



Statistics

## Entropic optimal transport is maximum-likelihood deconvolution



*Le transport optimal entropique correspond à l'estimateur du maximum de vraisemblance en déconvolution*

Philippe Rigollet, Jonathan Weed

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA

### ARTICLE INFO

#### Article history:

Received 7 September 2018

Accepted after revision 6 November 2018

Available online 13 November 2018

Presented by the Editorial Board

### ABSTRACT

We give a statistical interpretation of entropic optimal transport by showing that performing maximum-likelihood estimation for Gaussian deconvolution corresponds to calculating a projection with respect to the entropic optimal transport distance. This structural result gives theoretical support for the wide adoption of these tools in the machine learning community.

© 2018 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

### R É S U M É

Cette note donne un interprétation statistique du transport optimal entropique : on montre que l'estimateur du maximum de vraisemblance en déconvolution gaussienne correspond à la projection de la loi empirique des données au sens de la distance définie par le transport optimal entropique. Ce résultat structurel donne une justification théorique, qui soutient l'adoption massive de ces outils par la communauté de l'apprentissage automatique.

© 2018 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

## 1. Introduction

Optimal transport is a fundamental notion with arising in several branches of mathematics, including probability, analysis and statistics. More recently, it has found new applications in computational domains such as machine learning and image processing [2,6,12,23,30,36,37,41]. This newfound utility was largely fueled by algorithmic advances allowing optimal transport distances to be computed quickly between large scale discrete distributions [31]. At the heart of these algorithmic techniques is the idea of entropic penalization, which has been leveraged to obtain near-linear-time approximation schemes for optimal transport distances [1,13,42]. Hereafter, we refer to this technique as *entropic optimal transport*. This line of well

E-mail addresses: rigollet@math.mit.edu (P. Rigollet), jweed@mit.edu (J. Weed).

established computational research stands in sharp contrast with our *statistical* understanding of regularization for optimal transport, which is still in its infancy [18]. Entropic regularization has been shown to play a central statistical role in a variety of problems related to model selection [14–16,24,32–34] but our knowledge of its effect on optimal transport is currently limited to experimental evidence [13,31] without theoretical support.

In this note, we give a statistical interpretation of entropic optimal transport, showing that under some modeling assumptions, it corresponds to the objective function in maximum-likelihood estimation for deconvolution problems involving additive Gaussian noise. This interpretation provides a first indication that optimal transport problems where data is subject to Gaussian observation error should be handled with entropic regularization. Moreover, our results indicate that in the same context, a relaxed version of optimal transport should be preferred, as it is equivalent to maximum-likelihood estimation even in absence of said modeling assumptions.

## 2. Entropic optimal transport

Throughout we denote by  $\gamma$  a probability measure on  $\mathbb{R}^d \times \mathbb{R}^d$  and by  $\|\cdot\|$  the Euclidean norm over  $\mathbb{R}^d$ . Given such a measure, we denote by  $\pi_X\gamma$  and  $\pi_Y\gamma$  the two measures on  $\mathbb{R}^d$  obtained by projecting onto the first and second component, respectively. Given probability measures  $\mu$  and  $\nu$  on  $\mathbb{R}^d$ , define

$$\mathcal{M}(\mu, \nu) := \{\gamma : \pi_X\gamma = \mu, \pi_Y\gamma = \nu\} \quad \text{and} \quad \mathcal{M}(\mu) := \{\gamma : \pi_Y\gamma = \mu\}.$$

We also recall the definition of Kullback–Leibler (KL) divergence between probability measures  $\mu$  and  $\nu$ :

$$D(\mu \parallel \nu) = \begin{cases} \int \log\left(\frac{d\mu}{d\nu}\right) d\mu, & \text{if } \mu \ll \nu \\ \infty, & \text{otherwise.} \end{cases}$$

**Definition 2.1.** The *entropic optimal transport distance* between  $\mu$  and  $\nu$  is

$$W_{\sigma^2}(\mu, \nu) := \min_{\gamma \in \mathcal{M}(\mu, \nu)} \int \frac{1}{2} \|x - y\|^2 d\gamma(x, y) + \sigma^2 I(\gamma), \tag{1}$$

where  $I(\gamma)$  is the *mutual information* defined by

$$I(\gamma) := D(\gamma \parallel \pi_X\gamma \otimes \pi_Y\gamma).$$

Note that when  $\sigma = 0$ , this corresponds the squared 2-Wasserstein distance between probability measures over  $\mathbb{R}^d$  [38]. When  $\sigma > 0$ ,  $W_{\sigma^2}$  no longer satisfies the axioms of a (squared) distance, but it still possesses useful distance-like properties [13]. We employ the term “distance” for all values of  $\sigma$  for terminological consistency. When  $\mu$  and  $\nu$  are discrete measures, the minimizer of (1) is also discrete and agrees with the minimizer of

