



Wasserstein information matrix

Wuchen Li¹  · Jiaxi Zhao¹

Received: 29 November 2020 / Revised: 19 September 2022 / Accepted: 22 January 2023 /

Published online: 14 February 2023

© The Author(s), under exclusive licence to Springer Nature Singapore Pte Ltd. 2023

Abstract

We study information matrices for statistical models by the L^2 -Wasserstein metric. We call them Wasserstein information matrices (WIMs), which are analogs of classical Fisher information matrices. We introduce Wasserstein score functions and study covariance operators in statistical models. Using them, we establish Wasserstein–Cramer–Rao bounds for estimations and explore their comparisons with classical results. We next consider the asymptotic behaviors and efficiency of estimators. We derive the online asymptotic efficiency for Wasserstein natural gradient. Besides, we establish a Poincaré efficiency for Wasserstein natural gradient of maximal likelihood estimation. Several analytical examples of WIMs are presented, including location-scale families, independent families, rectified linear unit (ReLU) generative models.

Keywords Wasserstein information matrix · Wasserstein score function · Wasserstein–Cramer–Rao inequality · Wasserstein online efficiency · Poincaré efficiency

1 Introduction

Fisher information matrix plays essential roles in statistics, physics, and differential geometry with applications in machine learning [1, 2, 6, 10, 12]. In statistics, it is a fundamental quantity for the estimation theory, including both design and analysis of estimators. Especially, the maximal likelihood principle is a well-known example, which connects the Fisher information matrix to another concept, named score

Communicated by Nihat Ay.

✉ Wuchen Li
wuchen@mailbox.sc.edu

Jiaxi Zhao
zjx98math@gmail.com

¹ University of South Carolina, Columbia, SC, USA

functions. They frequently arise in statistical efficiency and sufficiency problems, especially for Cramer–Rao bound and Fisher-efficiency.

Fisher information matrix is closely related to Fisher-Rao metric in information geometry [3]. It uses the Fisher information matrix to study divergence functions and their invariance properties [3]. Furthermore, the Fisher information matrix is also useful for statistical learning problems. In particular, the natural gradient method [2] rectifies the gradient direction by the Fisher information matrix. It is shown that the Fisher natural gradient method is asymptotically online Fisher-efficient.

Besides, optimal transport introduces the other metric in probability space [21, 28, 29], in particular, Wasserstein-2 metric [14, 25]. Different from information geometry, it encodes the geometry of sample space into the definition of metric in probability space. Nowadays, it is known that the Wasserstein metric intrinsically connects the Kullback–Leibler (KL) divergence with Fisher information functional [25], known as de Bruijn identities [31]. Many concentration inequalities such as log-Sobolev inequalities and Poincaré inequalities arise naturally [26].

Despite various studies of optimal transport in full probability space, not much is known in parametric models, which play crucial roles in parametric statistics and estimation. Fundamental questions arise: *Is there a statistical theory based on optimal transport? Compared to Fisher information matrices and Fisher statistics, what are counterparts of information matrices, score functions, Cramer–Rao bounds, and online efficiencies of natural gradient methods in Wasserstein statistics? Moreover, can this theory provide statistical tools for machine learning models, especially for generative models?*

In this paper, following key ideas in [15], we positively answer the above questions by introducing a Wasserstein information matrix (WIM). We derive the WIM by pulling back the Wasserstein metric from a immense probability space to finite-dimensional parametric statistical models [16, 17]. We show that the WIM defines Wasserstein score functions with a Wasserstein covariance operator. Based on them, a Wasserstein–Cramer–Rao bound is derived. Furthermore, combining WIM with Wasserstein score functions, we recover an asymptotic efficiency property of the online Wasserstein natural gradient methods.

Meanwhile, by comparing both Wasserstein and Fisher information matrices, we naturally prove several concentration inequalities such as log-Sobolev inequalities and Poincaré inequalities within statistical models. Extending the study in full probability space, we further study the Ricci-Information-Wasserstein (RIW) criterion for log-Sobolev inequalities and Poincaré inequalities in statistical models. We provide several examples in analytic probability families. Those geometric calculations establish a new efficiency property named Poincaré efficiency. This is concerned with dynamics where the Wasserstein natural gradient works on Fisher score functions (related to maximal likelihood estimators). We prove convergence rate analysis for these dynamics. Several numerical experiments are provided to confirm our conclusions.

Lastly, we demonstrate that the WIM provides a clear statistical theory for complicated models coming from machine learning approaches, especially generative models. For example, we carefully study a one-dimensional probability family generated by push-forward maps based on the ReLU function. We demonstrate that the WIM still exists in this family while the classical Fisher information matrix does not

exist. In other words, it is suitable to introduce a statistical theory based on WIMs. It can be a theoretical background for models in machine learning.

In literature, many works have attempted to use tools from optimal transport and information geometry to study statistical problems. [7] designs new estimators for parametric inference using Wasserstein distance. This idea is utilized in approximating Bayesian computation. The authors apply Wasserstein distance to measure the similarity between synthetic and observed data sets. Compared to them, we focus on the study of estimation and efficiency of WIMs and expect it could have potential interplay with Wasserstein estimators. In [9], they design a generalized information matrix based on a maximum mean discrepancy while we majorly focus on information matrices generated by the Wasserstein metric and study related statistical properties. Most closely, [27] defines a Wasserstein covariance by applying a closed-form formula for the one-dimensional Wasserstein metric. Our approach further extends this idea into general parametric models. Also, [30] defines several new divergence functions by combining knowledge from both optimal transport and information geometry. Here we focus on the statistical properties of WIMs in statistical models. Furthermore, Wasserstein natural gradient method has been widely studied in optimization techniques with machine learning applications [5, 11, 19, 20, 22]. Here we focus on statistical theory and study its associated online efficiency. Compared with classical online Fisher-efficiency results in [2, 24], our results can deal with general information matrices. In particular, we discover a new efficiency property named Poincaré efficiency. In the same spirit, [4] further studies the asymptotic behavior of Wasserstein estimators using WIM on the one-dimensional location-scale family.

The paper is organized as follows. In Sect. 2, we establish the definition of the WIM. We present it analytically for several well-known probability families. We provide an explicit example of WIMs for ReLU generative models. Under this model, we show that the WIM exists while the Fisher information matrix does not exist. In Sect. 3, with the introduction of the Wasserstein covariance, the Wasserstein–Cramer–Rao inequality is established. In Sect. 4, we derive both Wasserstein efficiency and Poincaré efficiency for Wasserstein natural gradient method.

2 Wasserstein information matrix and score functions

In this section, we present Wasserstein information matrices (WIMs) and score functions. Several analytical studies are presented.

Let $\mathcal{X} \subset \mathbb{R}^n$ be a compact manifold with smooth Riemannian structure. This is the sample space associated with statistical inference problems. Let $\mathcal{P}^{2,\infty}(\mathcal{X})$ denote the space of probability distributions over \mathcal{X} which are absolutely continuous w.r.t. the Lebesgue measure on \mathcal{X} and have smooth positive density functions. Moreover, one further requires these distributions to have bounded second moments, which can be calculated via the embedding of \mathcal{X} into \mathbb{R}^n . For those who do not have much familiarity with the Riemannian geometry, \mathcal{X} can be simply regarded as the whole Euclidean space \mathbb{R}^n . Given a metric tensor g on $\mathcal{P}^{2,\infty}(\mathcal{X})$, we call $(\mathcal{P}^{2,\infty}(\mathcal{X}), g)$ density manifold. It is a manifold in the sense of [13, 21]. Consider a parameter space $\Theta \subset \mathbb{R}^d$ and a parameterization function

Table 1 In this table, we present Wasserstein, Fisher information matrices for probability families

Probability family	Wasserstein information matrix	Fisher information matrix
Uniform:		$G_F(a, b)$ not well-defined
$p(x; a, b) = \frac{1}{b-a} \mathbf{1}_{(a,b)}(x)$	$G_W(a, b) = \frac{1}{3} \begin{pmatrix} 1 & \frac{1}{2} \\ \frac{1}{2} & 1 \end{pmatrix}$	
Gaussian:		$G_F(\mu, \sigma) = \begin{pmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{2}{\sigma^2} \end{pmatrix}$
Exponential:	$G_W(\mu, \sigma) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$	$G_F(m, \lambda)$ not well-defined
$p(x; m, \lambda) = \lambda e^{-\lambda(x-m)}$	$G_W(m, \lambda) = \begin{pmatrix} 1 & \frac{1}{\lambda^2} \\ \frac{1}{\lambda^2} & \lambda^4 \end{pmatrix}$	
Laplacian:	$G_W(m, \lambda) = \begin{pmatrix} 1 & 0 \\ 0 & \frac{2}{\lambda^4} \end{pmatrix}$	$G_F(m, \lambda) = \begin{pmatrix} \lambda^2 & 0 \\ 0 & \frac{1}{\lambda^2} \end{pmatrix}$
Location-scale:	$G_W(\lambda, m) = \begin{pmatrix} \mathbb{E}_{\lambda, m} x^2 - 2m\mathbb{E}_{\lambda, m} x + m^2 & 0 \\ 0 & \frac{\lambda^2}{\lambda^4 p} \end{pmatrix}$	$G_F(\lambda, m) = \begin{pmatrix} \frac{1}{\lambda^2} \left(1 + \int_{\mathbb{R}} \left(\frac{(x-m)^2 p^2}{\lambda^2 p} + \frac{(x-m) p'}{\lambda} \right) dx \right) \int_{\mathbb{R}} \frac{(x-m) p'}{\lambda^3 p} dx \\ \int_{\mathbb{R}} \frac{(x-m) p^2}{\lambda^3 p} dx \end{pmatrix}$
Independent:	$G_W(\theta) = G_W^1(\theta) + G_W^2(\theta)$	$G_F(\theta) = G_F^1(\theta) + G_F^2(\theta)$
$p(x, y; \theta) = p_1(x; \theta)p_2(y; \theta)$		
ReLU push-forward:	$G_W(\theta) = F(\theta)$,	
$p(x; \theta) = f_{\theta*}p(x)$,		$G_F(\theta)$ not well-defined
f_{θ} θ -parameterized ReLUs, Ex. 8.	F cdf of $p(x)$, $F(y) = \int_{-\infty}^y p(x) dx$.	

Table 2 In this table, we continue to list the entropy functional, the Fisher information functionals, log-Sobolev inequalities for probability families

Family	Entropy functional	Fisher-information functional	Log-Sobolev inequality($LSI(\alpha)$)
Gaussian	$H(p_{\mu,\sigma}) = -\frac{1}{2} \log 2\pi - \log \sigma - \frac{1}{2},$ $\text{D}\text{KL}(p_{\mu,\sigma} \ p_{\mu_*,\sigma_*}) = -\log \sigma + \log \sigma_* - \frac{1}{2}$ $+ \frac{\sigma^2 + (\mu - \mu_*)^2}{2\sigma_*^2}.$	$\tilde{I}(p_{\mu,\sigma}) = \frac{1}{\sigma^2},$ $\tilde{I}(p_{\mu,\sigma} \ p_{\mu_*,\sigma_*}) = \frac{(\mu - \mu_*)^2}{4\sigma_*^4} + \left(-\frac{1}{\sigma} + \frac{\sigma}{\sigma_*}\right)^2.$	$\text{D}\text{KL}(p_{\mu,\sigma} \ p_{\mu_*,\sigma_*}) < \frac{1}{2\alpha} \tilde{I}(p_{\mu,\sigma} \ p_{\mu_*,\sigma_*}),$ $\mu, \sigma > 0.$
Laplacian	$H(p_{\lambda,m}) = -1 + \log \lambda - \log 2,$ $\text{D}\text{KL}(p_{m,\lambda} \ p_{m_*,\lambda_*}) = -1 + \log \lambda - \log \lambda_*$ $+ \lambda_* m - m_* + \frac{\lambda_* e^{-\lambda m - m_* }}{\lambda}.$	$\tilde{I}(p_{\lambda,m}) = \frac{\lambda^2}{2},$ $\tilde{I}(p_{m,\lambda} \ p_{m_*,\lambda_*}) = \lambda_*^2 \left(1 - e^{-\lambda m - m_* }\right)^2$ $+ \frac{\left((\lambda_* m - m_* + 1) \lambda_* e^{-\lambda_* m - m_* } - \lambda\right)^2}{2}.$	$\text{D}\text{KL}(p_{\lambda,m} \ p_{\lambda_*,m_*}) < \frac{1}{2\alpha} \tilde{I}(p_{\lambda,m} \ p_{\lambda_*,m_*}),$ $m \in \mathbb{R}, \lambda > 0.$

$$p: \Theta \rightarrow \mathcal{P}^{2,\infty}(\mathcal{X}), \quad \theta \mapsto p_\theta$$

which can also be viewed as $p: \mathcal{X} \times \Theta \rightarrow \mathbb{R}$, since we assume the distributions have smooth density functions. The image of Θ under the mapping p , i.e. $p(\Theta)$ is named a statistical model. Suppose that f, g are two vector-valued functions on \mathcal{X} , we denote by $\langle f, h \rangle = \int_{\mathcal{X}} (f(x), h(x)) dx$ the $L^2(\mathcal{X})$ inner product, where dx refers to the Lebesgue measure on \mathcal{X} . And we denote by $(v, w) = v \cdot w$ the (pointwise) Euclidean inner product of two vectors.

2.1 Information matrix

From now on, we will omit the superscript of the density manifold and simply denote it by \mathcal{P} . Since probability distributions are restricted to a special family, one can identify the tangent space of density manifold \mathcal{P} at any particular distribution ρ with $C_0^\infty(\mathcal{X})$:

$$T_p \mathcal{P}(\mathcal{X}) \simeq C_0^\infty(\mathcal{X}) = \left\{ f \in C^\infty(\mathcal{X}) \mid \int_{\mathcal{X}} f dx = 0 \right\},$$

where $C^\infty(\mathcal{X})$ is the space of smooth functions on the \mathcal{X} . And functions in its dual space are determined up to an additive constant, i.e. $f \sim h$ if $f = h + a, a \in \mathbb{R}$. We denote the equivalent class by $[f]$. And the pairing between tangent spaces and cotangent spaces is merely $\langle f, h \rangle = \int_{\mathcal{X}} f h dx$. Geometrically speaking, since the metric tensor is a positive definite bilinear form, using the associated invertible operator g to rise(resp. lower) indices, there exists a canonical isomorphism from tangent(resp. cotangent) space to cotangent(resp. tangent) space, namely

$$\begin{aligned} g_p : T_p \mathcal{P}(\mathcal{X}) &\simeq T_p^* \mathcal{P}(\mathcal{X}), \quad f \mapsto [g_p(f)], \\ \text{s.t. } g_p(f, h) &= \langle g_p(f), h \rangle \quad \forall h \in T_p \mathcal{P}(\mathcal{X}). \end{aligned}$$

The reader should recognize that in above formula g_p represents metric tensor in $g_p(\cdot, \cdot)$, which provides the inner product of two tangent (cotangent) vectors. While it represents the invertible transformation between tangent and cotangent space in $g_p(f)$. Thus, the metric tensor can actually be viewed as an operator between these two spaces. Notice that tangent spaces of a statistical model Θ can be viewed as subspaces of that of density manifold, i.e. we have inclusion

$$T_p \Theta \hookrightarrow T_p \mathcal{P}(\mathcal{X}).$$

While taking the dual of this inclusion we get projection from the cotangent space of the density manifold to that of the statistical model

$$T_p^* \mathcal{P}(\mathcal{X}) \rightarrow T_p^* \Theta.$$

We first review metric tensors on parameter space and connect them with information matrices.

Definition 1 (Statistical information matrix) Consider the density manifold $(\mathcal{P}^{2,\infty}(\mathcal{X}), g)$ with a metric tensor g , and a smoothly parametrized statistical model p_θ with parameter $\theta \in \Theta \subset \mathbb{R}^d$. Then the pull-back metric $G \in \mathbb{R}^{d \times d}$ of g onto this parameter space Θ is given by

$$G(\theta) = g_{p_\theta}(\nabla_\theta p_\theta, \nabla_\theta p_\theta) = \left\langle \nabla_\theta p_\theta, g_{p_\theta} \nabla_\theta p_\theta \right\rangle.$$

Denote $G(\theta) = (G(\theta)_{ij})_{1 \leq i, j \leq d}$, then

$$G(\theta)_{ij} = g_{p_\theta}(\partial_{\theta_i} p_\theta, \partial_{\theta_j} p_\theta) = \int_{\mathcal{X}} \frac{\partial}{\partial \theta_i} p(x; \theta) \left(g_{p_\theta} \frac{\partial}{\partial \theta_j} p \right)(x; \theta) dx.$$

We name g statistical metric, and call G statistical information matrix.

Next, we reinterpret the metric tensor in the dual coordinates, i.e. cotangent space.

Definition 2 (Score function) Denote $\Phi_i : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$, $i = 1, \dots, n$ satisfying

$$\Phi_i(x; \theta) = \left[g_p \left(\frac{\partial}{\partial \theta_i} p(x; \theta) \right) \right].$$

We call Φ_i , $i = 1, 2, \dots, n$ score functions associated with the statistical information matrix G and are equivalent classes in $C(\mathcal{X})/\mathbb{R}$. The representatives in equivalent classes are determined by the following normalization condition

$$\mathbb{E}_{p_\theta} \Phi_i = 0, \quad i = 1, \dots, n. \quad (1)$$

Then the statistical information matrix satisfies

$$G(\theta)_{ij} = \int_{\mathcal{X}} \Phi_i(x; \theta) \left(g_{p_\theta}^{-1} \Phi_j \right)(x; \theta) dx.$$

Remark 1 The existence of the score function depends on the metric tensor g_p . Meanwhile, we assume that the inverse transformation g_p^{-1} is well-defined and can be used in calculating the inner product in the definition.

Remark 2 The normalization condition is an artificial condition. It fixes a representative for the score function in the equivalent class. And we assume that score functions are always square integrable w.r.t. p_θ , which is satisfied by most classical parametric families. This postulation is assumed in the following paper.

