theta) \| q) - D_{\text{KL}}(p(\theta^*) \| q)}$ and $\Psi(\infty) = \lim_{t \to \infty} \Psi(t) = d_W(\theta, \theta^*)$. We claim that $\Psi(t)$ is nondecreasing. If so, then $\Psi(0) \leq \Psi(\infty)$, which proves the result.

To show $\Psi(t)$ is nondecreasing, we shall prove that

$$\frac{d}{dt}^{+}\Psi(t) = \lim \sup_{h \to 0+} \frac{\Psi(t+h) - \Psi(t)}{h} \le 0.$$

Here we assume $\theta(t) \neq \theta^*$, otherwise $\Psi(t+h) = \Psi(t)$ for any h, which shows the upper derivative zero.

On the one hand, by triangle inequality,

$$|d_W(\theta, \theta_t) - d_W(\theta, \theta_{t+h})| \le d_W(\theta_t, \theta_{t+h}),$$

so that

$$\lim \sup_{h \to 0+} \frac{d_W(\theta_t, \theta_{t+h})}{h} = \sqrt{g_W(\operatorname{grad}_W \operatorname{D}_{\operatorname{KL}}(p(\theta_t) \| q), \operatorname{grad}_W \operatorname{D}_{\operatorname{KL}}(p(\theta_t) \| q))} = \sqrt{\mathcal{I}(p(\theta_t) \| q)}.$$
(16)

On the other hand, since $\theta(t) \neq \theta^*$, then

$$\begin{split} &\sqrt{\frac{2}{\kappa}} \frac{d}{dt} \sqrt{\mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q) - \mathcal{D}_{\mathrm{KL}}(p(\theta^*) \| q)} \\ = &- \frac{g_W(\operatorname{grad}_W \mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q), \operatorname{grad}_W \mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q))}{\sqrt{2\kappa(\mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q) - \mathcal{D}_{\mathrm{KL}}(p(\theta^*) \| q))}} \\ = &- \frac{\mathcal{I}(p(\theta_t) \| q)}{\sqrt{2\kappa(\mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q) - \mathcal{D}_{\mathrm{KL}}(p(\theta^*) \| q))}}. \end{split}$$

From (14), we have

$$\sqrt{\frac{2}{\kappa}} \frac{d}{dt} \sqrt{\mathcal{D}_{\mathrm{KL}}(p(\theta_t) \| q) - \mathcal{D}_{\mathrm{KL}}(p(\theta^*) \| q)} \le -\sqrt{\mathcal{I}(p(\theta_t) \| q)}.$$
(17)

From (16) and (17), we have $\frac{d}{dt}^+ \Psi(t) = \limsup_{h \to 0+} \frac{\Psi(t+h) - \Psi(t)}{h} \leq 0$, which finishes the proof.

(iii) From the definition of $\operatorname{Ric}(\Theta, I, \mathbf{p}) \geq \kappa$, then $\operatorname{Hess}_W \operatorname{D}_{\operatorname{KL}}(p(\theta) || q) \succeq \kappa G_W$. Denote θ_t be a geodesic curve of least energy in manifold (Θ, g_W) , joining $\theta_0 = \theta$ and $\theta_1 = \theta^*$. Thus

$$d_W(\theta, \theta^*) = \sqrt{g_W(\frac{d\theta_t}{dt}, \frac{d\theta_t}{dt})}.$$

From the Taylor expansion on the (Θ, g_W) , we have

$$D_{\rm KL}(p(\theta^*) \| q) = D_{\rm KL}(p(\theta) \| q) + \frac{d}{dt}|_{t=0} D_{\rm KL}(p(\theta_t) \| q) + \int_0^1 (1-t) \frac{d^2}{dt^2} D_{\rm KL}(p(\theta_t) \| q) dt.$$