$$\min_{\gamma \in \mathcal{M}(\mu, \nu)} \int \frac{1}{2} \|x - y\|^2 d\gamma(x, y) - \sigma^2 H(\gamma), \tag{2}$$

where  $H$  is the standard Shannon entropy:

$$H(\gamma) := \sum_{ij} \gamma_{ij} \log \frac{1}{\gamma_{ij}},$$

where  $\gamma_{ij} := \gamma(x_i, y_j)$  for  $(x_i, y_j) \in \text{supp}(\gamma)$ . Note that (2) is the definition proposed by [13] for discrete measures, whereas (1) corresponds to the appropriate generalization studied in [20] for measures that are not necessarily discrete. It is not hard to check that  $W_{\sigma^2}(\mu, \nu) < \infty$  for any pair  $(\mu, \nu)$  possessing finite second moments, since the independent coupling  $\mu \otimes \nu$  in the minimization that appears in (1) leads to a finite objective value.

A maximum-likelihood interpretation of entropic optimal transport is already known in the context of a large-deviation principle for Brownian motion [26]. In this context, given two distributions  $\mu$  and  $\nu$  (viewed as the positions of particles at times  $t = 0$  and  $t = 1$ ), Schrödinger [40] gave a heuristic argument motivated by statistical physics establishing that the law of independent particles undergoing Brownian motion conditioned on having initial and final distributions  $\mu$  and  $\nu$  respectively is induced by the solution to the optimal transport problem with entropic regularization [26]. While suggestive, this interpretation does not hold any immediate implications for estimation problems where only data is available rather than distributions  $\mu$  and  $\nu$ . Below, we introduce the classical deconvolution model for corrupted observations and show that in this context, entropic optimal transport is precisely the maximum-likelihood estimator.

### 3. Deconvolution

Let  $\mathcal{P}$  be a given family of probability distributions over  $\mathbb{R}^d$  with finite second moments and let  $P^*$  be an unknown distribution, also with finite second moment. The deconvolution problem consists in estimating  $P^*$  on the basis of corrupted observations  $Y_1, \dots, Y_n$ , where

$$Y_i = X_i + Z_i, \quad i = 1, \dots, n, \quad (3)$$

and the errors  $Z_1, \dots, Z_n$  are independent of  $X_1, \dots, X_n$ . For identifiability purposes, the random variables  $\{Z_i\}$  are assumed to be independent copies of a random variable  $Z$  with known distribution:  $Z \sim \mathcal{N}(0, \sigma^2)$  where the variance  $\sigma^2$  is known.

In this context, the distribution of  $Y_i$  admits a density  $\varphi_\sigma \star dP^*$  with respect to  $\lambda$  where, for any  $P \in \mathcal{P}$ , we define

$$\varphi_\sigma \star dP(y) = \int \varphi_\sigma(y - x) dP(x) \quad (4)$$

and  $\varphi_\sigma$  denotes the density of  $Z \sim \mathcal{N}(0, \sigma^2)$ . Under these assumptions, we call (3) the *Gaussian deconvolution model*.

Deconvolution is a classical question of nonparametric statistics [7,8,17] and is core to mixture models [28] as well as statistical models with measurement errors [9]. As such, it has received significant attention from the statistics literature. More recently, it was shown that deconvolution has strong methodological and mathematical connections to optimal transport in the context of a problem known as uncoupled regression [35].

A natural candidate to estimate  $P^*$  is the maximum-likelihood estimator (MLE)  $\hat{P}$  defined by

$$\hat{P} = \operatorname{argmax}_{P \in \mathcal{P}} \sum_{i=1}^n \log \varphi_\sigma \star dP(Y_i). \quad (5)$$

The statistical properties of the MLE are well known and have been established under general conditions on the class  $\mathcal{P}$  [5, 22]. In the next section, we show that entropic optimal transport is in fact implementing  $\hat{P}$ .

### 4. Entropic optimal transport is maximum-likelihood deconvolution

In this section, we adopt the Gaussian deconvolution model (3) of the previous section. The extension to other distributions for the corruption errors  $\{Z_i\}$  is postponed to Section 5.

Our main result involves families of distributions satisfying a particular closure condition.

**Definition 4.1.** A class  $\mathcal{P}$  of probability measures is said to be *closed under domination* if  $Q \ll P$  for some  $P \in \mathcal{P}$  implies that  $Q \in \mathcal{P}$ .

This definition appears to be new. However, many families of distributions common in statistics are closed under domination. For example the class of all measures, the class of all measures absolutely continuous with respect to some reference measure  $\sigma$ , the class of discrete measures, and the set of measures whose support is finite or contains at most  $k$  points all possess this property. A class  $\mathcal{P}$  of probability measures may always be augmented to be closed under domination by adding to it the set of probability measures  $\bigcup_{P \in \mathcal{P}} \{Q : Q \ll P\}$ . The extension to families not closed under domination is considered in Section 5.