In conclusion, we introduce two formulations of metric tensor based on tangent and cotangent vector, whose relation can be summarized via

$$\begin{aligned} G(\theta)_{ij} &= \left\langle \nabla_{\theta_i} p_\theta, g_{p_\theta} \nabla_{\theta_j} p_\theta \right\rangle \\ &= \left\langle g_{p_\theta} \nabla_{\theta_i} p_\theta, g_{p_\theta}^{-1} g_{p_\theta} \nabla_{\theta_j} p_\theta \right\rangle \\ &= \left\langle \Phi_i, g_{p_\theta}^{-1} \Phi_j \right\rangle. \end{aligned}$$

Example 1 One important choice of metric is the Fisher-Rao metric, i.e.

$$g_p : T_p \mathcal{P}(\mathcal{X}) \simeq T_p^* \mathcal{P}(\mathcal{X}), \quad f \mapsto \left[\frac{f}{p} \right],$$

$$g_p^{-1} : T_p^* \mathcal{P}(\mathcal{X}) \simeq T_p \mathcal{P}(\mathcal{X}), \quad [f] \mapsto p(f - E_p f).$$

From the assumption, every probability distribution in \mathcal{P} has positive density, the above operator is well-defined. In this case, the statistical information matrix satisfies

$$G_F(\theta)_{ij} = \left\langle \frac{\partial}{\partial \theta_i} p_\theta, \frac{1}{p_\theta} \frac{\partial}{\partial \theta_j} p_\theta \right\rangle = \int_{\mathcal{X}} \frac{\frac{\partial}{\partial \theta_i} p(x; \theta) \frac{\partial}{\partial \theta_j} p(x; \theta)}{p(x; \theta)} dx.$$

And score functions of Fisher information matrix form

$$\Phi_i^F(x; \theta) = \frac{1}{p(x; \theta)} \frac{\partial}{\partial \theta_i} p(x; \theta) = \frac{\partial}{\partial \theta_i} \log p(x; \theta),$$

where the normalization condition holds automatically. In terms of score functions, the Fisher information matrix forms

$$G_F(\theta)_{ij} = \int_{\mathcal{X}} \Phi_i^F(x; \theta) \left(g_F(p)^{-1} \Phi_j^F \right)(x; \theta) dx$$

$$= \int_{\mathcal{X}} \frac{\partial}{\partial \theta_i} \log p(x; \theta) \frac{\partial}{\partial \theta_j} \log p(x; \theta) p(x; \theta) dx$$

$$= \mathbb{E}_{p_\theta} \left(\frac{\partial}{\partial \theta_i} \log p(x; \theta) \frac{\partial}{\partial \theta_j} \log p(x; \theta) \right).$$

To avoid ambiguity, we use $g_F(p)$, $g_F(p)^{-1}$ to denote the transformations associated with the Fisher metric. In literature, $\Phi_i^F(x; \theta) = \frac{\partial}{\partial \theta_i} \log p(x; \theta)$ is exactly the (Fisher) *score function*; while $G_F(\theta)$ is the *Fisher information matrix*. They play important roles in estimation, efficiency and Cramer–Rao bound.

Remark 3 The definition of Fisher score functions can be given in classical statistics as the gradient of the log-likelihood function w.r.t. parameters θ . We bring about another interpretation as an object on cotangent space associated with the Fisher-Rao metric. This fact is not limited to Fisher metric. Suppose we have a family of canonical tangent vector fields $\frac{\partial}{\partial \theta_i} p_\theta$ on statistical models, whenever there is a metric g_{p_θ} on this manifold, we can define score functions associated with it as

$$\Phi_i(x; \theta) = g(p_\theta) \frac{\partial}{\partial \theta_i} p_\theta(x; \theta).$$

From above facts, we observe that statistical concepts are related to the metric tensor in density manifold pull-back onto parameter spaces. In particular, classical statistics relates to the Fisher-Rao metric. The pull-back metric tensor forms an information

matrix while dual variables define score functions. In this paper, we derive these notations in the other important metric, namely Wasserstein-2 metric.

2.2 Wasserstein information matrix

The other statistical metric, namely Wasserstein metric tensor forms

$$g_W(p) = (-\Delta_p)^{-1}, \quad \text{where } \Delta_p = \nabla \cdot (p \nabla).$$

Here Δ_p is a weighted elliptic operator. When p satisfies suitable conditions, such as smooth and positive everywhere, a known fact that the operators Δ_p^{-1} exists, given by

$$\begin{aligned} \Delta_p^{-1} : C_0^\infty(\mathcal{X}) &\rightarrow C^\infty(\mathcal{X})/\mathbb{R}; \\ \Delta_p : C^\infty(\mathcal{X})/\mathbb{R} &\rightarrow C_0^\infty(\mathcal{X}). \end{aligned}$$

The pull-back G_W of g_W is given by

$$G_W(\theta)_{ij} = \langle \frac{\partial}{\partial \theta_i} p_\theta, (-\Delta_{p_\theta})^{-1} \frac{\partial}{\partial \theta_j} p_\theta \rangle.$$

Similar to the Fisher information matrix, we can rewrite G_W by dual coordinates. Denote

$$\Phi_i^W(x; \theta) = (-\Delta_{p_\theta})^{-1} \frac{\partial}{\partial \theta_i} p(x; \theta).$$

Then

$$\begin{aligned} G_W(\theta)_{ij} &= \left\langle \frac{\partial}{\partial \theta_i} p_\theta, (-\Delta_{p_\theta})^{-1} \frac{\partial}{\partial \theta_j} p_\theta \right\rangle \\ &= \left\langle \Phi_i^W, (-\Delta_{p_\theta}) \Phi_j^W \right\rangle \\ &= \int_{\mathcal{X}} (\nabla_x \Phi_i^W(x; \theta), \nabla_x \Phi_j^W(x; \theta)) p(x; \theta) dx \\ &\equiv \mathbb{E}_{p_\theta} \left[\nabla_x \Phi_i^W(x; \theta) \cdot \nabla_x \Phi_j^W(x; \theta) \right], \end{aligned}$$

where the last equality holds by integration by parts w.r.t. x .

We summarize the above fact into the following definition.

Definition 3 (Wasserstein information matrix & score function) Denote $G_W(\theta) \in \mathbb{R}^{d \times d}$ as

$$G_W(\theta)_{ij} = \mathbb{E}_{p_\theta} \left[\nabla_x \Phi_i^W(x; \theta) \cdot \nabla_x \Phi_j^W(x; \theta) \right],$$

where \cdot refers to the inner product of vector and $\Phi_i^W: \mathcal{X} \times \Theta \rightarrow \mathbb{R}$ satisfies

$$-\nabla_x \cdot (p(x; \theta) \nabla_x \Phi_i^W(x; \theta)) = \frac{\partial}{\partial \theta_i} p(x; \theta), \quad \mathbb{E}_{p_\theta} \Phi_i^W = 0, \quad i = 1, 2, \dots, d.$$

We name functions $\Phi_i^W(x; \theta) = \left((-\Delta_{p_\theta})^{-1} \frac{\partial}{\partial \theta_i} p_\theta \right)(x; \theta)$ *Wasserstein score functions*, and call the matrix $G_W(\theta)$ the *Wasserstein information matrix*.

Remark 4 It is little bit confusing that both $\partial_x, \partial_\theta$ appear in the definition of WIM and score functions. One simple method to distinguish between them is to look at the function they act on. ∂_x, ∇_x should act on score function $\Phi(x; \theta)$ while $\partial_\theta, \nabla_\theta$ act on density function $p(x; \theta)$. In this paper, We focus on the WIM on the statistical models, e.g. Gaussian, mixture Gaussian, or even infinite dimensional non-parametric model. The score functions will have enough regularity to take gradients, at least for most parametric models. And the above definition can be well applied to this settings.

Remark 5 This definition of information matrices is motivated by an intrinsic connection among distances, divergence functions, and metrics. Without loss of generality, working with a smooth family of probability densities $p(x; \theta)$ with $\theta \in \Theta = \mathbb{R}^d$ and a given perturbation $\Delta\theta \in T_\theta \Theta$, consider following Taylor expansions in term of $\Delta\theta$:

$$\begin{aligned} D_{KL}(p(\theta) \| p(\theta + \Delta\theta)) &= \frac{1}{2} \Delta\theta^\top G_F(\theta) \Delta\theta + o((\Delta\theta)^2), \\ W_2(p(\theta + \Delta\theta), p(\theta))^2 &= \Delta\theta^\top G_W(\theta) \Delta\theta + o((\Delta\theta)^2), \end{aligned} \quad (2)$$

where we used the identification of $T_\theta \Theta$ as \mathbb{R}^d . Here $D_{KL}(\cdot \| \cdot)$ represents the Kullback–Leibler (KL) divergence, a.k.a. the relative entropy functional

$$D_{KL}(p(\theta) \| p(\theta + \Delta\theta)) = \int_{\mathcal{X}} p(x; \theta) \log \frac{p(x; \theta)}{p(x; \theta + \Delta\theta)} dx.$$

While W_2^2 denotes the squared L^2 -Wasserstein distance defined by

$$W_2(p(\theta), p(\theta + \Delta\theta))^2 = \inf_{\pi \in \Pi(p(\theta), p(\theta + \Delta\theta))} \left\{ \int_{\mathcal{X} \times \mathcal{X}} d_{\mathcal{X}}(x, y)^2 d\pi(x, y) \right\}, \quad (3)$$

where $\Pi(p(\theta), p(\theta + \Delta\theta))$ refers to the set of couplings between $p(\theta)$, $p(\theta + \Delta\theta)$ and $d_{\mathcal{X}}$ is a distance function defined in \mathcal{X} . Thus our approach parallels classical Fisher statistics in the following sense: the Fisher information matrix approximates the KL divergence, which relates to the Fisher distance in Fisher geometry [1, 6], while WIM approximates the Wasserstein distance in Wasserstein geometry. Meanwhile, our approach can be viewed as exploring metric aspect of the Wasserstein statistics. For example, it can be used to study the Wasserstein estimators [7].

We next study several basic properties of WIMs and score functions. As the Wasserstein metric is associated with taking derivatives in the sample space, we assume all

the tangent and cotangent vectors appearing later to have enough regularities, which hold for most statistical models discussed in this paper. We first illustrate a relation between Wasserstein and Fisher score functions.

Proposition 4 (Poisson equation) *Wasserstein score functions $\Phi_i^W(x; \theta)$ satisfy the following Poisson equation*

$$\nabla_x \log p(x; \theta) \cdot \nabla_x \Phi_i^W(x; \theta) + \Delta_x \Phi_i^W(x; \theta) = -\frac{\partial}{\partial \theta_i} \log p(x; \theta). \quad (4)$$

Proof Notice the fact that

$$\begin{aligned} (\Delta_{p_\theta}) \Phi_i^W(x; \theta) &= \nabla_x \cdot (p(x; \theta) \nabla_x \Phi_i^W(x; \theta)) \\ &= \nabla_x p(x; \theta) \cdot \nabla_x \Phi_i^W(x; \theta) + p(x; \theta) \Delta_x \Phi_i^W(x; \theta). \end{aligned}$$

Then the Wasserstein score function $\Phi_i^W(x)$ satisfies

$$\nabla_x p(x; \theta) \cdot \nabla_x \Phi_i^W(x; \theta) + p(x; \theta) \Delta_x \Phi_i^W(x; \theta) = -\frac{\partial}{\partial \theta_i} p(x; \theta).$$

Divide the above equation on both sides by $p(x; \theta)$

$$\frac{1}{p(x; \theta)} \left\{ \nabla_x p(x; \theta) \cdot \nabla_x \Phi_i^W(x; \theta) + p(x; \theta) \Delta_x \Phi_i^W(x; \theta) \right\} = -\frac{1}{p(x; \theta)} \frac{\partial}{\partial \theta_i} p(x; \theta),$$

i.e.

$$\frac{1}{p(x; \theta)} \nabla_x p(x; \theta) \cdot \nabla_x \Phi_i^W(x; \theta) + \Delta_x \Phi_i^W(x; \theta) = -\frac{1}{p(x; \theta)} \frac{\partial}{\partial \theta_i} p(x; \theta).$$

Since $\frac{1}{p(x; \theta)} \nabla_x p(x; \theta) = \nabla_x \log p(x; \theta)$ and $\frac{1}{p(x; \theta)} \frac{\partial}{\partial \theta_i} p(x; \theta) = \frac{\partial}{\partial \theta_i} \log p(x; \theta)$, we prove the property (4). \square

We then demonstrate that Wasserstein score functions and information matrices can also be decomposed into a summation of separable functions in independent models.

Proposition 5 (Separability) *If $p(x; \theta)$ is an independence model, i.e.*

$$p(x; \theta) = \prod_{k=1}^n p_k(x_k; \theta), \quad x_k \in \mathcal{X}_k, \quad x = (x_1, \dots, x_n).$$

Then there exists a set of functions $\Phi_i^{W,k}: \mathcal{X}_k \times \Theta \rightarrow \mathbb{R}$, $i = 1, 2, \dots, \dim \Theta$, $k = 1, 2, \dots, n$, such that

$$\Phi_i^W(x; \theta) = \sum_{k=1}^n \Phi_i^{W,k}(x_k; \theta). \quad (5)$$

In addition, the WIM is separable, i.e.

$$G_W(\theta) = \sum_{k=1}^n G_W^k(\theta),$$

where $(G_W^k(\theta))_{ij} = \mathbb{E}_{p_k} \left(\frac{\partial}{\partial x} \Phi_i^{W,k}(x; \theta), \frac{\partial}{\partial x} \Phi_j^{W,k}(x; \theta) \right)$.

Proof The proof follows from proposition 4. Suppose one can write the solution in form of (5), then equation (4) forms

$$\sum_{k=1}^n \left\{ \frac{\partial}{\partial x_k} \log p_k(x_k; \theta_k) \frac{\partial}{\partial x_k} \Phi_i^{W,k}(x_k; \theta) + \frac{\partial^2}{\partial x_k^2} \Phi_i^{W,k}(x_k; \theta) - \frac{\partial}{\partial \theta_i} \log p_k(x_k; \theta) \right\} = 0.$$

From the separable method for solving the Poisson equation, we derive

$$\frac{\partial}{\partial x_k} \log p_k(x_k; \theta_k) \frac{\partial}{\partial x_k} \Phi_i^{W,k}(x_k; \theta) + \frac{\partial^2}{\partial x_k^2} \Phi_i^{W,k}(x_k; \theta_k) - \frac{\partial}{\partial \theta_i} \log p_k(x_k; \theta) = 0.$$

We finish the first part of the proof. In addition,

$$\begin{aligned} (G_W(\theta))_{ij} &= \mathbb{E}_{p_\theta} \left(\frac{\partial}{\partial x} \Phi_i^W(x; \theta), \frac{\partial}{\partial x} \Phi_j^W(x; \theta) \right) \\ &= \mathbb{E}_{p_\theta} \left(\sum_k \left(\frac{\partial}{\partial x} \Phi_i^{W,k}(x; \theta), \frac{\partial}{\partial x} \Phi_j^{W,k}(x; \theta) \right) \right) \\ &= \sum_k \mathbb{E}_{p_k} \left(\frac{\partial}{\partial x} \Phi_i^{W,k}(x; \theta), \frac{\partial}{\partial x} \Phi_j^{W,k}(x; \theta) \right) \\ &= \sum_k (G_W^k(\theta))_{ij}. \end{aligned}$$

□

We next list some analytical solutions for WIMs and score functions in 1-d case. See related studies in [28] (c.f. Ch 2.2).

Proposition 6 (One-dimensional sample space) *If $\mathcal{X} \subset \mathbb{R}^1$, Wasserstein score functions satisfy*

$$\Phi_i^W(x; \theta) = - \int_{\mathcal{X} \cap (\infty, x]} \frac{1}{p(z; \theta)} \frac{\partial}{\partial \theta_i} F(z; \theta) dz, \quad (6)$$

where $F(x; \theta) = \int_{\mathcal{X} \cap (\infty, x]} p(y; \theta) dy$ is the cumulative distribution function. And the WIM satisfies

$$G_W(\theta)_{ij} = \mathbb{E}_{p_\theta} \left(\frac{\frac{\partial}{\partial \theta_i} F(x; \theta) \frac{\partial}{\partial \theta_j} F(x; \theta)}{p(x; \theta)^2} \right).$$

If the dimension of sample space \mathcal{X} is larger than 1, exact solutions of Wasserstein score functions and information matrices depend on solutions of Poisson equation (4). We leave the derivation of general formulas for interested readers.

2.3 Analytic examples

We present several analytical examples of the WIM in one-dimensional sample space. The derivation is given in Sect. A.

Example 2 (Gaussian distribution) Consider Gaussian distribution families with mean value μ and standard variance $\sigma > 0$, i.e. $p(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2}(x-\mu)^2}$. Wasserstein score functions satisfy

$$\Phi_\mu^W(x; \mu, \sigma) = x - \mu, \quad \Phi_\sigma^W(x; \mu, \sigma) = \frac{(x - \mu)^2 - \sigma^2}{2\sigma}.$$

And the WIM satisfies

$$G_W(\mu, \sigma) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Example 3 (Exponential distribution) Consider exponential distribution families $Exp(m, \lambda)$, i.e. $p(x; m, \lambda) = \mathbf{1}_{[m, \infty)}(x) \lambda e^{-\lambda(x-m)}$, where the function $\mathbf{1}_C$ is the indicator function for a set $C \subset \mathbb{R}$. Wasserstein score functions satisfy

$$\Phi_\lambda^W(x; m, \lambda) = \frac{(x - m)^2 - \frac{2}{\lambda^2}}{2\lambda}, \quad \Phi_m^W(x; m, \lambda) = x - m - \frac{1}{\lambda}.$$

And the WIM satisfies

$$G_W(m, \lambda) = \begin{pmatrix} 1 & \frac{1}{\lambda^2} \\ \frac{1}{\lambda^2} & \frac{1}{\lambda^4} \end{pmatrix}.$$

Example 4 (Laplacian distribution) Consider Laplacian distribution families $La(m, \lambda)$, i.e. $p(x; m, \lambda) = \frac{\lambda}{2} e^{-\lambda|x-m|}$. Wasserstein score functions satisfy

$$\Phi_\lambda^W(x; m, \lambda) = \frac{(x - m)^2 - \frac{2}{\lambda^2}}{2\lambda}, \quad \Phi_m^W(x; m, \lambda) = x - m.$$

Notice that score functions for exponential families and Laplacian families have similar formulas. And the WIM satisfies

$$G_W(m, \lambda) = \begin{pmatrix} 1 & 0 \\ 0 & \frac{2}{\lambda^4} \end{pmatrix}.$$

We will show below that the Laplacian family has an advantage that densities within this family have the same support. Thus this model is convenient for us to compare the WIM with the Fisher information matrix. See details in Sect. B.3.

Example 5 (Uniform distribution) Consider uniform distribution families within interval $[a, b]$, i.e. $p(x; a, b) = \frac{1}{b-a} \mathbf{1}_{[a,b]}(x)$. Wasserstein score functions satisfy

$$\begin{aligned} \Phi_a^W(x; a, b) &= \frac{x(a + b - x)}{(b - a)} - \frac{b^2 + a^2 + 4ab}{6}, \\ \Phi_b^W(x; a, b) &= \frac{b(x - 2a)}{(b - a)} - \frac{b^2 - 3ab}{2}. \end{aligned}$$

And the WIM satisfies

$$G_W(a, b) = \frac{1}{3} \begin{pmatrix} 1 & \frac{1}{2} \\ \frac{1}{2} & 1 \end{pmatrix}.$$

Example 6 (Wigner semicircle distribution) Consider semicircle distribution families, i.e. $p(x; m, R) = \mathbf{1}_{[-R+m, R+m]}(x) \frac{2}{\pi R^2} \sqrt{R^2 - (x - m)^2}$. Wasserstein score functions satisfy

$$\Phi_R^W(x; m, R) = \frac{1}{R} \left(\frac{(x - m)^2}{2} - \frac{R^2}{8} \right), \quad \Phi_p^W(x; m, R) = x - m.$$

And the WIM satisfies

$$G_W(m, R) = \begin{pmatrix} 1 & 0 \\ 0 & \frac{1}{4} \end{pmatrix}.$$

Example 7 (Independent model) Consider an independent model as $X \sim p_1(x; \theta)$, and $Y \sim p_2(x; \theta)$, and $(X, Y) \sim p(x, y; \theta)$, then

$$p(x, y; \theta) = p_1(x; \theta)p_2(y; \theta).$$

Denote Wasserstein score functions (resp. WIM) for statistical model $X \sim p_1(x; \theta)$, $Y \sim p_2(x; \theta)$ as $\Phi_1^W(x; \theta)$, $\Phi_2^W(x; \theta)$ ($G_1^1(x; \theta)$, $G_2^2(x; \theta)$) respectively. Then, Wasserstein score functions for this model $(X, Y) \sim p(x, y; \theta)$ satisfy

$$\Phi^W(x, y; \theta) = \Phi_1^W(x; \theta) + \Phi_2^W(y; \theta),$$

because of the additivity of expectation $\mathbb{E}_{p_\theta} \Phi^W(x, y; \theta) = \mathbb{E}_{p_\theta} \Phi_1^W(x; \theta) + \mathbb{E}_{p_\theta} \Phi_2^W(y; \theta) = 0$. And the WIM satisfies

$$G_W(\theta) = G_W^1(\theta) + G_W^2(\theta).$$

The proof follows directly from proposition 5.