We note that

$$\begin{aligned} \frac{d}{dt}|_{t=0} \mathcal{D}_{\mathrm{KL}}(p(\theta_t)||q) &= g_W(\operatorname{grad}_W \mathcal{D}_{\mathrm{KL}}(p(\theta_t)||q), \frac{d\theta_t}{dt})|_{t=0} \\ &\geq -\sqrt{g_W(\operatorname{grad}_W \mathcal{D}_{\mathrm{KL}}(p(\theta_t)||q), \operatorname{grad}_W \mathcal{D}_{\mathrm{KL}}(p(\theta_t)||q))}|_{t=0} \sqrt{g_W(\frac{d\theta_t}{dt}, \frac{d\theta_t}{dt})}|_{t=0} \\ &= -\sqrt{\mathcal{I}(p(\theta)||q)} d_W(\theta, \theta^*), \end{aligned}$$

18

and

$$\begin{split} \int_0^1 (1-t) \frac{d^2}{dt^2} \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q) dt &= \int_0^1 (1-t) g_W(\operatorname{Hess}_W \operatorname{D}_{\mathrm{KL}}(p(\theta_t) \| q) \cdot \frac{d\theta_t}{dt}, \frac{d\theta_t}{dt}) dt \\ &\geq \int_0^1 \kappa (1-t) g_W(\frac{d\theta_t}{dt}, \frac{d\theta_t}{dt}) dt \\ &= \frac{\kappa}{2} d_W(\theta, \theta^*)^2. \end{split}$$

Combining the above formulas, we prove the result.

4.4. Computing the Ricci curvature lower bound and convergence rate. In this section, we design an algorithm for Ricci curvature lower bound κ .

We first approximate κ by RIW condition in Theorem 8. In other words, we compute formulas for (9) via

$$\kappa = \text{Smallest eigenvalue of } G_W(\theta)^{-1} \Big\{ G_F(\theta) + \sum_{a \in I} \Big(d_{\theta\theta} p_a(\theta) \log \frac{p_a(\theta)}{q} - \Gamma^{W,a}(\theta) \frac{d}{d_{\theta_a}} \operatorname{D}_{\mathrm{KL}}(p(\theta) \| q) \Big) \Big\}.$$

where $d_{\theta\theta}p_a(\theta)$, $\Gamma^{W,a}(\theta)$, $\frac{d}{d\theta_a} D_{\mathrm{KL}}(p(\theta) || q)$ are computed by numerical differentiation.

In practice, we also compute a uniform convergence rate $K \ge \kappa$ as the smallest ratio of $\frac{d}{dt} D_{\text{KL}}(p(\theta_t) || q)$ and $\frac{d^2}{dt^2} D_{\text{KL}}(p(\theta_t) || q)$ along the gradient flow (8) for any initial conditions. I.e.

$$K = \min_{\theta_0 \in \Theta} \frac{1}{2T} \frac{D_{\rm KL}(p(\theta_{2T}) \| q) - 2 D_{\rm KL}(p(\theta_T) \| q) + D_{\rm KL}(p(\theta_0) \| q)}{D_{\rm KL}(p(\theta_T) \| q) - D_{\rm KL}(p(\theta_0) \| q)},$$

where T is a given short time, θ_T is the solution of (11) with initial condition θ_0 . Whenever K > 0, it is always the tight bound for functional inequalities in Proposition 10.

Convergence rate

Input: Sample initial conditions $\{\theta_0^s\}_{s=1}^{|S|}$; Target distribution q; A suitable initial step size h > 0; A short terminal time T > 0. **Output**: Approximation K of the uniform convergence rate;

$$\begin{split} & \text{for } s \in \{1, \cdots, |S|\} \\ & \text{for } k = 1, 2, \dots, 2T/h \\ & \theta_{k+1}^s = \theta_k^s - hG_W(\theta_k^s)^{-1} \nabla_\theta \operatorname{D}_{\mathrm{KL}}(p(\theta_k^s) \| q) \ ; \\ & \text{end} \\ & \text{end} \\ & K = \min_{s \in \{1, \cdots, |S|\}} \frac{1}{2T} \frac{\operatorname{D}_{\mathrm{KL}}(p(\theta_{2T}^s) \| q) - 2\operatorname{D}_{\mathrm{KL}}(p(\theta_T^s) \| q) + \operatorname{D}_{\mathrm{KL}}(p(\theta_0^s) \| q)}{\operatorname{D}_{\mathrm{KL}}(p(\theta_T^s) \| q) - \operatorname{D}_{\mathrm{KL}}(p(\theta_0^s) \| q)}. \end{split}$$

19

5. Examples

In this section, we illustrate some of the concepts introduced in the previous sections by means of evaluating them on a simple class of exponential family models. We illustrate the effects from the choice of the ground metric on sample space in relation to the choice of the statistical model, and the relationships between the Ricci curvature lower bound and the rates of convergence in learning.