We are now in a position to state our main result: a structural representation of the maximum-likelihood estimator  $\hat{P}$  in terms of entropic optimal transport.

**Theorem 4.2.** Let  $\mathcal{P}$  be a class of probability measures that is closed under domination and assume the Gaussian deconvolution model (3). Then the maximum-likelihood estimator  $\hat{P}$  over  $\mathcal{P}$  defined in (5) satisfies:

$$\hat{P} = \operatorname{argmin}_{P \in \mathcal{P}} W_{\sigma^2} \left( P, \frac{1}{n} \sum_{i=1}^n \delta_{y_i} \right).$$

In other words, the maximum-likelihood estimator  $\hat{P}$  is the projection of the empirical measure  $\frac{1}{n} \sum_{i=1}^n \delta_{y_i}$  onto  $\mathcal{P}$  with respect to the entropic optimal transport distance  $W_{\sigma^2}$ .

The projection estimator  $\operatorname{argmin}_{P \in \mathcal{P}} W_{\sigma^2} \left( P, \frac{1}{n} \sum_{i=1}^n \delta_{y_i} \right)$  has been employed in the machine learning community [21,29] as a smoothed version of a minimum Kantorovich distance estimator [3] more suitable for optimization. Theorem 4.2 shows that this estimator has a statistical interpretation in addition to its computational benefits.

As noted above, in the special case when  $\sigma^2 = 0$ , the quantity  $W_{\sigma^2}$  reduces to the squared 2-Wasserstein distance  $W$ . In the context of Gaussian mixture models, when  $\mathcal{P}$  is the class of probability distributions supported on at most  $k$  points,

solving  $\operatorname{argmin}_{P \in \mathcal{P}} W(P, \frac{1}{n} \sum_{i=1}^n \delta_{y_i})$  corresponds to performing a “hard” clustering of the data by minimizing the  $k$ -means objective. It is known, however, that hard  $k$ -means clustering does not lead to consistent estimation of the centroids in a mixture of Gaussians model, whereas consistent estimation *can* be achieved with the MLE, which induces a relaxed “soft” clustering [25]. Theorem 4.2 implies that replacing  $W$  by  $W_{\sigma^2}$  precisely corresponds to this relaxation.

**Proof.** Write for simplicity  $\ell_P := \log \varphi_\sigma \star dP$ . By (4), we have

$$\ell_P(y_i) = C + \log \int \exp\left(-\frac{1}{2\sigma^2} \|x - y_i\|^2\right) dP(x),$$

where  $C$  is a constant not depending on  $y_i$  or  $P$ . The Gibbs variational principle [10, Equation (5.2.1)] then implies that

$$\ell_P(y_i) = C - \frac{1}{\sigma^2} \min_{Q_i} \left\{ \frac{1}{2} \mathbb{E}_{x \sim Q_i} \|x - y_i\|^2 + \sigma^2 D(Q_i \| P) \right\},$$

where the minimization is taken over probability measures on  $\mathbb{R}^d$ . Then, by definition of the MLE, we have

$$\hat{P} = \operatorname{argmin}_{P \in \mathcal{P}} \min_{Q_1, \dots, Q_n} \frac{1}{n} \sum_{i=1}^n \left[ \frac{1}{2} \mathbb{E}_{x \sim Q_i} \|x - y_i\|^2 + \sigma^2 D(Q_i \| P) \right].$$

Next, given any set of  $n$  distributions  $\{Q_1, \dots, Q_n\}$  on  $\mathbb{R}^d$ , we can define the joint probability measure  $\bar{\gamma}$  on  $\mathbb{R}^d \times \{y_1, \dots, y_n\}$  by

$$\bar{\gamma} := \frac{1}{n} \sum_{i=1}^n Q_i \otimes \delta_{y_i}.$$

Note that  $\pi_Y \bar{\gamma} = U$ , the uniform distribution on  $\{y_1, \dots, y_n\}$ . Conversely, for any joint probability measure  $\gamma$  on  $\mathbb{R}^d \times \mathbb{R}^d$  satisfying  $\pi_Y \gamma = U$ , we can decompose  $\gamma$  as  $\frac{1}{n} \sum_{i=1}^n Q_i \otimes \delta_{y_i}$  for some  $Q_1, \dots, Q_n$ . This bijection between sets of  $n$  probability measures  $\{Q_1, \dots, Q_n\}$  on  $\mathbb{R}^d$  and joint measures  $\gamma \in \mathcal{M}(U)$  satisfies the equality

$$\frac{1}{n} \sum_{i=1}^n \left[ \frac{1}{2} \mathbb{E}_{x \sim Q_i} \|x - y_i\|^2 + \sigma^2 D(Q_i \| P) \right] = \frac{1}{2} \mathbb{E}_{(x,y) \sim \gamma} \|x - y\|^2 + \sigma^2 D(\bar{\gamma} \| P \otimes U).$$