In above discussions, all examples are based on location-scale families, which will be derived carefully in Sect. A.2. We show that location-scale families are totally geodesic submanifolds in Wasserstein geometry.

2.4 WIM in generative models

In this section, we study the WIM for generative models using ReLU function, which is given by

$$\sigma(x) = \begin{cases} 0, & x \leq 0, \\ x, & x > 0. \end{cases}$$

Generative models are powerful in machine learning [18]. It applies the reparameterization trick (known as push-forward relation) to conduct efficient sampling. In practice, one often applies the ReLU as a push-forward function (7). For this reason, we call this kind of models ReLU push-forward family. The push-forward measure $f_* p$ is defined as

$$\int_A f_* p dx = \int_{f^{-1}(A)} p dx, \quad \forall A \subset \mathbb{R}. \quad (7)$$

To keep derivations simple, we consider one-dimensional cases with a given distribution $p_0(x)$, $x \in \mathbb{R}$. And its cumulative distribution function is denoted by $F_0(x)$.

Example 8 (ReLU push-forward family) We use a family of ReLU functions f_θ parameterized by θ to generate a push-forward family

$$p : \Theta \simeq \mathbb{R} \rightarrow \mathcal{P}(\mathbb{R}) : \quad \theta \mapsto p_\theta,$$

$$p_\theta(x) = p(x; \theta) = (f_{\theta*} p_0)(x), \quad f_\theta(x) = \sigma(x - \theta) = \begin{cases} 0, & x \leq \theta, \\ x - \theta, & x > \theta. \end{cases}$$

The WIM of p_θ satisfies

$$G_W(\theta) = 1 - F_0(\theta). \quad (8)$$

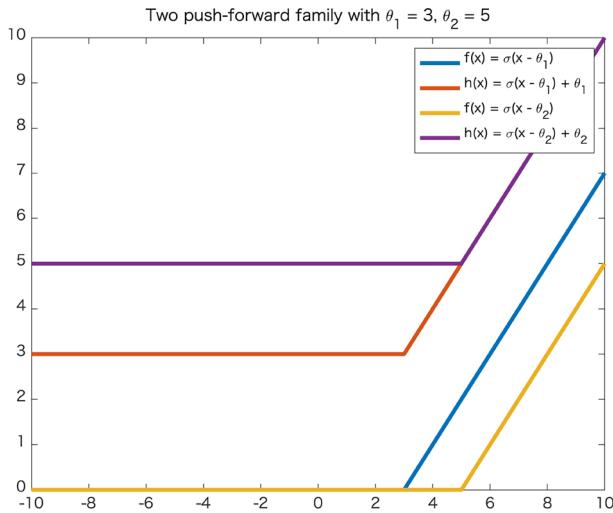


Fig. 1 This figure plots two examples of push-forward families with parameters chosen as $\theta_1 = 3, \theta_2 = 5$

We can also consider another family of ReLU maps to push forward the source distribution. This family is given by

$$p : \Theta \simeq \mathbb{R} \rightarrow \mathcal{P}(\mathbb{R}) : \theta \mapsto p_\theta$$

$$p_\theta(x) = p(x; \theta) = (h_{\theta*} p_0)(x), \quad h_\theta(x) = \sigma(x - \theta) + \theta = \begin{cases} \theta, & x \leq \theta, \\ x, & x > \theta. \end{cases}$$

The WIM of p_θ satisfies

$$G_W(\theta) = F_0(\theta). \quad (9)$$

A figure illustrating these two families is provided below.

Proof To calculate WIMs of this model, we cannot use previous approaches of score functions, since it is not smooth enough. Instead, we utilize the idea stated in remark 5. Namely, we use the relation (2) between Wasserstein distance and WIMs to compute the latter.

Consider the following two push-forward distributions given by

$$(f_{\theta+\Delta\theta*} p_0)(x) = F_0(\theta + \Delta\theta) \delta_0 + p_0(\cdot + \theta + \Delta\theta)_{[0, \infty)},$$

$$(f_{\theta*} p_0)(x) = F_0(\theta) \delta_0 + p_0(\cdot + \theta)_{[0, \infty)},$$

where δ_0 refers to the Dirac measure concentrating at point 0. And $p_0(\cdot + \theta)_{[0, \infty)}$ represents the measure $\tilde{p}(x) = p_0(x + \theta)$ restricting to the interval $[0, \infty)$. Using monotonicity of transportation plan in 1-d, we conclude that its restriction on $(0, \infty)$

transports measure on x to $x + \Delta\theta$. And it remains to transport the Dirac measure centered at 0 to the remained place. The transportation cost is given by

$$\begin{aligned} W_2^2(f_{\theta*}p_0, f_{\theta+\Delta\theta*}p_0) &= \int_0^\infty p_0(x + \theta + \Delta\theta) (\Delta\theta)^2 dx + \int_0^{\Delta\theta} x^2 p_0(x + \theta) dx \\ &= (\Delta\theta)^2 (1 - F_0(\theta + \Delta\theta)) + O((\Delta\theta)^3), \end{aligned} \quad (10)$$

where the third equality holds by

$$\int_0^{\Delta\theta} p_0(x + \theta) dx = O(\Delta\theta).$$

Notice in formula (10), we decompose the transportation cost into two parts: the first one is concerned with the cost on the right part of 0, while the second one considers transporting Dirac measure at 0 to the remained part. Since the WIM is an infinitesimal approximation of the Wasserstein distance, i.e. equation (2). The conclusion (8) follows.

For the other family, derivations follow the same method as before. Specifically, we have

$$\begin{aligned} W_2^2(h_{\theta*}p_0, h_{\theta+\Delta\theta*}p_0) &= \int_0^{\Delta\theta} x^2 p_0(x + \theta) dx + (\Delta\theta)^2 F_0(\theta) \\ &= (\Delta\theta)^2 F_0(\theta) + O((\Delta\theta)^3), \end{aligned}$$

where we again decompose the transportation cost into two parts. The first one is absolutely continuous w.r.t the Lebesgue measure, while the second one contains a Dirac measure. \square

Remark 6 Notice that density functions in ReLU push-forward family can be singular. Thus the Fisher information matrix, which depends on an explicit formula of density functions, namely

$$G_F(\theta)_{ij} = \int_{\mathcal{X}} \frac{\partial}{\partial \theta_i} \log p(x; \theta) \frac{\partial}{\partial \theta_j} \log p(x; \theta) p(x; \theta) dx$$

fails to exist in these models. On the contrary, as we have shown in the above example, the WIM still exists. This property shows that the WIM can provide statistical studies for generative models, while the Fisher information matrix in classical statistics can not. Moreover, compared to the information matrix definition, the remark 5 can handle a broader regime of statistical models that lack regularity or even contain singular Dirac measures. The ability to analyze generative models is what we believe to be an advantage of WIM over FIM. On the other hand, WIM can not handle all the parametric families. As pointed out in [23], due to its lack of the second moment, Cauchy family does not have a well-defined WIM.

3 Wasserstein estimation

In this section, we define the Wasserstein covariance and establish the Wasserstein–Cramer–Rao bound. Based on these concepts, we introduce a notion of efficiency in Wasserstein statistics. Several examples based on the previous section are provided. From now on, we will avoid the discussion on general density manifold but rather focus on finite dimensional statistical models associated with an information matrix structure.

3.1 Estimation and efficiency

Firstly, we generalize the definition of covariance matrix for a given metric tensor g on probability space. Denote $\langle f, h \rangle_g$ as the inner product of cotangent vectors f, h in the metric g :

$$\langle f, h \rangle_g = \langle f, g(p)^{-1}h \rangle.$$

Definition 7 (Information covariance matrix) Given a statistical model Θ with metric g , and statistics $T \in \mathbb{R}^m$, $\tilde{T} \in \mathbb{R}^n$, the information covariance matrix $\text{Cov}_\theta^g[T, \tilde{T}] \in \mathbb{R}^{m \times n}$ is defined by

$$(\text{Cov}_\theta^g[T, \tilde{T}])_{ij} := \langle T_i, \tilde{T}_j \rangle_g,$$

where $i = 1, 2, \dots, m$, $j = 1, 2, \dots, n$ and T_i, \tilde{T}_j are random variables as function of x . Denote the information covariance matrix as

$$(\text{Var}_\theta^g[T])_{ij} = (\text{Cov}_\theta^g[T])_{ij} = (\text{Cov}_\theta^g[T, T])_{ij} = \langle T_i, T_j \rangle_g,$$

for $i, j = 1, 2, \dots, m$.

Example 9 (Fisher covariance) Given two statistics T_1, T_2 , we view them as cotangent vectors in space $C(\mathcal{X})/\mathbb{R}$. Hence their Fisher inner product is defined as

$$\langle T_1, T_2 \rangle_{g_F} = \int_{\mathcal{X}} (T_1 - \mathbb{E}_{p_\theta}[T_1]) (T_2 - \mathbb{E}_{p_\theta}[T_2]) p_\theta dx.$$

Here choosing the function $T_1 - \mathbb{E}_{p_\theta}[T_1]$ as the representative of $[T_1]$ is consistent with the normalization requirement (1). Thus Fisher covariance (resp. variance) reduces to the original definition of the covariance (resp. variance) in probability theory. And the classical Cramer–Rao bound is given by

$$\text{Cov}_\theta^F[T(x)] \succeq \nabla_\theta \mathbb{E}_{p_\theta}[T(x)]^\top G_F(\theta)^{-1} \nabla_\theta \mathbb{E}_{p_\theta}[T(x)],$$

where $G_F(\theta)$ is the Fisher information matrix. In 1-d cases, the above forms

$$\text{Var}_\theta[T(x)] \geq \frac{(\nabla_\theta \mathbb{E}_{p_\theta} T(x))^2}{G_F(\theta)},$$

where we use classical notion of variance here.

We next focus on the Wasserstein covariance operator.

Definition 8 (Wasserstein covariance) Given a statistical model Θ , denote the Wasserstein covariance as follows

$$\text{Cov}_\theta^W[T_1, T_2] = \mathbb{E}_{p_\theta} \left[\nabla_x T_1(x) \cdot \nabla_x T_2(x)^\top \right],$$

where T_1, T_2 are random variables as functions of x and the expectation is taken w.r.t. $x \sim p_\theta$. Here \cdot refers to matrix product. Denote the Wasserstein variance

$$\text{Var}_\theta^W[T] = \text{Cov}_\theta^W[T] = \mathbb{E}_{p_\theta} \left[\nabla_x T(x) \cdot \nabla_x T(x)^\top \right].$$

Remark 7 Intuitively speaking, since any statistics is a function of sample x , the Wasserstein covariance measure the sensitivity of the statistics w.r.t. the underlying sample space structure. Based on this interpretation, one possible application of Wasserstein statistics is robust statistics. Indeed, Wasserstein distribution robust optimization [8] has been widely used. More theoretic understanding and derivation are left to future work.

Now we can state the main theorem of this section.

Theorem 9 (Wasserstein–Cramer–Rao inequality) *Given any set of statistics $T = (T_1, \dots, T_m) : \mathcal{X} \rightarrow \mathbb{R}^m$, where m is the number of the statistics, define $\nabla_\theta \mathbb{E}_{p_\theta} [T(x)]^\top = (\nabla_\theta \mathbb{E}_{p_\theta} [T(x)]^\top)_{ij} = \frac{\partial}{\partial \theta_j} \mathbb{E}_{p_\theta} [T_i(x)]$. Then*

$$\text{Cov}_\theta^W[T(x)] \succeq \nabla_\theta \mathbb{E}_{p_\theta} [T(x)]^\top G_W(\theta)^{-1} \nabla_\theta \mathbb{E}_{p_\theta} [T(x)],$$

where the notion \succeq refers to that the difference of two matrices is positive semi-definite.

Proposition 10 (Covariance property) *Given the Wasserstein score function $\Phi_i^W(x; \theta)$ and any smooth statistic $T : \mathcal{X} \rightarrow \mathbb{R}$, then*

$$\frac{\partial}{\partial \theta_i} \mathbb{E}_{p_\theta} T(x) = \mathbb{E}_{p_\theta} \left[\nabla_x \Phi_i^W(x; \theta) \cdot \nabla_x T(x)^\top \right] = \langle \Phi_i^W, f \rangle_{g_W}.$$

Proof Notice the fact that

$$\begin{aligned}
 \frac{\partial}{\partial \theta_i} \mathbb{E}_{p_\theta} T(x) &= \frac{\partial}{\partial \theta_i} \int_{\mathcal{X}} T(x) p(x; \theta) dx \\
 &= \int_{\mathcal{X}} T(x) \frac{\partial}{\partial \theta_i} p(x; \theta) dx \\
 &= \int_{\mathcal{X}} T(x) \left(-\nabla_x \cdot (p(x; \theta) \nabla_x \Phi_i^W(x; \theta)) \right) dx \\
 &= \int_{\mathcal{X}} \nabla_x T(x) \cdot \nabla_x \Phi_i^W(x; \theta) p(x; \theta) dx,
 \end{aligned}$$

where the third equality comes from the definition of Wasserstein score functions, while the last equality holds by integration by parts formula in spatial domain. The conclusion holds. \square

Remark 8 This property is in contrast to Fisher score functions, i.e.

$$\begin{aligned}
 \frac{\partial}{\partial \theta_i} \mathbb{E}_{p_\theta} T(x) &= \mathbb{E}_{p_\theta} \left[T(x) \frac{\partial}{\partial \theta_i} \log p(x; \theta) \right] \\
 &= \text{Cov}_\theta \left[T(x), \frac{\partial}{\partial \theta_i} \log p(x; \theta) \right] = \langle \Phi_i^F, T \rangle_{g_F}.
 \end{aligned}$$

This is merely a dual relation between tangent and cotangent space in the density manifold.

Proof (Proof of Theorem 9) By the definition of semi-positive matrix, it suffices to prove that for arbitrary $v \in \mathbb{R}^m$, we have

$$v^\top \text{Cov}_\theta^W [T(x)] v \geq v^\top \nabla_\theta \mathbb{E}_{p_\theta} [T(x)]^\top G_W(\theta)^{-1} \nabla_\theta \mathbb{E}_{p_\theta} [T(x)] v.$$

Here we define $T_v = v^\top T$ as the statistic associated to the vector v . Then the LHS of above formula equals to the variance of T_v , i.e.

$$v^\top \text{Cov}_\theta^W [T(x)] v = \text{Cov}_\theta^W [T_v].$$

As we have mentioned before, score functions $\Phi_i^W, i = 1, 2, \dots, n$ span a subspace $T_{p(x; \theta)}^* \Theta$ of the cotangent space $T_{p(x; \theta)}^* \mathcal{P}(\mathcal{X})$ at each point $p(x; \theta)$ on the density manifold. Meanwhile, the statistic $T_v: \mathcal{X} \rightarrow \mathbb{R}$ can be viewed as a cotangent vector field on this statistical model. Since the subspace $T_{p(x; \theta)}^* \Theta$ at each point θ is a finite-dimensional subspace of the Hilbert space $T_{p(x; \theta)}^* \mathcal{P}(\mathcal{X})$ endowed with the Wasserstein inner product. Thus it is a closed linear subspace. By elementary theory of functional analysis, we have orthogonal projection operator \mathbf{P} as

$$\mathbf{P}: T_{p(x; \theta)}^* \mathcal{P}(\mathcal{X}) \rightarrow T_{p(x; \theta)}^* \Theta.$$

Since $\Phi_i^W, i = 1, 2, \dots, n$ span the whole subspace, we have

$$\langle \Phi_i^W, v - \mathbf{P}v \rangle_{g_W} = 0, \quad \forall v \in T_{p(x;\theta)}^* \mathcal{P}(\mathcal{X}).$$

Back to the theorem, we have

$$\text{Cov}_\theta^W[T_v] = \mathbb{E}_{p_\theta} \left[\nabla_x T_v(x) \cdot \nabla_x T_v(x)^\top \right] = \langle T_v, T_v \rangle_{g_W} \geq \langle \mathbf{P}T_v, \mathbf{P}T_v \rangle_{g_W},$$

where the last inequality holds by the property of the orthogonal projection operator.

Since \mathbf{P} is the projection onto the subspace $T_{p(x;\theta)}^* \Theta$ with a set of basis Φ_i^W , at each point θ , we can write the cotangent vector $\mathbf{P}T_v$ as a linear combination of Wasserstein score functions below

$$\mathbf{P}T_v = \sum_{i=1}^d t_i^\theta \Phi_i^W,$$

where the superscript of t_i^θ indicates the dependency on point θ . Plugging this linear combination into the Wasserstein metric, we get

$$\begin{aligned} \langle \mathbf{P}T_v, \mathbf{P}T_v \rangle_{g_W} &= \sum_{i=1}^d t_i^\theta \langle \mathbf{P}T_v, \Phi_i^W \rangle_{g_W} \\ &= \sum_{i,k=1}^d t_k^\theta \delta_i^k \langle \mathbf{P}T_v, \Phi_i^W \rangle_{g_W} \\ &= \sum_{i,j,k=1}^d t_k^\theta g_{kj} g^{ij} \langle \mathbf{P}T_v, \Phi_i^W \rangle_{g_W} \\ &= \sum_{i,j=1}^d \langle \mathbf{P}T_v, \Phi_i^W \rangle_{g_W} \left(G_W(\theta)^{-1} \right)_{ij} \langle \mathbf{P}T_v, \Phi_i^W \rangle_{g_W} \\ &= \nabla_\theta \mathbb{E}_{p_\theta} [T_v(x)]^\top G_W(\theta)^{-1} \nabla_\theta \mathbb{E}_{p_\theta} [T_v(x)], \end{aligned}$$

where $G_W(\theta)^{-1}$ is the inverse matrix of the WIM, g_{kj} , g^{ij} are elements of matrix G_W , G_W^{-1} respectively, and the third equality holds by the fact $\sum_j g_{kj} g^{ij} = \delta_i^k$. The last equality is guaranteed by proposition 10. Combining the above calculation and the comparison between the inner product of T_v and $\mathbf{P}T_v$, we obtain the desired result. \square

Given the above theorem, we can define the Wasserstein efficiency as follows.

Definition 11 For an estimator $T(x)$, it is Wasserstein efficient if and only if it attains the Wasserstein–Cramer–Rao bound, namely

$$\text{Var}_\theta^W[T(x)] = \nabla_\theta \mathbb{E}_{p_\theta} [T(x)]^\top G_W(\theta)^{-1} \nabla_\theta \mathbb{E}_{p_\theta} [T(x)].$$

Next, we provide some examples of efficient statistics in classical models for Wasserstein and Fisher statistics.