Example 1 (Ricci curvature for a one-dimensional exponential family on three states). We study how the Ricci curvature changes with the choice of a probability model and with the choice of the ground metric on sample space. In order to obtain a picture as complete as possible, we consider the small setting of three states and one dimensional exponential families.

Consider the sample space $I = \{1, 2, 3\}$ with a fully connected graph with edges $E = \{(1, 2), (2, 3), (1, 3)\}$, and weights $\omega = (\omega_{12}, \omega_{23}, \omega_{13})$. The probability simplex is a triangle

$$\mathcal{P}(I) = \left\{ (p_i)_{i=1}^3 \in \mathbb{R}^3 : \sum_{i=1}^3 p_i = 1, \quad p_i \ge 0 \right\}.$$

We consider statistical manifolds of the form

$$p(\theta) = \frac{1}{Z(\theta)} (e^{\theta c_1}, e^{\theta c_2}, e^{\theta c_3}),$$

with sufficient statistic $c = (c_1, c_2, c_3) \in \mathbb{R}^3$, parameter $\theta \in \Theta = [\theta_{\min}, \theta_{\max}] \subset \mathbb{R}^1$, and partition function $Z(\theta) = \sum_{i=1}^3 e^{\theta c_i}$. These are exponential families specified by the choice of the sufficient statistic c. Here, addition of constants is immaterial. Multiplicative scaling by non-zero numbers does not change the model. For better comparability, we always choose c to have norm one.

In particular, these models can be indexed by the projective line, which for simplicity we can represent by a half circle, or an angle.

We fix a uniform reference measure $q = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. The KL divergence then takes the form

$$D_{KL}(p||q) = \sum_{i=1}^{3} p_i \log \frac{p_i}{q_i} = \sum_{i=1}^{3} p_i \log p_i + \log 3.$$

We evaluate the Ricci curvature lower bound for 30 different exponential families and 10 different choices of the ground metric. We choose the sufficient statistics as evenly spaced points on a radius 1 half circle, and set the parameter domain as $\Theta = [-2, 2]$.

The results are shown in Figure 4. The left panel estimates K as the minimum rate of convergence of the Wasserstein gradient flow of the KL divergence, over a grid of 10 different initial conditions on the parameter domain. As can be seen, the convergence is faster, the better ω connects the end points of the exponential family.



FIGURE 4. Lower bound on the Ricci curvature for one-dimensional exponential families on three states. Each simplex corresponds to a different choice of $\omega = (\omega_{12}, \omega_{23}, \omega_{13})$, indicated at the bottom. Within each simplex there are 30 different exponential families (which are curves) with sufficient statistics of norm one and parameter domain $\Theta = [-1, 1]$. The color of each exponential family corresponds to the value of K estimated as the minimum convergence rate (left panel), and the value of κ as the minimum eigenvalue of the Hessian (right panel), over the parameter domain. Blue corresponds to lower and yellow to higher values. We give a direct comparison of Kand κ in Figure 5.

The right panel estimates κ as the minimum eigenvalue of the Hessian operator of the KL divergence over a grid of parameter values in the domain. Figure 5 gives a direct comparison of the estimates obtained from convergence rates and the Hessian. As can be seen, the Hessian is always a lower bound of the convergence rate, which reflects Proposition 9.