We can therefore leverage this bijection to rewrite  $\hat{P}$  as

$$\hat{P} = \operatorname{argmin}_{P \in \mathcal{P}} \min_{\gamma \in \mathcal{M}(U)} \left\{ \frac{1}{2} \mathbb{E}_{(x,y) \sim \gamma} \|x - y\|^2 + \sigma^2 D(\gamma \| P \otimes U) \right\}.$$

By Lemma 6.1,

$$D(\gamma \| P \otimes U) = I(\gamma) + D(\pi_X \gamma \| P),$$

where we have used that  $D(\pi_Y \gamma \| U) = 0$  for any  $\gamma \in \mathcal{M}(U)$ . We obtain

$$\hat{P} = \operatorname{argmin}_{P \in \mathcal{P}} V(P),$$

where

$$V(P) := \min_{\gamma \in \mathcal{M}(U)} \left\{ \frac{1}{2} \mathbb{E}_{(x,y) \sim \gamma} \|x - y\|^2 + \sigma^2 I(\gamma) + \sigma^2 D(\pi_X \gamma \| P) \right\}. \tag{6}$$

Next, for any  $P \in \mathcal{P}$ , denote by  $\gamma_P$  the coupling that achieves the minimum in the definition (1) of  $W_{\sigma^2}(P, U)$  and observe that since  $D(\pi_X \gamma_P \| P) = 0$ , we get

$$V(P) \leq \frac{1}{2} \mathbb{E}_{(x,y) \sim \gamma_P} \|x - y\|^2 + \sigma^2 I(\gamma_P) = W_{\sigma^2}(P, U) < \infty,$$

since  $P$  has finite second moment. We now show that the functions  $P \mapsto V(P)$  and  $P \mapsto W_{\sigma^2}(P, U)$  achieve their minimum over  $\mathcal{P}$  at the same  $P$ . To that end, observe that since  $V(P) < \infty$ , the minimum in the definition (6) of  $V$  may be restricted to couplings  $\gamma$  such that  $\pi_X \gamma \ll P$ , since otherwise  $D(\pi_X \gamma \| P)$  is infinite. Thus,

$$\min_{P \in \mathcal{P}} V(P) = \min_{P \in \mathcal{P}} \min_{\substack{\gamma \in \mathcal{M}(U) \\ \pi_X \gamma \ll P}} \left\{ \frac{1}{2} \mathbb{E}_{(x,y) \sim \gamma} \|x - y\|^2 + \sigma^2 I(\gamma) + \sigma^2 D(\pi_X \gamma \| P) \right\} \tag{7}$$

$$\begin{aligned} &\geq \min_{P \in \mathcal{P}} \min_{\substack{\gamma \in \mathcal{M}(U) \\ \pi_X \gamma \ll P}} \left\{ \frac{1}{2} \mathbb{E}_{(x,y) \sim \gamma} \|x - y\|^2 + \sigma^2 I(\gamma) \right\} \\ &= \min_{P \in \mathcal{P}} \min_{\gamma \in \mathcal{M}(P, U)} \left\{ \frac{1}{2} \mathbb{E}_{(x,y) \sim \gamma} \|x - y\|^2 + \sigma^2 I(\gamma) \right\} \\ &= \min_{P \in \mathcal{P}} W_{\sigma^2}(P, U), \end{aligned} \tag{8}$$

where in the inequality we used that  $D(\pi_X \gamma \| P) \geq 0$ . To justify the equality in (8), we use the fact that  $\mathcal{P}$  is closed under domination, which implies that

$$\begin{aligned} \{\gamma : \exists P \in \mathcal{P} \text{ s.t. } \gamma \in \mathcal{M}(U), \pi_X \gamma \ll P\} &= \{\gamma : \gamma \in \mathcal{M}(U), \pi_X \gamma \in \mathcal{P}\} \\ &= \{\gamma : \exists P \in \mathcal{P} \text{ s.t. } \gamma \in \mathcal{M}(P, U)\}. \end{aligned}$$

Since  $\frac{1}{2} \mathbb{E}_{(x,y) \sim \gamma} \|x - y\|^2 + \sigma^2 I(\gamma)$  depends only on  $\gamma$  and the constraints on  $\gamma$  are the same, the equality holds.

In light of the fact that  $\min_{P \in \mathcal{P}} V(P) = \min_{P \in \mathcal{P}} W_{\sigma^2}(P, U)$ , we conclude that  $\hat{P} = \operatorname{argmin}_{P \in \mathcal{P}} W_{\sigma^2}\left(P, \frac{1}{n} \sum_{i=1}^n \delta_{y_i}\right)$ , as claimed.