Example 10 (Gaussian distribution) Recall that given a Gaussian distribution with mean value μ and standard variance σ , Wasserstein score functions satisfy

$$\Phi_\mu^W(x; \mu, \sigma) = x - \mu, \quad \Phi_\sigma^W(x; \mu, \sigma) = \frac{(x - \mu)^2 - \sigma^2}{2\sigma},$$

with the WIM

$$G_W(\mu, \sigma) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Thus by the criterion, we know that a statistics is Wasserstein efficient iff it can be written as linear combinations of Wasserstein score functions. Statistics only depend on samples x_i and do not depend on parameters μ, σ . Hence they must be of the form

$$ax^2 + bx + c = 2a\sigma\Phi_\sigma^W + (2a\mu + b)\Phi_\mu^W + c + a\mu^2 + b\mu - a\sigma^2,$$

where c is a constant. Wasserstein efficient statistics are degree 2 polynomials of x .

While Fisher score functions are given by

$$\Phi_\mu^F(x; \mu, \sigma) = \frac{x - \mu}{\sigma^2}, \quad \Phi_\sigma^F(x; \mu, \sigma) = \frac{(x - \mu)^2}{\sigma^3} - \frac{1}{\sigma}.$$

And the Fisher information matrix satisfies

$$G_F(\mu, \sigma) = \begin{pmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{2}{\sigma^2} \end{pmatrix}.$$

Although we have different score functions in Wasserstein and Fisher statistics, it turns out that efficient statistics associated with these two information matrices coincide. Still, Fisher and Wasserstein information matrices provide us with different Cramer–Rao bounds. The Wasserstein–Cramer–Rao bound is better if σ is small while worse if σ is large.

Example 11 (Exponential distribution) Given an exponential distribution, Wasserstein score functions satisfy

$$\Phi_\lambda^W(x; m, \lambda) = \frac{(x - m)^2 - \frac{2}{\lambda^2}}{2\lambda}, \quad \Phi_m^W(x; m, \lambda) = x - m - \frac{1}{\lambda},$$

and the WIM satisfies

$$G_W(m, \lambda) = \begin{pmatrix} 1 & \frac{1}{\lambda^2} \\ \frac{1}{\lambda^2} & \frac{2}{\lambda^4} \end{pmatrix}.$$

Similarly to Gaussian cases, Wasserstein sufficient statistics are also quadratic functions of variables x . Fisher score functions read

$$\Phi_\lambda^F(x; m, \lambda) = m - x + \frac{1}{\lambda}, \quad \Phi_m^F(x; m, \lambda) \text{ not well defined.}$$

Meanwhile, the Fisher information matrix is also ill-behaved. This is in contrast to the well-definedness of both Wasserstein score functions and WIMs. This example provides a situation where Wasserstein statistics are better than the classical Fisher statistics.

4 Wasserstein natural gradient works efficiently

In this section, we study Wasserstein dynamics in terms of sampling and estimation processes. As a consequence, we prove Wasserstein asymptotic efficiencies of the natural gradient algorithm. And we refer it as Wasserstein efficiency. Meanwhile, another efficiency that we name Poincaré efficiency is introduced. It is connected to Poincaré inequalities and log-Sobolev inequalities, which are discussed in Sect. B.

In the beginning, we review the natural gradient algorithm. We aim to estimate an un-known distribution in a probability family $p(x; \theta)$ with unknown parameters $\theta \in \Theta$. Assume the optimal parameter θ_* exists such that $p(x; \theta_*)$ coincides with the target distribution. Given a set of i.i.d. samples $x_t \sim p(x; \theta_*)$, $t \in \mathbb{Z}_+$, we utilize a general online natural gradient algorithm to solve this problem, i.e.

$$\theta_{t+1} = \theta_t - \frac{1}{t} \nabla_\theta^W l(x_t, \theta_t). \quad (11)$$

In the above formula, θ_t is an updating state variable, $\frac{1}{t}$ in the RHS is an adaptive factor. And ∇_θ^W is the Wasserstein (natural) gradient of the loss function l w.r.t. θ

$$\nabla_\theta^W l(x_t, \theta_t) = G_W^{-1} \nabla_\theta l(x_t, \theta_t).$$

We first define the efficiency of the natural gradient algorithm, which generalizes the notion discussed in [2]. Denote the Wasserstein covariance matrix of estimator θ_t by

$$V_t = \mathbb{E}_{p_{\theta_*}} \left[\nabla_x(\theta_t - \theta_*) \cdot \nabla_x(\theta_t - \theta_*)^T \right],$$

where $\nabla_x(\theta_t - \theta_*)$ is the matrix given by

$$\nabla_x(\theta_t - \theta_*) = \begin{pmatrix} \frac{\partial(\theta_t - \theta_*)_1}{\partial x_1} & \frac{\partial(\theta_t - \theta_*)_1}{\partial x_2} & \dots & \frac{\partial(\theta_t - \theta_*)_1}{\partial x_n} \\ \frac{\partial(\theta_t - \theta_*)_2}{\partial x_1} & \frac{\partial(\theta_t - \theta_*)_2}{\partial x_2} & \dots & \frac{\partial(\theta_t - \theta_*)_2}{\partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial(\theta_t - \theta_*)_n}{\partial x_1} & \frac{\partial(\theta_t - \theta_*)_n}{\partial x_2} & \dots & \frac{\partial(\theta_t - \theta_*)_n}{\partial x_n} \end{pmatrix},$$

and \cdot is matrix product. It turns out that the element of the covariance matrix is given by

$$\mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x (\theta_t - \theta_*)^T \right]_{ij} = \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*)_i \cdot \nabla_x (\theta_t - \theta_*)_j \right].$$

Here \cdot refers to inner product of gradient vectors. And subscripts $p(\cdot; \theta_*)$ refer to take expectation on the set of samples $x_t \sim p(\cdot; \theta_*), t \in \mathbb{Z}_+$.

Definition 12 The Wasserstein natural gradient is asymptotic efficient if

$$V_t = \frac{1}{t} G_W^{-1}(\theta_*) + O\left(\frac{1}{t^2}\right).$$

We first state a general variance updating equation for this dynamics. Then, we specify two different loss functions, namely, Fisher scores and Wasserstein scores. And we discuss convergence properties of these two cases separately.

Theorem 13 (Variance updating equation of the Wasserstein natural gradient) *For any function $l(x, \theta)$ that satisfies the condition $\mathbb{E}_{p_\theta} l(x, \theta) = 0$, consider the asymptotic behavior of the Wasserstein online natural gradient (11). That is, assume priorly*

$$\mathbb{E}_{p_{\theta_*}} \left[(\theta_t - \theta_*)^2 \right], \mathbb{E}_{p_{\theta_*}} \left[|\nabla_x (\theta_t - \theta_*)|^2 \right] = o(1), \forall t.$$

Then, the Wasserstein covariance matrix V_t updates according to the following equation

$$V_{t+1} = V_t + \frac{1}{t^2} G_W^{-1}(\theta_*) \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (l(x_t, \theta_*)) \cdot \nabla_x (l(x_t, \theta_*)^T) \right] \left(G_W^{-1}(\theta_*) \right) - \frac{2V_t}{t} \mathbb{E}_{p_{\theta_*}} [\nabla_\theta l(x_t, \theta_*)] G_W^{-1}(\theta_*) + o\left(\frac{V_t}{t}\right) + o\left(\frac{1}{t^2}\right).$$

Remark 9 In general, it will be shown that such a simple updating equation merely attributes to properties of information matrices. Specifically, any statistical information matrices with separability property w.r.t. independent variables have this form of updating equation. For the WIM, this is already established in proposition 5. And for the Fisher information matrix, this is a property of expectation of independent variables. Further results such as efficiency of the natural gradient can be established with the same procedure below.

The proof is technical and we leave it to Sect. C.1. Here, we show several important cases of Theorem 13.

4.1 Wasserstein natural gradient for Wasserstein scores

We consider the following dynamics in this subsection

$$\nabla_\theta^W l(x_t, \theta_t) = G_W^{-1} \Phi^W(x_t; \theta_t). \quad (12)$$

Concerned with this dynamics, we have the following corollary.

Corollary 14 (Wasserstein natural gradient efficiency) *For the dynamics*

$$\theta_{t+1} = \theta_t - \frac{1}{t} G_W^{-1}(\theta_t) \Phi^W(x_t; \theta_t),$$

the Wasserstein covariance updates according to

$$V_{t+1} = V_t + \frac{1}{t^2} G_W^{-1}(\theta_*) - \frac{2}{t} V_t + o\left(\frac{1}{t^2}\right) + o\left(\frac{V_t}{t}\right).$$

Then, the online Wasserstein natural gradient algorithm is Wasserstein efficient, that is

$$V_t = \frac{1}{t} G_W^{-1}(\theta_*) + O\left(\frac{1}{t^2}\right). \quad (13)$$

Proof of Corollary 14 If we choose function $f(x, t)$ to be Wasserstein scores Φ_i^W , we will have following simplification

$$\mathbb{E}_{p_{\theta_*}} \left[\nabla_x \left(\Phi^W(x_t; \theta_*) \right) \cdot \nabla_x \left(\Phi^W(x_t; \theta_*)^T \right) \right] = G_W(\theta_*),$$

since Φ^W is the dual coordinate of the statistical model. We also have

$$\mathbb{E}_{p_{\theta_*}} \left[\nabla_{\theta} \Phi^W(x_t; \theta_*) \right] = -G_W(\theta_*),$$

which is given by differentiating $\mathbb{E}_{p_{\theta_*}} [\Phi^W(x_t; \theta_*)] = \mathbf{0}$ by θ , i.e.

$$\begin{aligned} \mathbf{0} &= \nabla_{\theta} \mathbb{E}_{p_{\theta_*}} \left[\Phi^W(x_t; \theta_*) \right] \\ &= \nabla_{\theta} \left[\int_{\mathcal{X}} p(x; \theta_*) \Phi^W(x; \theta_*) dx \right] \\ &= \int_{\mathcal{X}} \nabla_{\theta} p(x; \theta_*) \Phi^W(x; \theta_*) dx + \int_{\mathcal{X}} p(x; \theta_*) \nabla_{\theta} \Phi^W(x; \theta_*) dx \\ &= G_W + \int_{\mathcal{X}} p(x; \theta_*) \nabla_{\theta} \Phi^W(x; \theta_*) dx, \end{aligned}$$

where the last equality holds because of the pairing between tangent vector and cotangent vector. And the final updating equation for the Wasserstein covariance reduces to

$$V_{t+1} = V_t + \frac{1}{t^2} G_W^{-1}(\theta_*) - \frac{2}{t} V_t + O\left(\frac{1}{t^3}\right) + o\left(\frac{V_t}{t}\right).$$

To further solve this updating equation, we expand $V_t = \frac{x}{t} + \frac{y}{t^2} + o\left(\frac{1}{t^2}\right)$ with constant x, y to be determined and plug into the equation (we ignore the term that is of order $o\left(\frac{1}{t^2}\right)$). Hence, we get

$$\frac{x}{t+1} + \frac{y}{(t+1)^2} + o\left(\frac{1}{t^2}\right) = \frac{x}{t} + \frac{y}{t^2} + o\left(\frac{1}{t^2}\right) + \frac{1}{t^2} G_W^{-1}(\theta_*) - \frac{2x}{t^2} + o\left(\frac{1}{t^2}\right),$$

which is equivalent to

$$\left(\frac{x}{t+1} - \frac{x}{t}\right) + \left(\frac{y}{(t+1)^2} - \frac{y}{t^2}\right) + \frac{2x}{t^2} - \frac{1}{t^2} G_W^{-1}(\theta_*) + o\left(\frac{1}{t^2}\right) = 0.$$

And we conclude that

$$x = G_W^{-1}(\theta_*).$$

Thus, we asymptotically have following estimation on the Wasserstein covariance concerned with this dynamics

$$V_t = \frac{1}{t} G_W^{-1}(\theta_*) + o\left(\frac{1}{t}\right).$$

Remark 10 At first, such a generalization to Wasserstein metric may seem unreasonable. We only use a fact that both of them are metrics on probability spaces. Different from Fisher scores $\Phi^F = \nabla_\theta l(x; \theta)$, Wasserstein scores Φ^W can not be written as gradients of some functions w.r.t. θ . There is no such “loss functions”. However, a key insight here is that, if in a second we assume that the statistical model Θ is exactly the density manifold $G_W(p_\theta) = g_W(p_\theta)$, $G_F(p_\theta) = g_F(p_\theta)$, i.e.

$$\begin{aligned} G_W^{-1}(p_\theta) \Phi^W(x; \theta) &= g_W(p_\theta) g_W^{-1}(p_\theta) \frac{\partial}{\partial \theta} p(x; \theta) = \nabla_\theta p(x_t; \theta_t) \\ &= g_F(p_\theta) g_F^{-1}(p_\theta) \frac{\partial}{\partial \theta} p(x; \theta) = G_F^{-1}(p_\theta) \Phi^F(x; \theta). \end{aligned}$$

Then both two dynamics can be written in the following way

$$\theta_{t+1} = \theta_t - \frac{1}{t} \nabla_\theta p(x_t; \theta_t).$$

4.2 Wasserstein natural gradient for Fisher scores

Another phenomenon appears when we consider the Wasserstein natural gradient applies to Fisher scores. Specifically, we use log-likelihood function as a loss function and apply WIM as a preconditioner. The dynamics concerned in this case is given by

$$\theta_{t+1} = \theta_t - \frac{1}{t} \nabla_{\theta}^W l(x_t, \theta_t).$$

We comment that $\nabla_{\theta} l(x_t, \theta_t) = \Phi^F(x_t, \theta_t)$ is both the Euclidean gradient of log-likelihood function l w.r.t. θ and the Fisher score. And the convergence analysis is shown in the following corollary.

Corollary 15 (Poincaré efficiency) *For the dynamics*

$$\theta_{t+1} = \theta_t - \frac{1}{t} \nabla_{\theta}^W l(x_t, \theta_t),$$

the Wasserstein covariance updates according to

$$V_{t+1} = V_t + \frac{1}{t^2} G_W^{-1}(\theta_*) \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\nabla_{\theta} l(x_t, \theta_*)) \cdot \nabla_x (\nabla_{\theta} l(x_t, \theta_*)^T) \right] G_W^{-1}(\theta_*) \\ - \frac{2}{t} V_t G_F(\theta_*) G_W^{-1}(\theta_*) + O\left(\frac{1}{t^3}\right) + o\left(\frac{V_t}{t}\right).$$

Now suppose that $\alpha = \sup\{a|G_F \succeq aG_W\}$. Then the dynamics is characterized by the following formula

$$V_t = \begin{cases} O(t^{-2\alpha}), & 2\alpha \leq 1, \\ \frac{1}{t} \left(2G_F G_W^{-1} - \mathbf{I} \right)^{-1} G_W^{-1}(\theta_*) \mathfrak{I} \left(G_W^{-1}(\theta_*) \right) + O\left(\frac{1}{t^2}\right), & 2\alpha > 1, \end{cases} \quad (14)$$

where

$$\mathfrak{I} = (\mathfrak{I})_{ij} = \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\nabla_{\theta_i} l(x_t, \theta_*)) \cdot \nabla_x (\nabla_{\theta_j} l(x_t, \theta_*)^T) \right]$$

Proof of Corollary 15 The result is obtained once we observe that

$$\mathbb{E}_{p_{\theta_*}} \left[\nabla_{\theta} \Phi^F(x_t, \theta_*) \right] = -G_F(\theta_*),$$

which follows exactly the same philosophy of corollary 14. We conclude that the Wasserstein covariance updates according to

$$V_{t+1} = V_t + \frac{1}{t^2} G_W^{-1}(\theta_*) \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\nabla_{\theta} l(x_t, \theta_*)) \cdot \nabla_x (\nabla_{\theta} l(x_t, \theta_*)^T) \right] \left(G_W^{-1}(\theta_*) \right) \\ - \frac{2}{t} V_t G_F(\theta_*) G_W^{-1}(\theta_*) + O\left(\frac{1}{t^3}\right) + o\left(\frac{V_t}{t}\right).$$

Next, we solve this dynamics asymptotically. We denote $G_F(\theta_*) G_W^{-1}(\theta_*) = B$ and

$$G_W^{-1}(\theta_*) \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\nabla_{\theta} l(x_t, \theta_*)) \cdot \nabla_x (\nabla_{\theta} l(x_t, \theta_*)^T) \right] \left(G_W^{-1}(\theta_*) \right) = C.$$

Now by elementary linear algebra, we know that the matrix $B = G_F(\theta_*)G_W^{-1}(\theta_*)$ is similar to the matrix $G_W^{-\frac{1}{2}}G_FG_W^{-\frac{1}{2}}$. Hence their eigenvalues coincide. By definition, α is the smallest eigenvalue of the symmetric matrix $G_W^{-\frac{1}{2}}G_FG_W^{-\frac{1}{2}}$. Thus we conclude that the smallest eigenvalues of the matrix B are also α . Suppose first that $2\alpha < 1$, we consider the following expansion of matrix V_t as

$$V_t = \frac{A_1}{t^q} + \frac{A_2}{t^{q+1}} + o\left(\frac{1}{t^{q+1}}\right), \quad A_1, A_2 = O(1).$$

Plug the above equation to both sides of the updating equation, we find

$$\begin{aligned} & \frac{A_1}{(t+1)^q} + \frac{A_2}{(t+1)^{q+1}} + o\left(\frac{1}{t^{q+1}}\right) \\ &= \frac{A_1}{t^q} + \frac{A_2}{t^{q+1}} + o\left(\frac{1}{t^{q+1}}\right) - \frac{2A_1B}{t^{q+1}} + \frac{C}{t^2} + O\left(\frac{1}{t^3}\right). \end{aligned}$$

Using the Lagrange's mean value theorem, we have

$$\frac{A}{t^q} - \frac{A}{(t+1)^q} = \frac{qA}{(t+v)^{q+1}} = \frac{qA}{t^{q+1}} + o\left(\frac{1}{t^q}\right), \quad v \in [0, 1].$$

Substituting back to the above equation, we get

$$\mathbf{0} = \frac{A_1(q\mathbf{I} - 2B)}{t^{q+1}} + \frac{C}{t^2} + o\left(\frac{1}{t^{q+1}}\right) + O\left(\frac{1}{t^3}\right).$$

We cannot have q strictly greater than 1, for then the most significant term in the RHS will be $\frac{C}{t^2} \neq \mathbf{0}$ which contradicts to the LHS. Thus if we have $q < 2\alpha < 1$, the matrix $q\mathbf{I} - 2B$ will be negative definite, and we cannot have $A_1(q\mathbf{I} - 2B) = \mathbf{0}$ unless A_1 equals to 0. Consequently, the index q should be greater than or equal to 2α . And we have that asymptotically

$$V_t = O\left(\frac{1}{t^{2\alpha}}\right).$$

While for the situation such that $2\alpha > 1$, we expand $V_t = \frac{A_1}{t} + \frac{A_2}{t^2} + o\left(\frac{1}{t^2}\right)$ with constant A_1, A_2 to be determined and obtain

$$\frac{A_1}{t+1} + \frac{A_2}{(t+1)^2} + o\left(\frac{1}{t^2}\right) = \frac{A_1}{t} + \frac{A_2}{t^2} + o\left(\frac{1}{t^2}\right) + \frac{C}{t^2} - \frac{2A_1B}{t^2} + o\left(\frac{1}{t^2}\right).$$

The constant A_1 can be fixed by considering the coefficient of the term $\frac{1}{t^2}$ for both sides with conclusion

$$A_1 = (2B - \mathbf{I})^{-1} C.$$

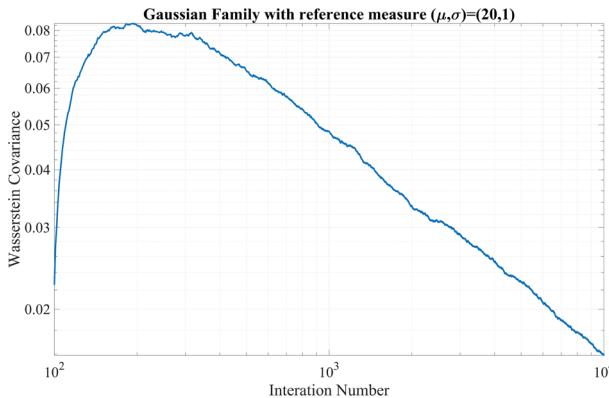


Fig. 2 The Wasserstein–Cramer–Rao Type Convergence Rate. We take the reference measure in KL-divergence to be Gaussian $\mathcal{N}(20, 1)$ where the parameter $\mu_* = 20$ is the optimal point we aim to estimate. Since we have $\frac{1}{\sigma_*^2} = 1 > \frac{1}{2}$, the Cramer–Rao type convergence holds. The result is calculated by running 1e3 samples for 1e4 iterations

Here, invertibility of the matrix $2B - \mathbf{I}$ is guaranteed by the fact that eigenvalues of $2B$ are all greater than 1, thus the matrix $2B - \mathbf{I}$ is indeed positive definite. \square

The convergence behavior of this dynamics relies largely on the smallest eigenvalue of the matrix $G_F G_W^{-1}$. This is in great similarity with the RIW condition for Poincaré inequality in statistical models [17]. This inspires us to name such efficiency Poincaré efficiency. For discussions and calculations on Poincaré inequalities in statistical models, see details in the Sec. B. We also illustrate some results of two efficiencies in Gaussian family, whose proof is given in the Sect. C.2.