If the parameter domain is smaller, the Hessian gives a closer bound to the rate of convergence. If, on the contrary, the parameter domain is larger, the gaps between the Hessian and the convergence rates tend to be larger. Larger parameters correspond to distributions closer to the boundary of the simplex. We illustrate these effects in the Appendix, where we provide figures with different choices of Θ (Figures 6 and 7), and also comparing the Hessian and rates of convergence at individual parameter values (Figure 8).

6. DISCUSSION

To summarize, we introduced a notion of Ricci curvature lower bound for parametric statistical models and illustrated its possible relevance in the context of parameter estimation and learning. This notion is based on the geodesic convexity of the KL divergence in Wasserstein geometry. Following the program from [16], we hope that this paper continues to strengthen the interactions between information geometry and Wasserstein geometry.



FIGURE 5. This figure compares the values of K and κ from Figure 4. Each subplot corresponds to one choice of ω , indicated at the top, with x axis corresponding to the 30 different exponential families. As can be seen, the curvature κ obtained as the smallest Hessian eigenvalue (red) is, indeed, always a lower bound of the convergence rate K (blue).

The Ricci curvature lower bound depends on the target distribution, the statistical model, and the ground metric on sample space. We think that this notion can serve to capture the general properties of learning in different models, and hence that it can serve to guide the design of statistical models (e.g., the graph of a graphical model or the connectivity structure of a neural network) and the ground metric. Our experiments show that an adequate choice of the two, in conjunction, can significantly increase the rates of convergence in learning. On the other hand, the Ricci curvature depends on both, the information and the Wasserstein metric tensors. An interesting question arises; namely to find the statistical interpolation of such a connection.

We note that the Ricci curvature lower bound is a global notion over the probability model. This is advantageous to provide a uniform analysis, but it can also lead to difficulties, especially when the models include points near the boundary of the simplex, where the behavior is not as regular. Our experiments indicate that restricting the parameter domain to a region bounded away from the boundary of the simplex allows us to closely track the rates of convergence. Another challenge is that, being a global quantity, the computation can be challenging. Nonetheless, we point out that computing the curvature in terms of the Hessian is much cheaper than estimating the learning rates empirically. We have focused on discrete sample spaces, which allowed us to obtain an intuitive and transparent picture of the relationships that derive form this theory. However, we expect that the derivations extend naturally to the case of continuous sample spaces.

Another interesting line of investigation is the following. Our definitions are based on the KL divergence and the Wasserstein and Fisher metric tensors. In principle, it is possible to derive analogous definitions for other metric structures. In particular, one can consider the family of f-divergences. Such an analysis could allow us to compare different learning paradigms.