**5. Extensions**

5.1. Relaxed transport

Traditional parametric classes of distributions  $\mathcal{P}$  are often not closed under domination. For example, the one-dimensional scale/location family with template density  $\varphi$  with respect to the Lebesgue measure  $\text{Leb}$  on  $\mathbb{R}$  is defined by

$$\mathcal{P} = \left\{ P : \frac{dP}{d\text{Leb}}(\cdot) = \frac{1}{\tau} \varphi\left(\frac{\cdot - \mu}{\tau}\right), \mu \in \mathbb{R}, \tau > 0 \right\}.$$

Clearly  $\mathcal{P}$  is not closed under domination. In such cases, Theorem 4.2 can fail to hold as illustrated by Proposition 5.1 below. However, it follows from (7) in the proof of Theorem 4.2 that the following representation for the MLE *always* holds:

$$\hat{P} = \operatorname{argmin}_{P \in \mathcal{P}} W_{\sigma^2}^{\text{rel}}\left(P, \frac{1}{n} \sum_{i=1}^n \delta_{y_i}\right),$$

where  $W_{\sigma^2}^{\text{rel}}$  denotes the *relaxed entropic optimal transport* distance defined for any probability measures  $\mu, \nu$  by

$$W_{\sigma^2}^{\text{rel}}(\mu, \nu) = \min_{\substack{\gamma \in \mathcal{M}(\nu) \\ \pi_X \gamma \ll \mu}} \int \frac{1}{2} \|x - y\|^2 d\gamma(x, y) + \sigma^2 [I(\gamma) + D(\pi_X \gamma \| \mu)].$$

This result indicates that it may be preferable to use relaxed transport in statistical contexts. Relaxing the marginal constraints in the optimal transport problem is an idea which has attracted significant recent interest [11,19,27] after it was first formally proposed in [4] under the name “unbalanced transport” to generalize optimal transport to apply to nonnegative measures with different total mass. Relaxed optimal transport has since been used to improve robustness to sampling noise in statistical applications [39].

We now exhibit a simple example of a class  $\mathcal{P}$  that is not closed under domination and for which Theorem 4.2 fails to hold. For any  $\sigma > 0$ , let  $\mathcal{P} = \{P_1, P_2\}$  where  $P_1$  and  $P_2$  are two probability measures on the real line defined respectively by

$$P_1 := \frac{1}{2}(\delta_0 + \delta_{4\sigma}) \quad \text{and} \quad P_2 := \frac{1}{2}(\delta_{2\sigma} + \delta_{6\sigma}).$$

Let  $X \sim P_1$  and let  $Y = X + Z$  where  $Z \sim \mathcal{N}(0, \sigma^2)$  is independent of  $X$ .

**Proposition 5.1.** *With probability at least .15, we have*

$$P_1 = \operatorname{argmin}_{P \in \mathcal{P}} W_{\sigma^2}(P, \delta_Y), \quad \text{and} \quad P_2 = \operatorname{argmax}_{P \in \mathcal{P}} \log \varphi_{\sigma} \star dP(Y).$$

*In other words, the maximum-likelihood estimator and the projection with respect to the entropic optimal transport distance do not agree.*

**Proof.** By rescaling, it suffices to consider the case  $\sigma^2 = 1$ . For each  $P \in \mathcal{P}$ , the set  $\mathcal{M}(P, \delta_Y)$  contains only the independent coupling  $P \otimes \delta_Y$ , for which the mutual information vanishes. Therefore,

$$W_{\sigma^2}(P_1, \delta_Y) = W_0(P_1, \delta_Y) = \frac{1}{4}(Y^2 + (Y - 4)^2) \quad W_{\sigma^2}(P_2, \delta_Y) = \frac{1}{4}((Y - 2)^2 + (Y - 6)^2),$$

so that  $P_1$  is the unique minimizer over  $\mathcal{P}$  of  $W_{\sigma^2}(P, \delta_Y)$  as long as  $Y < 3$ .

On the other hand,  $P_2$  is the unique maximum-likelihood estimator if

$$e^{-Y^2/2} + e^{-(Y-4)^2/2} < e^{-(Y-2)^2/2} + e^{-(Y-6)^2/2},$$

and it can be checked that this condition holds on the interval  $[1.01, 3)$ .

Therefore, if  $Y \in [1.01, 3)$ , then the claimed situation occurs. To conclude the proof, it suffices to observe that  $P_1(1.01 \leq Y < 3) \geq .5\mathbb{P}(1.01 \leq |Z| \leq 2.99) \geq .15$ .