Example 12 (Gaussian distribution) Suppose we have following dynamics in a Gaussian model $p(x; \mu, \sigma)$

$$\theta_{t+1} = \theta_t - \frac{1}{t} \nabla_{\theta}^W l(x_t, \theta_t), \quad x_t \sim p(x; \mu_*, \sigma_*).$$

The asymptotic behavior of the covariance matrix for the online Wasserstein natural gradient algorithm is given by

$$V_t = \begin{cases} O\left(t^{-\frac{2}{\sigma_*^2}}\right), & \frac{1}{\sigma_*^2} \leq \frac{1}{2}, \\ \frac{1}{t} \begin{pmatrix} \frac{1}{(2-\sigma_*^2)\sigma_*^2} & 0 \\ 0 & \frac{4}{(4-\sigma_*^2)\sigma_*^2} \end{pmatrix} + O\left(\frac{1}{t^2}\right), & \frac{1}{\sigma_*^2} > \frac{1}{2}. \end{cases}$$

We verify our theory by two numerical experiments. In two cases, we verify the Wasserstein efficiency and the Poincaré efficiency respectively. In the first experiment, we verify the constant G_W^{-1} appearing in asymptotic efficiency of the Wasserstein natural gradient. While for the other situation we verify the asymptotic exponential index α showing up in Poincaré efficiency.

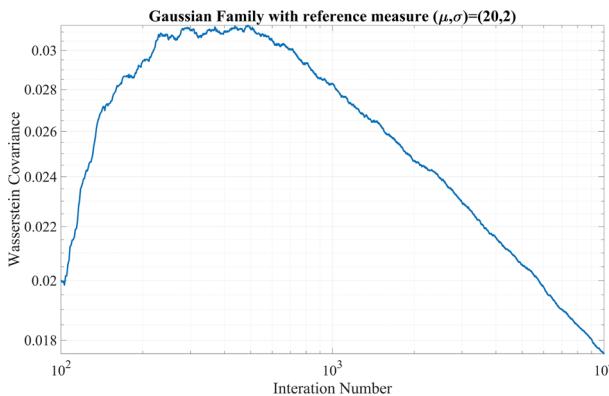


Fig. 3 Poincaré Type Convergence Rate. We take the reference measure in KL-divergence to be Gaussian $\mathcal{N}(20, 2)$ where the parameter $\mu_* = 20$ is the optimal point we aim to estimate. Since we have $\frac{1}{\sigma_*^2} = \frac{1}{4} < \frac{1}{2}$, the Poincaré type convergence holds. The result is calculated by running 1e3 samples for 1e4 iterations

5 Discussions

In this paper, we introduce the Wasserstein information matrix in statistical models. Similar to the study in information geometry, we turn the geometric aspect of the Wasserstein metric into statistics. We generalize the classical concepts such as score function, covariance operator, Cramer–Rao bound, and estimation to the Wasserstein statistics. Several explicit computable examples are provided, including the location-scale family and the ReLU push-forward family. Also, by comparing both Wasserstein and Fisher information matrices, some new efficiency concepts, such as Wasserstein efficiency and Poincaré efficiency have been introduced.

In the future, several natural questions between Fisher and Wasserstein statistics arise. For example, similar to the relation with Fisher information matrices and maximal likelihood estimators, what is the relation between the WIM and the Wasserstein distance estimator? Is there a canonical Wasserstein divergence function for the WIM? What is the corresponding Wasserstein maximal likelihood estimator? Meanwhile, we will apply the Wasserstein natural gradient to study stochastic gradient descent algorithms in statistical learning problems. Lastly, we have shown that the Wasserstein statistics provide rigorous statistical advantages in generative models over classical Fisher statistics. We will study the statistical properties of WIMs in detail for machine learning models.

Funding W. Li is supported by AFOSR MURI FA9550-18-1-0502, AFOSR YIP award 2023, and NSF RTG: 2038080.

Data Availability Statement Not Applicable.

Declarations

Conflict of interest There is no conflict of interest between authors W. Li and J.X. Zhao.

Appendix A. Proofs in Sect. 2

A.1. WIMs and score functions in analytic examples

Proof of WIMs in Gaussian families Since we have

$$\log p(x; \mu, \sigma) = -\frac{(x - \mu)^2}{2\sigma^2} - \log \sigma - \log \sqrt{2\pi},$$

taking derivative, we get

$$\begin{aligned} \frac{\partial}{\partial x} \log p(x; \mu, \sigma) &= -\frac{x - \mu}{\sigma^2}, \\ \frac{\partial}{\partial \mu} \log p(x; \mu, \sigma) &= \frac{x - \mu}{\sigma^2}, \quad \frac{\partial}{\partial \sigma} \log p(x; \mu, \sigma) = \frac{(x - \mu)^2}{\sigma^3} - \frac{1}{\sigma}. \end{aligned}$$

In this case, the Possion equation for Wasserstein score functions $(\Phi_\mu^W, \Phi_\sigma^W)$ forms

$$\begin{cases} -\frac{x - \mu}{\sigma^2} \cdot \frac{\partial}{\partial x} \Phi_\mu^W + \frac{\partial^2}{\partial x^2} \Phi_\mu^W = -\frac{x - \mu}{\sigma^2}, \\ -\frac{x - \mu}{\sigma^2} \cdot \frac{\partial}{\partial x} \Phi_\sigma^W + \frac{\partial^2}{\partial x^2} \Phi_\sigma^W = -\frac{(x - \mu)^2}{\sigma^3} + \frac{1}{\sigma}. \end{cases}$$

We simply check that $\Phi_\mu^W(x; \mu, \sigma) = x - \mu$ and $\Phi_\sigma^W(x; \mu, \sigma) = \frac{(x - \mu)^2 - \sigma^2}{2\sigma}$ are solutions, and they also satisfy the normalization condition $\mathbb{E}_{p_\theta} \Phi_i^W = 0$. Thus

$$\begin{aligned} G_W(\mu, \sigma)_{\mu\mu} &= \mathbb{E}_{p_{\mu, \sigma}} \left(\frac{\partial}{\partial x} \Phi_\mu^W, \frac{\partial}{\partial x} \Phi_\mu^W \right) = \mathbb{E}_{p_{\mu, \sigma}} 1 = 1, \\ G_W(\mu, \sigma)_{\mu\sigma} &= \mathbb{E}_{p_{\mu, \sigma}} \left(\frac{\partial}{\partial x} \Phi_\mu^W, \frac{\partial}{\partial x} \Phi_\sigma^W \right) = \mathbb{E}_{p_{\mu, \sigma}} \left(1 \cdot \left(-\frac{X - \mu}{2\sigma} \right) \right) = 0, \\ G_W(\mu, \sigma)_{\sigma\sigma} &= \mathbb{E}_{p_{\mu, \sigma}} \left(\frac{\partial}{\partial x} \Phi_\sigma^W, \frac{\partial}{\partial x} \Phi_\sigma^W \right) = \mathbb{E}_{p_{\mu, \sigma}} \left(\frac{X - \mu}{\sigma} \cdot \frac{X - \mu}{\sigma} \right) = 1. \end{aligned}$$

Proof of WIMs in exponential families We derive results using the closed-form solution in 1-d. The cumulative distribution function satisfies

$$F(x; m, \lambda) = \begin{cases} 1 - e^{-\lambda(x-m)} & x \geq m, \\ 0 & x < m. \end{cases}$$

Thus

$$\frac{\partial}{\partial \lambda} F(x; m, \lambda) = \begin{cases} (x - m)e^{-\lambda(x-m)} & x \geq m, \\ 0 & x < m. \end{cases}$$

$$\frac{\partial}{\partial m} F(x; m, \lambda) = \begin{cases} \lambda e^{-\lambda(x-m)} & x \geq m, \\ 0 & x < m. \end{cases}$$

Then

$$\Phi_\lambda^W(x; m, \lambda) = - \int_m^x \frac{1}{p(y; m, \lambda)} \frac{\partial}{\partial \lambda} F(y; m, \lambda) dy + C_1$$

$$= - \int_m^x \frac{(y - m)}{\lambda} dy + C_1 = \frac{(x - m)^2}{2\lambda} + C_1,$$

$$\Phi_m^W(x; m, \lambda) = - \int_m^x \frac{1}{p(y; m, \lambda)} \frac{\partial}{\partial m} F(y; m, \lambda) dy + C_2$$

$$= - \int_m^x dy + C_2 = (x - m) + C_2.$$

Using the normalization condition, we can decide integration constants appearing above. And inner products between score functions follow as

$$G_W(m, \lambda)_{\lambda\lambda} = \mathbb{E}_{p_{m,\lambda}} \left(\frac{\partial}{\partial x} \Phi_\lambda^W, \frac{\partial}{\partial x} \Phi_\lambda^W \right)$$

$$= \int_m^\infty \frac{(x - m)}{\lambda} \cdot \frac{(x - m)}{\lambda} \cdot \lambda e^{-\lambda(x-m)} dx$$

$$= \int_m^\infty \frac{(x - m)^2}{\lambda} e^{-\lambda(x-m)} dx = \frac{2}{\lambda^4},$$

$$G_W(m, \lambda)_{\lambda m} = \mathbb{E}_{p_{m,\lambda}} \left(\frac{\partial}{\partial x} \Phi_\lambda^W, \frac{\partial}{\partial x} \Phi_m^W \right) = \int_m^\infty \frac{(x - m)}{\lambda} \cdot \lambda e^{-\lambda(x-m)} dx = \frac{1}{\lambda^2},$$

$$G_W(m, \lambda)_{mm} = \mathbb{E}_{p_{m,\lambda}} \left(\frac{\partial}{\partial x} \Phi_m^W, \frac{\partial}{\partial x} \Phi_m^W \right) = \int_m^\infty \lambda e^{-\lambda(x-m)} dx = 1.$$

Proof of WIMs in uniform families The cumulative distribution function satisfies

$$F(x; a, b) = \begin{cases} 1 & x > b, \\ \frac{x-a}{b-a} & a \leq x \leq b, \\ 0 & x < a. \end{cases}$$

Thus when $x \in [a, b]$,

$$\frac{\partial}{\partial a} F(x; a, b) = \frac{x - b}{(b - a)^2}, \quad \frac{\partial}{\partial b} F(x; a, b) = \frac{a - x}{(b - a)^2}.$$

Then

$$\begin{aligned}\Phi_a^W(x; a, b) &= - \int_a^x \frac{1}{p(y; a, b)} \frac{\partial}{\partial a} F(y; a, b) dy + C_1 = \frac{(x-a)(a-2b+x)}{2(b-a)} + C_1, \\ \Phi_b^W(x; a, b) &= - \int_a^x \frac{1}{p(y; a, b)} \frac{\partial}{\partial b} F(y; a, b) dy + C_2 = \frac{(a-x)^2}{2(b-a)} + C_2,\end{aligned}$$

where integration constants C_1, C_2 can be decided via the normalization condition. Thus

$$\begin{aligned}G_W(a, b)_{aa} &= \mathbb{E}_{p_{a,b}} \left(\frac{\partial}{\partial x} \Phi_a^W, \frac{\partial}{\partial x} \Phi_a^W \right) = \frac{1}{3}, \\ G_W(a, b)_{ab} &= \mathbb{E}_{p_{a,b}} \left(\frac{\partial}{\partial x} \Phi_a^W, \frac{\partial}{\partial x} \Phi_b^W \right) = \frac{1}{6}, \\ G_W(a, b)_{bb} &= \mathbb{E}_{p_{a,b}} \left(\frac{\partial}{\partial x} \Phi_b^W, \frac{\partial}{\partial x} \Phi_b^W \right) = \frac{1}{3}.\end{aligned}$$

Proof of the WIM in semicircle families The cumulative distribution function satisfies

$$\begin{aligned}F(x+m; m, R) &= \int_{-R}^x \frac{2}{\pi R^2} \sqrt{R^2 - y^2} dy \\ &= \int_{-\frac{\pi}{2}}^{\arcsin(\frac{x}{R})} \frac{2}{\pi R^2} \sqrt{R^2 - R^2 \sin^2 t} d(R \sin t) \\ &= \int_{-\frac{\pi}{2}}^{\arcsin(\frac{x}{R})} \frac{2}{\pi R^2} R^2 (\cos t)^2 dt \\ &= \int_{-\frac{\pi}{2}}^{\arcsin(\frac{x}{R})} \frac{1}{2\pi} \frac{\cos(2t) + 1}{2} dt \\ &= \frac{1}{\pi} \left(\frac{\sin(2t)}{2} + t \right) \Big|_{-\frac{\pi}{2}}^{\arcsin \frac{x}{R}} \\ &= \frac{1}{\pi} \left\{ \frac{x\sqrt{R^2 - x^2}}{R^2} + \arcsin \frac{x}{R} + \frac{\pi}{2} \right\},\end{aligned}$$

where we use a transformation $y = R \sin t$. Thus

$$\begin{aligned}\frac{\partial}{\partial R} F(x+m; m, R) &= \frac{1}{\pi} \left\{ \frac{(x\sqrt{R^2 - x^2})' R^2 - 2Rx\sqrt{R^2 - x^2}}{R^4} + \left(\arcsin \frac{x}{R} \right)' \right\} \\ &= \frac{1}{\pi} \left\{ \frac{xR(R^2 - x^2)^{-\frac{1}{2}} R^2 - 2Rx\sqrt{R^2 - x^2}}{R^4} - \frac{x}{R\sqrt{R^2 - x^2}} \right\} \\ &= - \frac{2x\sqrt{R^2 - x^2}}{\pi R^3}.\end{aligned}$$

Thus

$$\begin{aligned}\Phi_R^W(x + m; m, R) &= - \int_{-R}^x \frac{1}{p(y; m, R)} \frac{\partial}{\partial R} F(y; m, R) dy + C \\ &= \int_{-R}^x \frac{y}{R} dy + C \\ &= \frac{1}{R} \left(\frac{x^2}{2} - \frac{R^2}{2} \right) + C.\end{aligned}$$

The calculation of the score function associated with the parameter p is the same as before. And we conclude

$$\Phi_p^W(x; m, R) = x - m.$$

Thus

$$\begin{aligned}G_W(m, R)_{mm} &= \mathbb{E}_{p_{m,R}} \left(\frac{\partial}{\partial x} \Phi_m^W, \frac{\partial}{\partial x} \Phi_m^W \right) = 1, \\ G_W(m, R)_{mR} &= \mathbb{E}_{p_{m,R}} \left(\frac{\partial}{\partial x} \Phi_R^W, \frac{\partial}{\partial x} \Phi_m^W \right) = \frac{1}{R^2} \mathbb{E}_{m_R} (x - m) = 0, \\ G_W(m, R)_{RR} &= \mathbb{E}_{p_{m,R}} \left(\frac{\partial}{\partial x} \Phi_R^W, \frac{\partial}{\partial x} \Phi_R^W \right) = \frac{1}{R^2} \mathbb{E}_{m_R} (x - m)^2 = \frac{1}{4}.\end{aligned}$$

A.2. The location-scale family

Example 13 (Location-scale families) Consider a location-scale family as follows. Given a probability density function $p(x)$ with $\int_{\mathbb{R}} p(x) dx = 1$, we define density functions of a location-scale family with a location parameter m , and a scale parameter λ as

$$p(x; m, \lambda) = \frac{1}{\lambda} p\left(\frac{x - m}{\lambda}\right), \quad \lambda > 0.$$

Most of previously discussed examples belong to this family, except that we do not use location and scale parameters in their parameterizations. We present some geometric formulas in this setting. We further require the original density function to be symmetric according to the location parameter m , i.e. $p(x) = p(2m - x)$. Notice that a simple corollary of this assumption is $\mathbb{E}_{p_{m,\lambda}} x = m$.