References

- [1] S. Amari. *Information Geometry and Its Applications*. Number volume 194 in Applied mathematical sciences. Springer, Japan, 2016.
- [2] N. Ay, J. Jost, H. V. Lê, and L. J. Schwachhöfer. Information Geometry. Springer, Cham, 2017.
- [3] D. Bakry and M. Émery. Diffusions hypercontractives. Séminaire de probabilités de Strasbourg, 19:177–206, 1985.
- [4] Y. Chen and W. Li. Natural gradient in Wasserstein statistical manifold. 2018.
- [5] I. Csiszár and P. C. Shields. Information theory and statistics: A tutorial. Commun. Inf. Theory, 1(4):417–528, Dec. 2004.
- [6] M. Erbar and M. Fathi. Poincaré, modified logarithmic Sobolev and isoperimetric inequalities for Markov chains with non-negative Ricci curvature. *Journal of Functional Analysis*, 274(11):3056–3089, 2018.
- [7] M. Erbar, C. Henderson, G. Menz, and P. Tetali. Ricci curvature bounds for weakly interacting Markov chains. *Electronic Journal of Probability*, 22, 2017.
- [8] M. Erbar and E. Kopfer. Super Ricci flows for weighted graphs. arXiv:1805.06703 [math], 2018.
- M. Erbar and J. Maas. Ricci Curvature of Finite Markov Chains via Convexity of the Entropy. Archive for Rational Mechanics and Analysis, 206(3):997–1038, 2012.
- [10] M. Erbar, J. Maas, and P. Tetali. Discrete Ricci Curvature bounds for Bernoulli-Laplace and Random Transposition models. Annales de la faculté des sciences de Toulouse Mathématiques, 24(4):781–800, 2015.
- M. Fathi and J. Maas. Entropic Ricci curvature bounds for discrete interacting systems. The Annals of Applied Probability, 26(3):1774–1806, 2016.
- [12] B. Hua, J. Jost, and S. Liu. Geometric Analysis Aspects of Infinite Semiplanar Graphs with Nonnegative Curvature. Journal f
 ür die reine und angewandte Mathematik (Crelles Journal), 2015(700):1–36, 2015.
- [13] J. Jost and S. Liu. Ollivier's Ricci Curvature, Local Clustering and Curvature-Dimension Inequalities on Graphs. Discrete Comput. Geom., 51(2):300–322, 2014.
- [14] J. D. Lafferty. The Density Manifold and Configuration Space Quantization. Transactions of the American Mathematical Society, 305(2):699–741, 1988.
- [15] W. Li. Geometry of probability simplex via optimal transport. arXiv:1803.06360 [math], 2018.
- [16] W. Li and G. Montúfar. Natural gradient via optimal transport I. arXiv:1803.07033, 2018.
- [17] Y. Lin, L. Lu, and S.-T. Yau. Ricci Curvature of Graphs. Tohoku Mathematical Journal, 63(4):605–627, 2011.
- [18] Y. Lin and S.-T. Yau. Ricci Curvature and Eigenvalue Estimate on Locally Finite Graphs. Mathematical Research Letters, 17(2):343–356, 2010.
- [19] J. Lott and C. Villani. Ricci Curvature for Metric-Measure Spaces via Optimal Transport. Annals of Mathematics, 169(3):903–991, 2009.
- [20] Y. Ollivier. Ricci Curvature of Markov Chains on Metric Spaces. Journal of Functional Analysis, 256(3):810–864, 2009.
- [21] Y. Ollivier and C. Villani. A Curved Brunn–Minkowski Inequality on the Discrete Hypercube, Or: What Is the Ricci Curvature of the Discrete Hypercube? SIAM Journal on Discrete Mathematics, 26(3):983–996, 2012.
- [22] F. Otto and C. Villani. Generalization of an Inequality by Talagrand and Links with the Logarithmic Sobolev Inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.
- [23] K.-T. Sturm. On the Geometry of Metric Measure Spaces. Acta Mathematica, 196(1):65–131, 2006.
- [24] C. Villani. Optimal Transport: Old and New. Number 338 in Grundlehren der mathematischen Wissenschaften. Springer, Berlin, 2009.



Appendix A. Additional figures to Example 1

FIGURE 6. Similar to Figure 4 but with $\Theta = [-1/2, 1/2]$. Note how on this tight parameter domain around $\theta = 0$ (the value of the reference measure), the Ricci curvature lower bound gives a very close lower bound on the minimum rate of convergence for each of the models. The middle shows the direct comparison of the two values across the 30 exponential families. The minimum rate of convergence is shown in blue, and the Hessian in red.



FIGURE 7. Similar to Figures 4, but with a larger parameter domain $\Theta = [-4, 4]$. On this relatively large parameter domain, the models contain points close to the boundary of the simplex, where the Hessian (and the Ricci curvature) can have large oscillations. In turn, we observe larger gaps to the minimum rate of convergence, compared with Figure 6.



FIGURE 8. Convergence rates and minimum Hessian eigenvalue at individual parameter choices. Here we fixed the ground metric $\omega = (\omega_{12}, \omega_{23}, \omega_{13}) = (1/2, 1/2, 0)$. Each subplot corresponds to one exponential family, with sufficient statistic indicated at the top. Within a region around $\theta = 0$ (the value of the reference measure), the minimum of the Hessian is closer to the convergence rates. In fact, the Hessian eigenvalue intersects the rate of convergence at $\theta = 0$. The Hessian at $\theta = 0$ is the asymptotic rate of convergence. The lower row zooms in the y axis of the upper row. For these exponential families, the convergence rates do not vary much across choices of the initial parameter value.