### 5.2. General noise distribution

In this section, we raise the question of non-Gaussian deconvolution that arises when the errors  $\{Z_i\}$  are not Gaussian. It turns out that a simple modification of our argument can be made to accommodate any noise distribution that admits a density  $f$  with respect to the Lebesgue measure on  $\mathbb{R}^d$ . In this context, the MLE takes the form

$$\hat{P} = \operatorname{argmax}_{P \in \mathcal{P}} \sum_{i=1}^n \log f \star dP(Y_i). \tag{9}$$

The use of the squared Euclidean norm  $\frac{1}{2\sigma^2} \|x - y\|^2$  in the objective (1) is tailored to Gaussian errors and may be replaced with  $-\log f(x - y)$ . After rescaling by  $\sigma^2$ , we define

$$W_f(\mu, \nu) := \min_{\gamma \in \mathcal{M}(\mu, \nu)} - \int \log f(x - y) d\gamma(x, y) + I(\gamma).$$

We assume in what follows that  $\log f(\cdot - y) \in L_1(P)$  for all  $y \in \mathbb{R}^d$  and  $P \in \mathcal{P}$ . The following proposition is stated without proof as it follows from exactly the same arguments as Theorem 4.2.

**Proposition 5.2.** *Let  $\mathcal{P}$  be a class of probability measures that is closed under domination and assume the deconvolution model (3) where  $Z_i$  has density  $f$  with respect to the Lebesgue measure. Then the maximum-likelihood estimator  $\hat{P}$  over  $\mathcal{P}$  defined in (9) satisfies:*

$$\hat{P} = \operatorname{argmin}_{P \in \mathcal{P}} W_f\left(P, \frac{1}{n} \sum_{i=1}^n \delta_{y_i}\right).$$

In other words, the maximum-likelihood estimator  $\hat{P}$  is the projection of the empirical measure  $\frac{1}{n} \sum_{i=1}^n \delta_{y_i}$  onto  $\mathcal{P}$  with respect to the entropic optimal transport distance  $W_f$ .

In the case where  $f(z) \propto \exp(-\|z\|_p^p)$ , the cost  $-\log f(x - y)$  corresponds to the  $\ell_p$  metric arising in the definition of the  $p$ -Wasserstein distance. Another intriguing example is the cost  $-\log \cos^2(\|x - y\| \wedge \pi/2)$ , which appears in the definition of the Wasserstein–Fisher–Rao [11] or Hellinger–Kantorovich [27] distance between positive measures. These formulations differ from ours in that they consider a version of relaxed transport which allows for misspecification of both marginals; nevertheless, our analysis suggests that inference involving these distances is likely to be robust to convolutional noise  $Z$  supported on the Euclidean ball of radius  $\pi/2$  around the origin with density

$$f(z) \propto \cos^2(\|z\|) \mathbb{1}(\|z\| \leq \frac{\pi}{2}).$$

The Wasserstein–Fisher–Rao/Hellinger–Kantorovich distance is motivated by a dynamic formulation of unbalanced transport and it is unclear whether the above noise distribution plays a special role in the context of deconvolution.

### 6. Additional lemmas

**Lemma 6.1.** *Let  $\gamma$  be a measure on  $\mathbb{R}^d \times \mathbb{R}^d$ , and let  $\alpha$  and  $\beta$  be probability measures on  $\mathbb{R}^d$ . Then*

$$D(\gamma \parallel \alpha \otimes \beta) = I(\gamma) + D(\pi_X \gamma \parallel \alpha) + D(\pi_Y \gamma \parallel \beta).$$

**Proof.** We assume  $\gamma \ll \pi_X \gamma \otimes \pi_Y \gamma \ll \alpha \otimes \beta$ , since otherwise both sides are infinite. Under this condition, we have

$$\begin{aligned} D(\gamma \parallel \alpha \otimes \beta) &= \int \log \frac{d\gamma}{d\alpha d\beta}(x, y) d\gamma(x, y) \\ &= \int \log \frac{d\gamma}{d\pi_X \gamma d\pi_Y \gamma}(x, y) d\gamma(x, y) + \int \log \frac{d\pi_X \gamma d\pi_Y \gamma}{d\alpha d\beta}(x, y) d\gamma(x, y) \\ &= I(\gamma) + \int \log \frac{d\pi_X \gamma}{d\alpha}(x, y) d\gamma(x, y) + \int \log \frac{d\pi_Y \gamma}{d\beta}(x, y) d\gamma(x, y) \\ &= I(\gamma) + \int \log \frac{d\pi_X \gamma}{d\alpha}(x) d\pi_X \gamma(x) + \int \log \frac{d\pi_Y \gamma}{d\beta}(y) d\pi_Y \gamma(y) \\ &= I(\gamma) + D(\pi_X \gamma \parallel \alpha) + D(\pi_Y \gamma \parallel \beta). \end{aligned}$$

## Acknowledgements

This work was supported by the National Science Foundation [DGE-1122374 to J.W.; DMS-1712596, TRIPODS-1740751 to P.R.], the Office of Naval Research [N00014-17-1-2147 to P.R.], the Chan Zuckerberg Initiative DAF, a donor advised fund of the Silicon Valley Community Foundation [2018-182642 to P.R.]; and the MIT Skoltech Program [MIT Skoltech Seed Fund to P.R.].