We use the closed-form solution for 1-d model to calculate the score function associated with the location parameter m . Thus we have

$$\begin{aligned}\frac{\partial}{\partial m} F(x; m, \lambda) &= \frac{\partial}{\partial m} \int_{-\infty}^x p(y; m, \lambda) dy = \frac{\partial}{\partial m} \int_{-\infty}^x \frac{1}{\lambda} p\left(\frac{y - m}{\lambda}\right) dy \\ &= -\frac{\partial}{\partial x} \int_{-\infty}^x \frac{1}{\lambda} p\left(\frac{y - m}{\lambda}\right) dy = -p(x; m, \lambda).\end{aligned}$$

Consequently, the score function associated to the parameter m satisfies

$$\Phi_m^W(x; m, \lambda) = - \int_m^x \frac{1}{p(y; m, \lambda)} \frac{\partial}{\partial m} F(y; m, \lambda) dy + C_1 = (x - m) + C_1,$$

where the integration constant C_1 is determined to be 0. Thus we have

$$G_W(m, \lambda)_{mm} = \mathbb{E}_{p_{m, \lambda}} \left(\frac{\partial}{\partial x} \Phi_m^W, \frac{\partial}{\partial x} \Phi_m^W \right) = 1.$$

For the scaling parameter λ , we use a method of optimal transportation map to determine its score function. Namely, for two smooth distributions p_1, p_2 which are absolutely continuous w.r.t. each other, their Wasserstein distance can be obtained by an optimal transportation map f , i.e.

$$f_* p_1 = p_2, \quad W_2^2(p_1, p_2) = \int_{\mathcal{X}} (f(x) - x)^2 p_1(x) dx.$$

Assume we have a tangent vector $\frac{\partial p}{\partial \theta}$ and a smooth path $p(t) \subset \mathcal{P}(\mathcal{X})$, $t \in [-\epsilon, \epsilon]$ with $p(0) = p_0, p'(0) = \frac{\partial p}{\partial \theta}$. Denote the optimal transportation map between $p(0), p(\theta)$ as $f(x, \theta)$. Then we have following relation between optimal transportation maps and the score function associated with tangent vector $\frac{\partial p}{\partial \theta}$

$$\frac{\partial}{\partial x} \Phi^W(x) = \lim_{\Delta \theta \rightarrow 0} \frac{f(x, \Delta \theta) - x}{\Delta \theta}.$$

First, we show that the optimal transportation map between distributions $p(x; m_1, \lambda_1)$ and $p(x; m_2, \lambda_2)$ is given by a linear map

$$l(x) = m_2 + \frac{(x - m_1)\lambda_2}{\lambda_1}.$$

As we are working in a location-scale family, it is easy to show that this map pushes $p(x; m_1, \lambda_1)$ forward to $p(x; m_2, \lambda_2)$, i.e. $l_* p_{m_1, \lambda_1} = p_{m_2, \lambda_2}$. Then, we have

$$l(x) = \frac{\partial}{\partial x} \left(m_2 (x - m_1) + \frac{(x - m_1)^2 \lambda_2}{2\lambda_1} \right).$$

The function in the bracket is a convex function. Therefore, $l(x)$ is exactly the optimal transportation map between these two distributions.

To calculate the score function correspondent to the tangent vector $\frac{\partial}{\partial \lambda}$, we consider following infinitesimal optimal transportation $p(x; m_1, \lambda_1) \rightarrow p(x; m_1, \lambda_1 + d\lambda)$. By discussions above, the optimal transportation map is given by

$$l(x) = m_1 + \frac{(x - m_1)(\lambda_1 + d\lambda)}{\lambda_1} = x + (x - m_1) \frac{d\lambda}{\lambda_1}.$$

Thus the gradient of the score function is given by

$$\frac{\partial}{\partial x} \Phi_{\lambda}^W(x; m_1, \lambda_1) = \frac{l(x) - x}{d\lambda} = \frac{(x - m_1)}{\lambda_1}.$$

The inner product of this tangent vector is given by

$$\begin{aligned} G_W(m, \lambda)_{\lambda\lambda} &= \mathbb{E}_{p_{m,\lambda}} \left(\frac{\partial}{\partial x} \Phi_{\lambda}^W, \frac{\partial}{\partial x} \Phi_{\lambda}^W \right) \\ &= \int_{\mathbb{R}} \left(\frac{x - m}{\lambda} \right)^2 p(x; m, \lambda) dx \\ &= \frac{\mathbb{E}_{p_{m,\lambda}} x^2 - 2m \mathbb{E}_{p_{m,\lambda}} x + m^2}{\lambda^2}. \end{aligned}$$

The gradient of the score function associated to the parameter λ (resp. m) is odd (resp. even) function when viewing as a function of $x - m$. We conclude that the integration of their product is zero, i.e.

$$G_W(m, \lambda)_{\lambda m} = \mathbb{E}_{p_{m,\lambda}} \left(\frac{\partial}{\partial x} \Phi_{\lambda}^W, \frac{\partial}{\partial x} \Phi_m^W \right) = \mathbb{E}_{p_{m,\lambda}} (x - m) = 0.$$

Consequently, WIMs of location-scale families are diagonal matrices, i.e.

$$G_W(m, \lambda) = \begin{pmatrix} 1 & 0 \\ 0 & \frac{\mathbb{E}_{p_{m,\lambda}} x^2 - 2m \mathbb{E}_{p_{m,\lambda}} x + m^2}{\lambda^2} \end{pmatrix}.$$

We next explain above closed-form solutions of WIMs by following proposition.

Proposition 16 *A location-scale family $p(x; m, \lambda)$ is a totally geodesic family in density manifold under Wasserstein metric.*

Proof It suffices to prove that for any two densities $\rho_1 = p(x; m_1, \lambda_1)$ and $\rho_2 = p(x; m_2, \lambda_2)$, a geodesic connecting them lies within this family. We compute the optimal transport map T associated with these two measures ρ_1, ρ_2 , that is

$$T = \operatorname{argmin}_{T_* \rho_1 = \rho_2} \int_{\mathbb{R}} (T(x) - x)^2 \rho_1(x) dx,$$

where T is a map that pushes density ρ_1 forward to density ρ_2 . It is known that a sufficient and necessary condition for an optimal map in 1-d case is that it is a monotone map, i.e. $(T(x) - T(y))(x - y) \geq 0$. And in a location-scale family, such map has a closed-form solution, namely

$$T(x) = \frac{\lambda_2(x - m_1)}{\lambda_1} + m_2.$$

The geodesic $\gamma(t) : [0, 1] \rightarrow \mathcal{P}(\mathbb{R})$ between ρ_1 and ρ_2 follows easily as below by the classical theory of optimal transport

$$\gamma(t) = (tx + (1-t)T(x))_* \rho_1,$$

where the push-forward map has a closed-form solution

$$\begin{aligned} tx + (1-t)T(x) &= tx + (1-t) \frac{\lambda_2(x - m_1)}{\lambda_1} + (1-t)m_2 \\ &= \frac{(t\lambda_1 + (1-t)\lambda_2)(x - m_1)}{\lambda_1} + (1-t)m_2 + tm_1. \end{aligned}$$

And by the same argument, $\gamma(t)$ lies in this location-scale family with parameters given by

$$\lambda_t = t\lambda_1 + (1-t)\lambda_2, \quad m_t = (1-t)m_2 + tm_1.$$

Thus we show that geodesics between any two densities in a location-scale family lie in this family. In other words, location-scale families are totally geodesic submanifolds in density manifold. \square

Remark 11 This result on totally geodesic of location-scale families is a generalization of the same result on Gaussian families in 1-d. Both proofs of these two cases rely on the fact that optimal transport maps in these families are linear.

Remark 12 For location-scale families, we also formulate its Fisher scores and Fisher information matrices for comparisons as

$$\begin{aligned} \Phi_m^F(x; m, \lambda) &= \frac{p'}{\lambda p}, \quad \Phi_\lambda^F(x; m, \lambda) = -\frac{1}{\lambda} - \frac{(x - m)p'}{\lambda^2 p}, \\ G_F(m, \lambda)_{\lambda\lambda} &= \int_{\mathbb{R}} p \left(\frac{\partial}{\partial \lambda} \log p \right)^2 dx = \int_{\mathbb{R}} p \left(-\frac{1}{\lambda} - \frac{(x - m)p'}{\lambda^2 p} \right)^2 dx \\ &= \frac{1}{\lambda^2} \left(1 + \int_{\mathbb{R}} \left(\frac{(x - m)^2 p'^2}{\lambda^2 p} + \frac{(x - m)p'}{\lambda} \right) dx \right), \\ G_F(m, \lambda)_{mm} &= \int_{\mathbb{R}} p \left(\frac{\partial}{\partial m} \log p \right)^2 dx = \frac{1}{\lambda^2} \int_{\mathbb{R}} \frac{p'^2}{p} dx, \\ G_F(m, \lambda)_{m\lambda} &= \int_{\mathbb{R}} p \left(\frac{\partial}{\partial m} \log p \right) \left(\frac{\partial}{\partial \lambda} \log p \right) dx \\ &= \int_{\mathbb{R}} p \left(-\frac{p'}{\lambda p} \right) \left(-\frac{1}{\lambda} - \frac{(x - m)p'}{\lambda^2 p} \right) dx \\ &= \int_{\mathbb{R}} \frac{(x - m)p'^2}{\lambda^3 p} dx. \end{aligned} \tag{15}$$

Above, we use p' to denote $\frac{\partial}{\partial x} p$ for simplicity. WIMs and Fisher information matrices are given by

$$G_W(m, \lambda) = \begin{pmatrix} 1 & 0 \\ 0 & \frac{\mathbb{E}_{p_{m,\lambda}} x^2 - 2m\mathbb{E}_{p_{m,\lambda}} x + m^2}{\lambda^2} \end{pmatrix},$$

$$G_F(m, \lambda) = \frac{1}{\lambda^2} \left(\int_{\mathbb{R}} \frac{p'^2}{p} dx \quad \int_{\mathbb{R}} \frac{(x-m)p'^2}{\lambda p} dx \right) \begin{pmatrix} 1 + \int_{\mathbb{R}} \left(\frac{(x-m)^2 p'^2}{\lambda^2 p} + \frac{(x-m)p'}{\lambda} \right) dx \end{pmatrix},$$

which illustrates that WIMs are simpler than Fisher information matrices in location-scale families.

Appendix B. Functional inequalities via information matrices

In this section, we explore connections between information matrices and functional inequalities such as log-Sobolev inequalities (LSIs) and Poincaré inequalities (PIs) in statistical models. In Sect. 4, we show that these inequalities are important for the study of statistical efficiency properties.

B.1. Classical functional inequalities

Before working in statistical models, we first give a summary of relations among PIs, LSIs and dynamical quantities on density manifold.

Consider the relative entropy (KL-divergence) defined on density manifold as

$$D_{\text{KL}}(\mu \parallel \nu) = \int_{\mathcal{X}} \log \frac{\mu(x)}{\nu(x)} \mu(x) dx, \quad \mu \in \mathcal{P}(\mathcal{X}).$$

We recall the definition of log-Sobolev inequality as below.

Definition 17 (Log-Sobolev inequality) A probability measure ν is said to satisfy a log-Sobolev inequality with constant $\alpha > 0$ (in short: LSI(α)) if we have

$$D_{\text{KL}}(\mu \parallel \nu) < \frac{1}{2\alpha} I(\mu \parallel \nu), \quad \mu \in \mathcal{P}(\mathcal{X}),$$

where the quantity $I(\mu \parallel \nu)$ is the so-called Fisher-information functional

$$I(\mu \parallel \nu) = \int_{\mathcal{X}} \left| \nabla_x \log \frac{\mu(x)}{\nu(x)} \right|^2 \mu(x) dx, \quad \mu \in \mathcal{P}(\mathcal{X}).$$

Remark 13 If we assume that μ is absolutely continuous w.r.t. the reference measure ν and define function h on \mathcal{X} as

$$\mu(x) = \frac{h(x)\nu(x)}{\int_{\mathcal{X}} h(x)\nu(x) dx},$$

then above definition of LSI translates to

$$\begin{aligned} & \left(D_{KL}(\mu\|\nu) \int_{\mathcal{X}} h(x)\nu(x)dx \right) \\ &= \int_{\mathcal{X}} h(x) \log h(x)\nu(x)dx - \left(\int_{\mathcal{X}} h(x)\nu(x)dx \right) \log \left(\int_{\mathcal{X}} h(x)\nu(x)dx \right) \\ & \frac{1}{2\alpha} \int_{\mathcal{X}} \frac{|\nabla_x h(x)|^2}{h(x)} \nu(x)dx = \frac{1}{2\alpha} \left(I(\mu\|\nu) \int_{\mathcal{X}} h(x)\nu(x)dx \right). \end{aligned}$$

The middle inequality is a more familiar definition of $LSI(\alpha)$. By linearizing above formula with $h = 1 + \epsilon f$, $\epsilon \rightarrow 0$, we get the classical definition of $PI(\alpha)$

$$\int_{\mathcal{X}} f^2(x)\nu(x)dx \leq \frac{1}{\alpha} \int_{\mathcal{X}} |\nabla_x f(x)|^2 \nu(x)dx, \quad \int_{\mathcal{X}} f(x)\nu(x)dx = 0.$$

Definition 18 (Poincaré inequalities) A probability measure ν is said to satisfy a Poincaré inequalities with constant $\alpha > 0$ (in short: $PI(\alpha)$) if we have

$$\int_{\mathcal{X}} f^2(x)\nu(x)dx \leq \frac{1}{\alpha} \int_{\mathcal{X}} |\nabla_x f(x)|^2 \nu(x)dx, \quad \forall f, \text{ s.t. } \int_{\mathcal{X}} f(x)\nu(x)dx = 0.$$

A sufficient criterion that guarantees LSIs and PIs is related to information matrices (operators or metrics in infinite dimension case) G_W .

Proposition 19 Denote $\text{Hess}_W D_{KL}(\mu\|\nu)$, $G_W(\mu)$ two bi-linear forms correspondent to Hessian of the relative entropy and Wasserstein metric.

- (1) Suppose $\text{Hess}_W D_{KL}(\mu\|\nu) - 2\alpha G_W(\mu)$ is a semi-positive definite bi-linear form on the Hilbert space $T_{\mu}\mathcal{P}(\mathcal{X})$, $\forall \mu \in \mathcal{P}(\mathcal{X})$. Then $LSI(\alpha)$ holds for ν .
- (2) Suppose $\text{Hess}_W D_{KL}(\nu\|\nu) - 2\alpha G_W(\nu)$ is a semi-positive definite bi-linear form on the Hilbert space $T_{\nu}\mathcal{P}(\mathcal{X})$. Then $PI(\alpha)$ holds for ν .

Proof First, we prove the result concerned with LSIs. We compute the gradient of the relative entropy w.r.t. Wasserstein metric, which is given by

$$\text{grad}_W D_{KL}(\mu\|\nu) = -\nabla \cdot \left(\mu \nabla \frac{\delta}{\delta \mu} D_{KL}(\mu\|\nu) \right) = -\nabla \cdot \left(\mu \nabla \log \frac{\mu(x)}{\nu(x)} \right),$$

where $\frac{\delta}{\delta \mu}$ refers to the L^2 functional derivative. Thus it is easy to obtain the relative entropy dissipation along the gradient flow as

$$\begin{aligned} & \frac{d}{dt} D_{KL}(\mu\|\nu) \\ &= -g_W \left(\text{grad}_W D_{KL}(\mu\|\nu), \text{grad}_W D_{KL}(\mu\|\nu) \right) \\ &= - \int_{\mathcal{X}} \left| \nabla_x \log \frac{\mu(x)}{\nu(x)} \right|^2 \mu(x)dx = -I(\mu\|\nu). \end{aligned} \tag{16}$$

Using the assumption, we have

$$\begin{aligned}
 \frac{d^2}{dt^2} D_{\text{KL}}(\mu_t \| \nu) &= \text{Hess}_W D_{\text{KL}}(\mu_t \| \nu) (\text{grad}_W D_{\text{KL}}(\mu_t \| \nu), \text{grad}_W D_{\text{KL}}(\mu_t \| \nu)) \\
 &\geq 2\alpha G_W(\mu_t) (\text{grad}_W D_{\text{KL}}(\mu_t \| \nu), \text{grad}_W D_{\text{KL}}(\mu_t \| \nu)) \\
 &= -2\alpha \frac{d}{dt} D_{\text{KL}}(\mu_t \| \nu),
 \end{aligned}$$

from which LSI(α) holds via integrating the above formula, i.e.

$$\begin{aligned}
 I(\mu_t \| \nu) &= I(\mu_t \| \nu) - I(\nu \| \nu) \\
 &= \int_t^\infty \left(\frac{d^2}{dt^2} D_{\text{KL}}(\mu_\tau \| \nu) \right) d\tau \\
 &\geq 2\alpha \int_t^\infty \left(-\frac{d}{dt} D_{\text{KL}}(\mu_\tau \| \nu) \right) d\tau \\
 &= 2\alpha (D_{\text{KL}}(\mu_t \| \nu) - D_{\text{KL}}(\nu \| \nu)) \\
 &= 2\alpha D_{\text{KL}}(\mu_t \| \nu),
 \end{aligned}$$

where we use the fact that this gradient flow μ_t converges to ν and $D_{\text{KL}}(\nu \| \nu) = I(\nu \| \nu) = 0$.

To prove the conclusion of Poincaré inequalities, we consider a path in density manifold, i.e $\rho(\epsilon) = \nu(1 + \epsilon f)$, $\int_{\mathcal{X}} f(x)\nu(x)dx = 0$. Since we have

$$\begin{aligned}
 D_{\text{KL}}(\rho(\epsilon) \| \nu) &= \frac{\epsilon^2}{2} \int_{\mathcal{X}} f^2(x)\nu(x)dx + o(\epsilon^2), \\
 -\frac{d}{dt} D_{\text{KL}}(\rho(\epsilon) \| \nu) &= I(\rho(\epsilon) \| \nu) = \epsilon^2 \int_{\mathcal{X}} |\nabla_x f(x)|^2 \nu(x)dx + o(\epsilon^2).
 \end{aligned}$$

Consequently, we obtain

$$\begin{aligned}
 &\frac{\int_{\mathcal{X}} f^2(x)\nu(x)dx}{\int_{\mathcal{X}} |\nabla_x f(x)|^2 \nu(x)dx} \\
 &= \frac{1}{2} \lim_{\epsilon \rightarrow 0} -\frac{D_{\text{KL}}(\rho(\epsilon) \| \nu)}{\frac{d}{d\epsilon} D_{\text{KL}}(\rho(\epsilon) \| \nu)} \\
 &= \frac{1}{2} \lim_{\epsilon \rightarrow 0} -\frac{\frac{d}{d\epsilon} D_{\text{KL}}(\rho(\epsilon) \| \nu)}{\frac{d^2}{d\epsilon^2} D_{\text{KL}}(\rho(\epsilon) \| \nu)} \\
 &= \frac{1}{2} \lim_{\epsilon \rightarrow 0} \frac{G_W(\rho(0)) \left(\frac{d}{d\epsilon} \rho(0), \frac{d}{d\epsilon} \rho(0) \right)}{\text{Hess}_W D_{\text{KL}}(\rho(0) \| \nu) \left(\frac{d}{d\epsilon} \rho(0), \frac{d}{d\epsilon} \rho(0) \right)} \\
 &\leq \frac{1}{\alpha},
 \end{aligned}$$

where we use L'Hopital's rule in second equality and the third equality holds because of the assumption that $\text{Hess}_W D_{\text{KL}}(\nu \parallel \nu) - 2\alpha G_W(\nu)$ is semi-definite. \square

Remark 14 With the help of (16), readers can recognize that LSI guarantees a global exponential convergence of the gradient flow of the relative entropy. Indeed, suppose μ_t is a gradient flow of $D_{\text{KL}}(\cdot \parallel \nu)$ starting from μ_0 , then we have

$$D_{\text{KL}}(\mu_t \parallel \nu) \leq e^{-2\alpha t} D_{\text{KL}}(\mu_0 \parallel \nu), \quad \mu_0 \in \mathcal{P}(\mathcal{X}) \text{ (LSI}(\alpha)\text{).}$$

While intuitively speaking, a PI can be viewed as an infinitesimal version of a LSI, that is to consider the dynamics in a neighborhood of the optimal value.

B.2. LSIs and PIs in families

Now, it is clear that PIs and LSIs are related to density manifold. We attempt to find those counterparts in statistical models, i.e. submanifolds.

Now, we fix a model $\Theta \subset \mathcal{P}(\mathcal{X})$ with metric given by G_W . The relative entropy is indeed a restriction of global functional to this family. And we furthermore require the reference measure ν to lie in this family, i.e. $\nu = p_{\theta_*}, \theta_* \in \Theta$. We use $\tilde{\cdot}$ to distinguish constraint cases (statistical models) from the global situation (density manifold). Recall that the Fisher information functional is merely the relative entropy dissipation along a gradient flow. Thus we have

$$\begin{aligned} \tilde{I}(p_{\theta_t} \parallel p_{\theta_*}) &= -\frac{d}{dt} D_{\text{KL}}(p_{\theta_t} \parallel p_{\theta_*}) \\ &= g_W(\text{grad}_W D_{\text{KL}}(p_{\theta} \parallel p_{\theta_*}), \text{grad}_W D_{\text{KL}}(p_{\theta} \parallel p_{\theta_*})) \\ &= (\nabla_{\theta} D_{\text{KL}})^T \left(\tilde{G}_W^{-1} \right)^T \tilde{G}_W \tilde{G}_W^{-1} \nabla_{\theta} D_{\text{KL}} \\ &= (\nabla_{\theta} D_{\text{KL}})^T \tilde{G}_W^{-1} \nabla_{\theta} D_{\text{KL}}, \end{aligned} \quad (17)$$

where we use a fact

$$\text{grad}_W D_{\text{KL}}(p_{\theta} \parallel p_{\theta_*}) = \tilde{G}_W^{-1} \nabla_{\theta} D_{\text{KL}}.$$

Definition 20 (LSI in family) Consider a statistical model $p : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$, a probability measure p_{θ_*} is said to satisfy LSI(α) in Θ with constant $\alpha > 0$ (in short: LSI(α)) if we have

$$D_{\text{KL}}(p_{\theta} \parallel p_{\theta_*}) < \frac{1}{2\alpha} \tilde{I}(p_{\theta} \parallel p_{\theta_*}), \quad \theta \in \Theta.$$

Using information matrices, we seek a sufficient condition for LSIs and PIs as proposition 19, i.e.

$$\text{Hess}_W D_{\text{KL}}(p_{\theta} \parallel p_{\theta_*}) \geq 2\alpha \tilde{G}_W(\theta),$$

where we have to take care that the Hessian on LHS is calculated in a submanifold instead of density manifold. Fisher information matrix also comes into this picture, via a decomposition of the Hessian term $\text{Hess}_W D_{\text{KL}}(p_\theta \| p_{\theta_*})$. This point is known as the Ricci-information-Wasserstein (RIW) condition.

Theorem 21 (RIW-condition) *The information matrices criterion for LSI(α) of distribution p_{θ_*} is given by*

$$G_F(\theta) + \nabla_\theta^2 p_\theta \log \frac{p_\theta}{p_{\theta_*}} - \Gamma^W \nabla_\theta D_{\text{KL}}(p_\theta \| p_{\theta_*}) \geq 2\alpha G_W(\theta),$$

where Γ^W s are Christoffel symbols in Wasserstein statistical model Θ , while for PI(α) of distribution p_{θ_*} can be written as

$$G_F(\theta) + \nabla_\theta^2 p_\theta \log \frac{p_\theta}{p_{\theta_*}} \geq 2\alpha G_W(\theta).$$

Remark 15 It can be seen that the condition for log-Sobolev inequalities is much more complicated than that of Poincaré inequalities. For LSIs require a global convexity of the entropy while PIs only correspond to local behavior at the minimum. The most significant change takes place in the Hessian term of entropy, where Wasserstein Christoffel symbols come in.

B.3. Examples in 1-d Family

Both LSIs and PIs can be proved by using Wasserstein and Fisher information matrices. Previously, we have done geometric computations on metric tensor and Hessian of the entropy. This prepares ingredients for us to establish inequalities in families of probability distributions. In this section, we utilize previous calculations to obtain concrete bounds on these functional inequalities.

Example 14 (Gaussian distribution) Recall that for a Gaussian distribution with mean value μ and standard variance σ , the Wasserstein and Fisher information matrices are given by

$$G_W(\mu, \sigma) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad G_F(\mu, \sigma) = \begin{pmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{2}{\sigma^2} \end{pmatrix}.$$

The entropy and the relative entropy defined on this model are provided by

$$H(\mu, \sigma) = -\frac{1}{2} \log 2\pi - \log \sigma - \frac{1}{2},$$

$$D_{\text{KL}}(\mu, \sigma \| p_*) = -\log \sigma + \log \sigma_* - \frac{1}{2} + \frac{\sigma^2 + (\mu - \mu_*)^2}{2\sigma_*^2}, \quad p_* \sim p_{\mu_*, \sigma_*}.$$

We can calculate Wasserstein gradients associated with these two functionals

$$\nabla_{\mu, \sigma}^W H(\mu, \sigma) = \begin{pmatrix} 0 \\ -\frac{1}{\sigma} \end{pmatrix}, \quad \nabla_{\mu, \sigma}^W D_{\text{KL}}(\mu, \sigma \| p_*) = \begin{pmatrix} \frac{\mu - \mu_*}{\sigma_*^2} \\ -\frac{1}{\sigma} + \frac{\sigma}{\sigma_*^2} \end{pmatrix},$$

with the correspondent Fisher information functionals as

$$\begin{aligned} \tilde{I}(\mu, \sigma) &= \frac{1}{\sigma^2}, \\ \tilde{I}(\mu, \sigma \| p_*) &= \frac{(\mu - \mu_*)^2}{\sigma_*^4} + \left(-\frac{1}{\sigma} + \frac{\sigma}{\sigma_*^2} \right)^2. \end{aligned}$$

Thus, the LSI(α) for Gaussian p_{μ_*, σ_*} is given by

$$\frac{(\mu - \mu_*)^2}{\sigma_*^4} + \left(-\frac{1}{\sigma} + \frac{\sigma}{\sigma_*^2} \right)^2 \geq 2\alpha \left(-\log \sigma + \log \sigma_* - \frac{1}{2} + \frac{\sigma^2 + (\mu - \mu_*)^2}{2\sigma_*^2} \right).$$

Next, we move onto the derivation of the RIW condition. It suffices to consider a relation between \tilde{G}_W , $\text{Hess}_W D_{\text{KL}}$ at each point in a statistical model. Recall the formula for Hessian in Riemannian geometry

$$(\text{Hess } f)_{ij} = \partial_i \partial_j f - \Gamma_{ij}^{k(W)} \partial_k f,$$

where Γ^W s are Christoffel symbols in Wasserstein geometry. In Wasserstein Gaussian model where the metric is Euclidean, Christoffel symbols vanish, i.e. $\Gamma^W = 0$. Thus we have

$$\text{Hess}_W H(\mu, \sigma) = \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{\sigma^2} \end{pmatrix}, \quad \text{Hess}_W D_{\text{KL}}(\mu, \sigma \| p_*) = \begin{pmatrix} \frac{1}{\sigma_*^2} & 0 \\ 0 & \frac{1}{\sigma^2} + \frac{1}{\sigma_*^2} \end{pmatrix}.$$

For a gradient flow of the relative entropy w.r.t. a Gaussian p_{μ_*, σ_*} , we conclude that

$$\text{Hess}_W D_{\text{KL}}(\mu, \sigma \| p_{\theta_*}) \geq \left(\frac{1}{\sigma_*^2} \right) G_W(\mu, \sigma),$$

since $G_W(\mu, \sigma)$ is exactly an identity matrix. In other words, the Gaussian p_{μ_*, σ_*} satisfies a LSI($\frac{1}{2\sigma_*^2}$) in a Gaussian model. Notice this result coincides with the one in global case, which is simple consequence of the fact that Gaussian family is totally geodesic in the Wasserstein density manifold.

Next, for the gradient flow of the entropy $H(\cdot)$, we do not have a satisfying constant α such that the Hessian condition proposition 19 holds. For $\text{Hess}_W H(\mu, \sigma)$ matrix has an eigenvalue 0. Despite of this, we have

$$\text{grad}_W H(\mu, \sigma) = G_W^{-1} \nabla_{\mu, \sigma} H(\mu, \sigma) = \nabla_{\mu, \sigma} H(\mu, \sigma),$$

whose μ component always vanishes. Thus the gradient direction of $H(\cdot)$ always coincides with σ direction, in which we have eigenvalue's bound: $\text{eig}_\sigma(H) \geq \frac{1}{\sigma^2} \text{eig}_\sigma(G_W)$. This refers to that eigenvalues of two matrices correspond to direction $\frac{\partial}{\partial \mu}$ have a bound. For LSIs, if the range of σ is the whole \mathbb{R} , then it is easy to see there will not exist a satisfying constant $\alpha > 0$ for $\text{LSI}(\alpha)$ to hold, i.e. $\frac{1}{\sigma^2} \geq 2\alpha, \forall \sigma \in \mathbb{R}$. However, if we restrict the range of σ to a bounded region such as $[-M, M]$, then $\text{LSI}(\frac{1}{2M^2})$ will hold.

Remark 16 Above calculation on gradient flows of the entropy does not establish $\text{LSI}(\alpha)$ for any specific distribution. It merely provides an example of using Hessian condition to study dynamical behaviors.

Example 15 (Laplacian distribution) Consider the case of Laplacian distribution, where

$$G_W(m, \lambda) = \begin{pmatrix} 1 & 0 \\ 0 & \frac{2}{\lambda^4} \end{pmatrix}, \quad G_F(m, \lambda) = \begin{pmatrix} \lambda^2 & 0 \\ 0 & \frac{1}{\lambda^2} \end{pmatrix},$$

from which we can calculate the Christoffel symbol as

$$\begin{aligned} \Gamma_{22}^{2(W)}(m, \lambda) &= \frac{g_{22}^{-1}}{2} (\partial_2 g_{22} + \partial_2 g_{22} - \partial_2 g_{22}) = \frac{g_{22}^{-1}}{2} \partial_2 g_{22} = \frac{\lambda^4}{4} \cdot \left(-\frac{8}{\lambda^5} \right) = -\frac{2}{\lambda}, \\ \Gamma_{ij}^{k(W)}(m, \lambda) &= 0 \quad \text{otherwise.} \end{aligned}$$

Following the same procedure we have done before, the entropy and the relative entropy w.r.t. p_{m_*, λ_*} defined on this model is provided by

$$\begin{aligned} H(m, \lambda) &= -1 + \log \lambda - \log 2, \\ \text{D}_{\text{KL}}(m, \lambda \| p_*) &= -1 + \log \lambda - \log \lambda_* + \lambda_* |m - m_*| + \frac{\lambda_* e^{-\lambda|m - m_*|}}{\lambda}, \end{aligned}$$

from which we can calculate Wasserstein gradients associated with two functionals

$$\begin{aligned} \nabla_{m, \lambda}^W H(m, \lambda) &= \begin{pmatrix} 0 \\ \frac{1}{\lambda} \end{pmatrix}, \\ \nabla_{m, \lambda}^W \text{D}_{\text{KL}}(m, \lambda \| p_*) &= \begin{cases} \begin{pmatrix} \lambda_* (1 - e^{-\lambda(m - m_*)}) \\ - \frac{(\lambda(m - m_*) + 1) \lambda_* e^{-\lambda(m - m_*)} - \lambda}{\lambda^2} \end{pmatrix}, & m > m_*, \\ \begin{pmatrix} -\lambda_* (1 - e^{-\lambda(m_* - m)}) \\ - \frac{(\lambda(m_* - m) + 1) \lambda_* e^{-\lambda(m_* - m)} - \lambda}{\lambda^2} \end{pmatrix}, & m < m_*, \end{cases} \end{aligned}$$

with the Fisher information functionals as

$$\begin{aligned}\widetilde{I}(m, \lambda) &= \frac{\lambda^2}{2}, \\ \widetilde{I}(m, \lambda \| p_*) &= \lambda_*^2 \left(1 - e^{-\lambda|m-m_*|}\right)^2 + \frac{((\lambda|m-m_*|+1)\lambda_* e^{-\lambda|m-m_*|} - \lambda)^2}{2}.\end{aligned}$$

Notice that the value of $\nabla_{m, \lambda}^W D_{\text{KL}}(m, \lambda \| p_*)$ is not well-defined at point $m = m_*$. However, what we considered is integral on the whole \mathbb{R} . Thus we can simply ignore its value at $m = m_*$. As before, LSI(α) is given by

$$\begin{aligned}\lambda_*^2 \left(1 - e^{-\lambda|m_*-m|}\right)^2 + \frac{((\lambda|m_*-m|+1)\lambda_* e^{-\lambda|m_*-m|} - \lambda)^2}{2} \\ \geq 2\alpha \left(-1 + \log \lambda - \log \lambda_* + \lambda_* |m - m_*| + \frac{\lambda_* e^{-\lambda|m-m_*|}}{\lambda}\right).\end{aligned}$$

And we find Hessians of the entropy and the relative entropy in (Θ, G_W) are given by

$$\begin{aligned}\text{Hess}_W H(m, \lambda) &= \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{\lambda^2} \end{pmatrix}, \\ \text{Hess}_W D_{\text{KL}}(m, \lambda \| p_*) &= \begin{pmatrix} \lambda \lambda_* e^{-\lambda|m-m_*|} & 0 \\ 0 & \frac{1}{\lambda^2} + \frac{\lambda_* e^{-\lambda|m-m_*|} (m_* - m)^2}{\lambda^3} \end{pmatrix}.\end{aligned}$$

Following the same analysis, we conclude that for gradient flows of the entropy $D_{\text{KL}}(m, \lambda)$, a LSI($\frac{\lambda^2}{4}$) holds. While for the relative entropy $D_{\text{KL}}(m, \lambda \| p_{m_*, \lambda_*})$, Hessian condition can be written as

$$\begin{pmatrix} \lambda \lambda_* e^{-\lambda|m-m_*|} & 0 \\ 0 & \frac{1}{\lambda^2} + \frac{\lambda_* e^{-\lambda|m-m_*|} (m_* - m)^2}{\lambda^3} \end{pmatrix} \geq \alpha \begin{pmatrix} 1 & 0 \\ 0 & \frac{2}{\lambda^4} \end{pmatrix},$$

which can be reformulated as

$$\alpha = \min_{m, \lambda} \left\{ \lambda \lambda_* e^{-\lambda|m-m_*|}, \frac{1}{2} \left(\lambda^2 + \lambda_* e^{-\lambda|m-m_*|} \lambda (m_* - m)^2 \right) \right\}.$$

From above formula, we conclude that in order to find a satisfying constant, it suffices to restrict the region of $m \in [-M, M]$, $\lambda \in [N, \infty)$. The distribution $\text{La}(m_*, \lambda_*)$ satisfies a LSI(α) in Laplacian family with α given above.

Example 16 (Independent model) For an independent family $p(x, y; \theta) = p_1(x; \theta) p(y; \theta)$, we have

$$G_W = G_W^1 + G_W^2, \quad G_F = G_F^1 + G_F^2.$$

The entropy and the relative entropy also have this separability property

$$\begin{aligned} H(\theta) &= H_1(\theta) + H_2(\theta), \\ D_{\text{KL}}(\theta \| p_*) &= D_{\text{KL}}^1(\theta \| p_{1*}) + D_{\text{KL}}^2(\theta \| p_{2*}), \\ \nabla_{\theta} D_{\text{KL}}(\theta \| p_*) &= \nabla_{\theta} D_{\text{KL}}^1(\theta \| p_{1*}) + \nabla_{\theta} D_{\text{KL}}^2(\theta \| p_{2*}). \end{aligned}$$

The Fisher information functional is given by

$$\begin{aligned} I(p_{\theta} \| p_*) &= \left(\nabla_{\theta} D_{\text{KL}}^1(\theta \| p_{1*}) + \nabla_{\theta} D_{\text{KL}}^2(\theta \| p_{2*}) \right)^T \left(G_W^1 + G_W^2 \right)^{-1} \\ &\quad \left(\nabla_{\theta} D_{\text{KL}}^1(\theta \| p_{1*}) + \nabla_{\theta} D_{\text{KL}}^2(\theta \| p_{2*}) \right), \end{aligned}$$

with LSI(α) given by

$$\begin{aligned} &(\nabla_{\theta} D_{\text{KL}}^1(\theta \| p_{1*}) + \nabla_{\theta} D_{\text{KL}}^2(\theta \| p_{2*}))^T (G_W^1 + G_W^2)^{-1} (\nabla_{\theta} D_{\text{KL}}^1(\theta \| p_{1*}) + \nabla_{\theta} D_{\text{KL}}^2(\theta \| p_{2*})) \\ &\geq 2\alpha \nabla_{\theta} D_{\text{KL}}^1(\theta \| p_{1*}) + \nabla_{\theta} D_{\text{KL}}^2(\theta \| p_{2*}). \end{aligned}$$

In conclusion, above examples introduce another way to prove functional inequalities as well as convergence rates of dynamics in probability families.