## References

- [1] J. Altschuler, J. Weed, P. Rigollet, Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration, in: I. Guyon, U. von Luxburg, S. Bengio, H.M. Wallach, R. Fergus, S.V.N. Vishwanathan, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems, 2017, 4–9 December 2017, Long Beach, CA, USA, 2017*, pp. 1961–1971.
- [2] M. Arjovsky, S. Chintala, L. Bottou, Wasserstein GAN, preprint, arXiv:1701.07875, 2017.
- [3] F. Bassetti, A. Bodini, E. Regazzini, On minimum Kantorovich distance estimators, *Stat. Probab. Lett.* 76 (12) (2006) 1298–1302, MR 2269358.
- [4] J.-D. Benamou, Numerical resolution of an “unbalanced” mass transport problem, *M2AN, Math. Model. Numer. Anal.* 37 (5) (2003) 851–868, MR 2020867.
- [5] P.J. Bickel, K.A. Doksum, *Mathematical Statistics: Basic Ideas and Selected Topics*, vol. 1, 2 ed., Prentice-Hall, 2006, updated printing.
- [6] N. Bonneel, M. van de Panne, S. Paris, W. Heidrich, Displacement interpolation using lagrangian mass transport, *ACM Trans. Graph.* 30 (6) (2011) 158:1–158:12.
- [7] C. Caillerie, F. Chazal, J. Dedecker, B. Michel, Deconvolution for the Wasserstein metric and geometric inference, *Electron. J. Stat.* 5 (2011) 1394–1423, MR 2851684.
- [8] R.J. Carroll, P. Hall, Optimal rates of convergence for deconvolving a density, *J. Amer. Stat. Assoc.* 83 (404) (1988) 1184–1186.
- [9] R.J. Carroll, D. Ruppert, L.A. Stefanski, C.M. Crainiceanu, Measurement error in nonlinear models, in: *A Modern Perspective*, second ed., in: *Monogr. Stat. Appl. Probab.*, vol. 105, Chapman & Hall/CRC, Boca Raton, FL, 2006.
- [10] O. Catoni, Statistical learning theory and stochastic optimization, in: *Lecture Notes from the 31st Summer School on Probability Theory held in Saint-Flour, July 8–25, 2001*, in: *Lect. Notes Math.*, vol. 1851, Springer-Verlag, Berlin, 2004, MR 2163920 (2006d:62004).
- [11] L. Chizat, G. Peyré, B. Schmitzer, F.-X. Vialard, Scaling algorithms for unbalanced transport problems, preprint, arXiv:1607.05816, 2016.
- [12] N. Courty, R. Flamary, D. Tuia, A. Rakotomamonjy, Optimal transport for domain adaptation, *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (9) (2017) 1853–1865.
- [13] M. Cuturi, Sinkhorn distances: lightspeed computation of optimal transport, in: C.J.C. Burges, L. Bottou, Z. Ghahramani, K.Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a Meeting, held December 5–8, 2013, Lake Tahoe, Nevada, United States, 2013*, pp. 2292–2300.
- [14] A. Dalalyan, A.B. Tsybakov, Aggregation by exponential weighting, sharp pac-bayesian bounds and sparsity, *Mach. Learn.* 72 (1) (2008) 39–61.
- [15] A.S. Dalalyan, A.B. Tsybakov, Mirror averaging with sparsity priors, *Bernoulli* 18 (3) (2012) 914–944, MR 2948907.
- [16] A.S. Dalalyan, A.B. Tsybakov, Sparse regression learning by aggregation and Langevin Monte-Carlo, *J. Comput. Syst. Sci.* 78 (5) (2012) 1423–1443, arXiv:0903.1223.
- [17] J. Fan, On the estimation of quadratic functionals, *Ann. Stat.* 19 (3) (1991) 1273–1294, MR 1126325 (92j:62006).
- [18] A. Forrow, J.-C. Hütter, M. Nitzan, G. Schiebinger, P. Rigollet, J. Weed, Statistical optimal transport via factored couplings, arXiv:1806.07348, 2018.
- [19] C. Frogner, C. Zhang, H. Mobahi, M. Araya-Polo, T.A. Poggio, Learning with a Wasserstein loss, in: C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7–12, 2015, Montreal, Quebec, Canada, 2015*, pp. 2053–2061.
- [20] A. Genevay, M. Cuturi, G. Peyré, F.R. Bach, Stochastic optimization for large-scale optimal transport, in: D.D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5–10, 2016, Barcelona, Spain, 2016*, pp. 3432–3440.
- [21] A. Genevay, G. Peyré, M. Cuturi, Learning generative models with sinkhorn divergences, in: A.J. Storkey, F. Pérez-Cruz (Eds.), *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9–11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, in: *Proceedings of Machine Learning Research*, vol. 84, PMLR, 2018, pp. 1608–1617.
- [22] E. Giné, R. Nickl, *Mathematical Foundations of Infinite-Dimensional Statistical Models*, Cambridge Ser. Statist. Probab. Math., vol. 40, Cambridge University Press, New York, 2016, MR 3588285.
- [23] W. Jitkrittum, Z. Szabó, K.P. Chwialkowski, A. Gretton, Interpretable distribution features with maximum testing power, in: D.D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems, 2016, December 5–10, 2016, Barcelona, Spain, 2016*, pp. 181–189.
- [24] A. Juditsky, P. Rigollet, A. Tsybakov, Learning by mirror averaging, *Ann. Stat.* 36 (5) (2008) 2183–2206, MR MR2458184.
- [25] M.J. Kearns, Y. Mansour, A.Y. Ng, An information-theoretic analysis of hard and soft assignment methods for clustering, in: D. Geiger, P.P. Shenoy (Eds.), *UAI '97: Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, Brown University, Providence, Rhode Island, USA, August 1–3, 1997, Morgan Kaufmann, 1997, pp. 282–293.

- [26] C. Léonard, A survey of the Schrödinger problem and some of its connections with optimal transport, *Discrete Contin. Dyn. Syst.* 34 (4) (2014) 1533–1574, MR 3121631.
- [27] M. Liero, A. Mielke, G. Savaré, Optimal entropy-transport problems and a new Hellinger–Kantorovich distance between positive measures, *Invent. Math.* 211 (3) (2018) 969–1117, MR 3763404.
- [28] B.G. Lindsay, Mixture models: theory, geometry and applications, in: *Regional Conference Series in Probability and Statistics*, vol. 5, Institute of Mathematical Statistics and American Statistical Association, Haywood CA and Alexandria VA, 1995.
- [29] G. Montavon, K.-R. Müller, M. Cuturi, Wasserstein training of restricted boltzmann machines, in: D.D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016*, December 5–10, 2016, Barcelona, Spain, 2016, pp. 3711–3719.
- [30] J. Mueller, T.S. Jaakkola, Principal differences analysis: interpretable characterization of differences between distributions, in: C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015*, December 7–12, 2015, Montreal, Quebec, Canada, 2015, pp. 1702–1710.
- [31] G. Peyré, M. Cuturi, Computational Optimal Transport, Tech. report, 2017.
- [32] P. Rigollet, Kullback–Leibler aggregation and misspecified generalized linear models, *Ann. Stat.* 40 (2) (2012) 639–665, MR 2933661.
- [33] P. Rigollet, A. Tsybakov, Exponential screening and optimal rates of sparse estimation, *Ann. Stat.* 39 (2) (2011) 731–771, MR 2816337.
- [34] P. Rigollet, A. Tsybakov, Sparse estimation by exponential weighting, *Stat. Sci.* 27 (4) (2012) 558–575.
- [35] P. Rigollet, J. Weed, Uncoupled isotonic regression via minimum Wasserstein deconvolution, arXiv:1806.10648, 2018.
- [36] A. Rolet, M. Cuturi, G. Peyré, Fast dictionary learning with a smoothed wasserstein loss, in: A. Gretton, C.C. Robert (Eds.), *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016*, Cadiz, Spain, May 9–11, 2016, in: *J. Mach. Learn. Res. Workshop Conf. Proc.*, vol. 51, JMLR.org, 2016, pp. 630–638.
- [37] Y. Rubner, C. Tomasi, L.J. Guibas, The earth mover’s distance as a metric for image retrieval, *Int. J. Comput. Vis.* 40 (2) (2000) 99–121.
- [38] F. Santambrogio, *Optimal Transport for Applied Mathematicians*, Birkhäuser, NY, 2015.
- [39] G. Schiebinger, J. Shu, M. Tabaka, B. Cleary, V. Subramanian, A. Solomon, S. Liu, S. Lin, P. Berube, L. Lee, et al., Reconstruction of developmental landscapes by optimal-transport analysis of single-cell gene expression sheds light on cellular reprogramming, *bioRxiv* (2017) 191056.
- [40] E. Schrödinger, Sur la théorie relativiste de l’électron et l’interprétation de la mécanique quantique, *Ann. Inst. Henri Poincaré* 2 (4) (1932) 269–310, MR 1508000.
- [41] J. Solomon, F. De Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, L. Guibas, Convolutional wasserstein distances: efficient optimal transportation on geometric domains, *ACM Trans. Graph.* 34 (4) (2015) 66.
- [42] A.G. Wilson, The use of entropy maximising models, in the theory of trip distribution, mode split and route split, *J. Transp. Econ. Policy* 3 (1) (1969) 108–126.