Appendix C. Proofs in Sect. 4

C.1. Proof of Theorem 13

Proof of Theorem 13 First, we postulate that ∇_x refers to the gradient w.r.t. x variable while ∇_{θ} refers to the gradient w.r.t. θ variable. We expand the function $l(x_t, \theta_t)$

$$l(x_t, \theta_t) = l(x_t, \theta_*) + \nabla_{\theta} l(x_t, \theta_*) (\theta_t - \theta_*) + O(|\theta_t - \theta_*|^2).$$

By substrating θ_* in both sides of the updating equation and plugging in the expansion above, we get

$$\begin{aligned} \theta_{t+1} - \theta_* &= (\theta_t - \theta_*) - \frac{1}{t} G_W^{-1}(\theta_t) (l(x_t, \theta_*) + \nabla_{\theta} l(x_t, \theta_*) (\theta_t - \theta_*)) \\ &\quad + O(|\theta_t - \theta_*|^2). \end{aligned}$$

Then, taking Wasserstein covariances of both sides, we get

$$\begin{aligned} V_{t+1} &= V_t + \frac{1}{t^2} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x \left(G_W^{-1}(\theta_t) l(x_t, \theta_t) \right) \cdot \nabla_x \left(l(x_t, \theta_t)^T G_W^{-1}(\theta_t) \right) \right] \\ &\quad - \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x \left(l(x_t, \theta_*)^T G_W^{-1}(\theta_t) \right) \right] + o \left(\frac{V_t}{t} \right) \\ &\quad - \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x \left((\theta_t - \theta_*)^T \nabla_\theta l(x_t, \theta_*)^T G_W^{-1}(\theta_t) \right) \right], \end{aligned}$$

where the last term corresponds to an expansion term $O(|\theta_t - \theta_*|^2)$ and we use an assumption that $\mathbb{E}_{p_{\theta_*}} [(\theta_t - \theta_*)^2], \mathbb{E}_{p_{\theta_*}} [|\nabla_x (\theta_t - \theta_*)|^2] = o(1)$. In above formula, we eliminate transpose symbols T on metric tensor G_W because of its symmetry. For the second term on the RHS, we have

$$\begin{aligned} &\frac{1}{t^2} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x \left(G_W^{-1}(\theta_t) l(x_t, \theta_t) \right) \cdot \nabla_x \left(l(x_t, \theta_t)^T G_W^{-1}(\theta_t) \right) \right] \\ &= \frac{1}{t^2} \mathbb{E}_{p_{\theta_*}} \left[G_W^{-1}(\theta_*) \nabla_x (l(x_t, \theta_*)) \cdot \nabla_x \left(l(x_t, \theta_*)^T \right) G_W^{-1}(\theta_*) \right] + o \left(\frac{1}{t^2} \right) \\ &= \frac{1}{t^2} G_W^{-1}(\theta_*) \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (l(x_t, \theta_*)) \cdot \nabla_x \left(l(x_t, \theta_*)^T \right) \right] G_W^{-1}(\theta_*) + o \left(\frac{1}{t^2} \right), \end{aligned}$$

where we use the following fact

$$\begin{aligned} &\mathbb{E}_{p_{\theta_*}} \left[\nabla_x \left(G_W^{-1}(\theta_t) l(x_t, \theta_t) \right) \cdot \nabla_x \left(l(x_t, \theta_t)^T G_W^{-1}(\theta_t) \right) \right] \\ &\quad - \mathbb{E}_{p_{\theta_*}} \left[G_W^{-1}(\theta_*) \nabla_x (l(x_t, \theta_*)) \cdot \nabla_x \left(l(x_t, \theta_*)^T \right) G_W^{-1}(\theta_*) \right] \\ &= O \left(\mathbb{E}_{p_{\theta_*}} [|\theta_t - \theta_*|] \right) = o(1). \end{aligned}$$

And the third term in the RHS can be reduced according have

$$\begin{aligned} &-\frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x \left(l(x_t, \theta_*)^T G_W^{-1}(\theta_t) \right) \right] \\ &= -\frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x \left(l(x_t, \theta_*)^T \right) G_W^{-1}(\theta_t) \right] \\ &\quad - \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot l(x_t, \theta_*)^T \nabla_x G_W^{-1}(\theta_t) \right] \\ &= 0, \end{aligned}$$

where the first term vanishes because $\nabla_x (\theta_t - \theta_*)$ only has non-vanishing components at x_1, \dots, x_{t-1} while $\nabla_x (l(x_t, \theta_*)^T)$ only has a non-vanishing component at x_t . Consequently their inner product vanishes everywhere. While the second term vanishes by considering each element of this matrix, we have

$$\begin{aligned}
& \left(\mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot l(x_t, \theta_*)^T \nabla_x G_W^{-1}(\theta_t) \right] \right)_{ij} \\
&= \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \nabla_x (\theta_t - \theta_*)_i \cdot \left(l(x_t, \theta_*)^T \nabla_x G_W^{-1}(\theta_t) \right)_j \\
&= \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*)_i \cdot \nabla_x \left(G_W^{-1}(\theta_t)_{kj} \right) l(x_t, \theta_*)_k^T \right] \\
&= \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*)_i \cdot \nabla_x \left(G_W^{-1}(\theta_t)_{kj} \right) \right] \mathbb{E}_{p_{\theta_*}} l(x_t, \theta_*)_k^T \\
&= 0,
\end{aligned}$$

where the third equality is guaranteed by the fact that $\theta_t - \theta_*$ is independent to $\nabla_\theta f(x_t, \theta_*)$ since θ_t, x_t are mutually independent. While the last equality holds by an assumption as

$$\mathbb{E}_{p_{\theta_*}} l(x_t, \theta_*) = 0.$$

For the last term, same as the analysis of the third term, we find

$$\begin{aligned}
& -\frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x \left((\theta_t - \theta_*)^T \nabla_\theta l(x_t, \theta_*)^T G_W^{-1}(\theta_t) \right) \right] \\
&= -\frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x \left((\theta_t - \theta_*)^T \right) \nabla_\theta l(x_t, \theta_*)^T G_W^{-1}(\theta_t) \right] \\
&\quad - \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot (\theta_t - \theta_*)^T \nabla_x \left(\nabla_\theta l(x_t, \theta_*)^T \right) G_W^{-1}(\theta_t) \right] \\
&\quad - \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot (\theta_t - \theta_*)^T \nabla_\theta l(x_t, \theta_*)^T \nabla_x \left(G_W^{-1}(\theta_t) \right) \right] \\
&= -\frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x \left((\theta_t - \theta_*)^T \right) \nabla_\theta l(x_t, \theta_*)^T G_W^{-1}(\theta_*) \right] + o\left(\frac{V_t}{t}\right) \\
&\quad - \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot (\theta_t - \theta_*)^T \nabla_\theta l(x_t, \theta_*)^T \nabla_x \left(G_W^{-1}(\theta_t) \right) \right],
\end{aligned}$$

where we again use the independent relation between $(\theta_t - \theta_*)$ and $f(x_t, \theta_*)$. The additional term appearing above, with the help that $\mathbb{E}_{p_{\theta_*}} [\nabla_\theta f(x_t, \theta_*)] = O(1)$, $\nabla_x \left(G_W^{-1}(\theta_t) \right) = \nabla_x \theta_t \nabla_\theta \left(G_W^{-1}(\theta_t) \right) = O(\nabla_x (\theta_t - \theta_*))$, can be further reduced to the form below

$$\begin{aligned}
& \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot (\theta_t - \theta_*)^T \nabla_\theta l(x_t, \theta_*) \nabla_x \left(\left(G_W^{-1}(\theta_t) \right) \right) \right] \\
&= \frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot (\theta_t - \theta_*)^T O(1) \nabla_x (\theta_t - \theta_*) O(1) \right] \\
&\leq \frac{O(1)}{t} \sqrt{\mathbb{E}_{p_{\theta_*}} [|\nabla_x (\theta_t - \theta_*)|^2] \mathbb{E}_{p_{\theta_*}} [(\theta_t - \theta_*)^2] \mathbb{E}_{p_{\theta_*}} [|\nabla_x (\theta_t - \theta_*)|^2]} \\
&= o\left(\frac{V_t}{t}\right).
\end{aligned}$$

And the last term finally reduces have

$$\begin{aligned} & -\frac{2}{t} \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (\theta_t - \theta_*) \cdot \nabla_x \left((\theta_t - \theta_*)^T \right) \nabla_\theta l(x_t, \theta_*) \left(G_W^{-1}(\theta_*) \right) \right] + o\left(\frac{V_t}{t}\right) \\ & = -\frac{2V_t}{t} \mathbb{E}_{p_{\theta_*}} [\nabla_\theta l(x_t, \theta_*)] G_W^{-1}(\theta_*) + o\left(\frac{V_t}{t}\right). \end{aligned}$$

Combining all the terms we have in hand, we derive the following updating equation for Wasserstein covariances during a natural gradient descent

$$\begin{aligned} V_{t+1} &= V_t + \frac{1}{t^2} G_W^{-1}(\theta_*) \mathbb{E}_{p_{\theta_*}} \left[\nabla_x (l(x_t, \theta_*)) \cdot \nabla_x \left(l(x_t, \theta_*)^T \right) \right] \left(G_W^{-1}(\theta_*) \right) \\ & - \frac{2V_t}{t} \mathbb{E}_{p_{\theta_*}} [\nabla_\theta l(x_t, \theta_*)] G_W^{-1}(\theta_*) + o\left(\frac{V_t}{t}\right) + o\left(\frac{1}{t^2}\right) + O\left(\frac{V_t}{t^2}\right). \end{aligned}$$

Remark 17 The most frequently used tools in this proof is a separability property, c.f. proposition 5. The key observation is that, for two statistics T_1, T_2 which depend on (independent) different variables, such as $T_1 = T_1(x_1, \dots, x_{t-1})$, $T_2 = T_2(x_t, \dots, x_{t+n})$ are “orthogonal” in both Wasserstein and Fisher metrics. Specifically, consider gradients of T_1, T_2 w.r.t. x , since they depend on different variables, thus

$$\text{Cov}^W [T_1, T_2] = \mathbb{E}_{p_{\theta_*}} [\nabla_x T_1 \cdot \nabla_x T_2] = 0.$$

This type of separability is a direct analog of the one in Fisher-Rao geometry, i.e.

$$\text{Cov}^F [T_1, T_2] = \mathbb{E}_{p_{\theta_*}} [T_1 T_2] = \mathbb{E}_{p_{\theta_*}} [T_1] \cdot \mathbb{E}_{p_{\theta_*}} [T_2] = 0.$$

C.2. Examples and numerical experiments of two efficiencies

Example 17 (Gaussian distribution) Consider the Gaussian distribution with mean value μ and standard variance σ as

$$p(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2}(x-\mu)^2}.$$

The WIM satisfies

$$G_W(\mu, \sigma) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

The Fisher information matrix satisfies

$$G_F(\mu, \sigma) = \begin{pmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{2}{\sigma^2} \end{pmatrix}.$$

Further, the matrix $G_F G_W^{-1}$ is given by

$$G_F(\mu, \sigma) G_W^{-1}(\mu, \sigma) = \begin{pmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{2}{\sigma^2} \end{pmatrix}.$$

And optimal parameters are given by μ_* , σ_* . Thus we have following conclusions on efficiency of the Fisher, Wasserstein natural gradients and Wasserstein natural gradient on Fisher score (maximal likelihood estimator).

The Wasserstein natural gradient is asymptotically efficient with an asymptotic Wasserstein covariance given by

$$V_t = \frac{1}{t} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + O\left(\frac{1}{t^2}\right).$$

The Fisher natural gradient is asymptotic efficient with an asymptotic classical covariance given by

$$V_t = \frac{1}{t} \begin{pmatrix} \sigma_*^2 & 0 \\ 0 & \frac{\sigma_*^2}{2} \end{pmatrix} + O\left(\frac{1}{t^2}\right).$$

An interesting thing is that the covariance matrix appears in the Wasserstein efficiency is independent of the optimal value. While in Fisher case, the asymptotic behavior depends a lot on the optimal parameter we obtain.

For the last case, in the Gaussian family, two metric tensors G_F , G_W can be simultaneously diagonalized, thus the situation is even simpler. We denote the smallest eigenvalue of $G_F G_W^{-1}$ as α , i.e.

$$\alpha = \frac{1}{\sigma_*^2}.$$

Furthermore, we have to figure out the term

$$\mathbb{E}_{p_{\mu*, \sigma*}} \left[\nabla_x \left(\nabla_{\mu*, \sigma*} l(x_t, \mu_*, \sigma_*) \right) \cdot \nabla_x \left(\nabla_{\mu*, \sigma*} l(x_t, \mu_*, \sigma_*)^T \right) \right],$$

that appears in the final result. In Gaussian family, since we have Fisher scores $\nabla_{\mu*, \sigma*} l(x; \theta) = \Phi^F(x, \mu_*, \sigma_*)$ as

$$\Phi_\mu^F(x; \mu, \sigma) = \frac{x - \mu}{\sigma^2}, \quad \Phi_\sigma^F(x; \mu, \sigma) = \frac{(x - \mu)^2}{\sigma^3} - \frac{1}{\sigma}.$$

Via calculation, we have

$$\begin{aligned}\mathbb{E}_{p_{\mu_*, \sigma_*}} \left[\nabla_x \Phi_\mu^F(x; \mu_*, \sigma_*) \cdot \nabla_x \left(\Phi_\mu^F(x; \mu_*, \sigma_*)^T \right) \right] &= \mathbb{E}_{p_{\mu_*, \sigma_*}} \left[\frac{1}{\sigma_*^4} \right] = \frac{1}{\sigma_*^4}, \\ \mathbb{E}_{p_{\mu_*, \sigma_*}} \left[\nabla_x \Phi_\mu^F(x; \mu_*, \sigma_*) \cdot \nabla_x \left(\Phi_\sigma^F(x; \mu_*, \sigma_*)^T \right) \right] &= \mathbb{E}_{p_{\mu_*, \sigma_*}} \left[\frac{1}{\sigma_*^2} \cdot \frac{2(x - \mu_*)}{\sigma_*^3} \right] \\ &= 0, \\ \mathbb{E}_{p_{\mu_*, \sigma_*}} \left[\nabla_x \Phi_\sigma^F(x; \mu_*, \sigma_*) \cdot \nabla_x \left(\Phi_\sigma^F(x; \mu_*, \sigma_*)^T \right) \right] &= \mathbb{E}_{p_{\mu_*, \sigma_*}} \left[\frac{4(x - \mu_*)^2}{\sigma_*^6} \right] \\ &= \frac{4}{\sigma_*^4},\end{aligned}$$

we conclude the middle term is given by

$$\mathfrak{I} = \mathbb{E}_{p_{\mu_*, \sigma_*}} \left[\nabla_x \left(\nabla_{\mu_*, \sigma_*} l(x_t, \mu_*, \sigma_*) \right) \cdot \nabla_x \left(\nabla_{\mu_*, \sigma_*} l(x_t, \mu_*, \sigma_*)^T \right) \right] = \begin{pmatrix} \frac{1}{\sigma_*^4} & 0 \\ 0 & \frac{4}{\sigma_*^4} \end{pmatrix}.$$

And when we have $\frac{2}{\sigma_*^2} > 1$, the inverse matrix of $2B - \mathbf{I}$ is given by

$$(2B - \mathbf{I})^{-1} = \begin{pmatrix} \frac{\sigma_*^2}{2-\sigma_*^2} & 0 \\ 0 & \frac{\sigma_*^2}{4-\sigma_*^2} \end{pmatrix}.$$

Consequently, the term appearing in the asymptotic behavior of the Poincaré efficiency is given by

$$\begin{aligned}\frac{1}{t} \left(2G_F G_W^{-1} - \mathbf{I} \right)^{-1} G_W^{-1}(\theta_*) \mathfrak{I} \left(G_W^{-1}(\theta_*) \right) &= \begin{pmatrix} \frac{\sigma_*^2}{2-\sigma_*^2} & 0 \\ 0 & \frac{\sigma_*^2}{4-\sigma_*^2} \end{pmatrix} \begin{pmatrix} \frac{1}{\sigma_*^4} & 0 \\ 0 & \frac{4}{\sigma_*^4} \end{pmatrix} \\ &= \begin{pmatrix} \frac{1}{(2-\sigma_*^2)\sigma_*^2} & 0 \\ 0 & \frac{4}{(4-\sigma_*^2)\sigma_*^2} \end{pmatrix}.\end{aligned}$$

Thus the asymptotic behavior of the Wasserstein covariance in the Wasserstein natural gradient of Fisher scores is given by

$$V_t = \begin{cases} O\left(t^{-\frac{2}{\sigma_*^2}}\right), & \frac{1}{\sigma_*^2} \leq \frac{1}{2}, \\ \frac{1}{t} \begin{pmatrix} \frac{1}{(2-\sigma_*^2)\sigma_*^2} & 0 \\ 0 & \frac{4}{(4-\sigma_*^2)\sigma_*^2} \end{pmatrix} + O\left(\frac{1}{t^2}\right), & \frac{1}{\sigma_*^2} > \frac{1}{2}. \end{cases}$$

References

1. Amari, S.: Differential-geometrical methods in statistics. Number 28 in Lecture Notes in Statistics. Springer, Berlin; New York, corr. 2nd print edition (1990)
2. Amari, S.: Natural gradient works efficiently in learning. *Neural Comput.* **10**(2), 251–276 (1998)
3. Amari, S.: Information Geometry and Its Applications. Number volume 194 in Applied mathematical sciences. Springer, Japan (2016)
4. Amari, S., Matsuda, T.: Wasserstein statistics in one-dimensional location-scale model. *Ann. Inst. Stat. Math.* **74**, 33–47 (2022)
5. Arbel, M., Gretton, A., Li, W., Montufar, G.: Kernelized wasserstein natural gradient. *Int. Conf. Learn. Represent.* (2020)
6. Ay, N., Jost, J., Ván Lé, H., Schwachhöfer, L.: Information Geometry, volume 64. Springer (2017)
7. Bernton, E., Jacob, P.E., Gerber, M., Robert, C.P.: On parameter estimation with the Wasserstein distance. *Inform. Inference A J. IMA* **8**(4), 657–676 (2019)
8. Blanchet, J., Murthy, K., Nguyen, V.A.: Statistical analysis of Wasserstein distributionally robust estimators. In: *Tutorials in Operations Research: Emerging Optimization Methods and Modeling Techniques with Applications*, pp 227–254. INFORMS (2021)
9. Briol, F.-X., Barp, A., Duncan, A.B., Girolami, M.: Statistical inference for generative models with maximum mean discrepancy (2019). arXiv preprint [arXiv:1906.05944](https://arxiv.org/abs/1906.05944)
10. Casella, G., Berger, R.L.: Statistical Inference, vol. 2. Duxbury, Pacific Grove (2002)
11. Chen, Y., Li, W.: Optimal transport natural gradient for statistical manifolds with continuous sample space. *Inform. Geom.* **3**, 1–32 (2020)
12. Cover, T.M., Thomas, J.A.: Elements of information theory. Wiley-Interscience, Hoboken, N.J, 2nd ed edition (2006)
13. Kriegel, A., Michor, P.W.: The convenient setting of global analysis, volume 53. American Mathematical Soc (1997)
14. Lafferty, J.D.: The density manifold and configuration space quantization. *Trans. Am. Math. Soc.* **305**(2), 699–741 (1988)
15. Li, W.: Transport information geometry: Riemannian calculus on probability simplex. *Inform. Geom.* **5**, 161–207 (2022)
16. Li, W., Montúfar, G.: Natural gradient via optimal transport. *Inform. Geom.* **1**, 181–214 (2018)
17. Li, W., Montúfar, G.: Ricci curvature for parametric statistics via optimal transport. *Inform. Geom.* **3**, 89–117 (2020)
18. Li, W., Liu, S., Zha, H., Zhou, H.: Parametric fokker-planck equation. *Geom. Sci. Inform.*, 715–724 (2019)
19. Li, W., Lin, A.T., Montúfar, G.: Affine natural proximal learning. *Geomet. Sci. Inform.*, 705–714 (2019)
20. Lin, A.T., Li, W., Osher, S., Montufar, G.: Wasserstein proximal of GANs. *Geom. Sci. Inform.*, 524–533 (2021)
21. Lott, J.: Some geometric calculations on Wasserstein space. *Commun. Math. Phys.* **277**, 423–437 (2008)
22. Mallasto, A., Hajie, T.D., Feragen, A.: A formalization of the natural gradient method for general similarity measures. *Geom. Sci. Inform.*, 599–607 (2019)
23. Nielsen, F.: On voronoi diagrams on the information-geometric Cauchy manifolds. *Entropy* **22**(7), 713 (2020)
24. Ollivier, Y.: Online natural gradient as a Kalman filter. *Electron. J. Stat.* **12**(2), 2930–2961 (2018)
25. Otto, F.: The geometry of dissipative evolution equations the porous medium equation. *Commun. Part. Differ. Equ.* **26**(1–2), 101–174 (2001)
26. Otto, F., Villani, C.: Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality. *J. Funct. Anal.* **173**(2), 361–400 (2000)
27. Petersen, A., Müller, H.-G.: Wasserstein covariance for multiple random densities. *Biometrika* **106**(2), 339–351 (2019)
28. Villani, C.: Topics in optimal transportation. Number 58. American Mathematical Soc., (2003)
29. Villani, C.: Optimal transport: old and new, volume 338. Springer (2008)
30. Wong, T.-K.L.: Logarithmic divergences from optimal transport and Rényi geometry. *Inform. Geom.* **1**(1), 39–78 (2018)

31. Zozor, S., Brossier, J.-M.: Debruijn identities: from shannon, Kullback–Leibler and Fisher to generalized φ -entropies, φ -divergences and φ -fisher informations. In: AIP Conference Proceedings **1641**, 522–529. AIP (2015)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.