
Debiasser Beware: Pitfalls of Centering Regularized Transport Maps

Aram-Alexandre Pooladian¹ Marco Cuturi² Jonathan Niles-Weed^{1,3}

Abstract

Estimating optimal transport (OT) maps (*a.k.a.* Monge maps) between two measures P and Q is a problem fraught with computational and statistical challenges. A promising approach lies in using the dual potential functions obtained when solving an entropy-regularized OT problem between samples P_n and Q_n , which can be used to recover an approximately optimal map. The negativity penalization in that scheme introduces, however, an estimation bias that grows with the regularization strength. A well-known remedy to debias such estimates, which has gained wide popularity among practitioners of regularized OT, is to center them, by subtracting auxiliary problems involving P_n and itself, as well as Q_n and itself. We do prove that, under favorable conditions on P and Q , debiasing can yield better approximations to the Monge map. However, and perhaps surprisingly, we present a few cases in which debiasing is provably detrimental in a statistical sense, notably when the regularization strength is large or the number of samples is small. These claims are validated experimentally on synthetic and real datasets, and should reopen the debate on whether debiasing is needed when using entropic OT.

1. Introduction

Estimating an optimal transport (OT) map from a source measure P to a target measure Q is an increasingly central issue in machine learning: such a map, when successfully learned from data, would in principle allow the generation of new samples from Q by pushing particles sampled from P . For example, if we are given full access to the source measure (typically a Gaussian) but only samples from the

target, this problem can be viewed with the lens of normalizing flows (Finlay et al., 2020; Grathwohl et al., 2018; Huang et al., 2021) with tight connections with the problem of estimating generative models (Arjovsky et al., 2017; Genevay et al., 2018a; Salimans et al., 2018). When both P and Q can only be accessed through i.i.d samples $X_1, \dots, X_n \sim P$ and $Y_1, \dots, Y_n \sim Q$, estimating such maps is even more challenging, yet increasingly relevant when, for instance, trying to infer cellular evolution from population measurements (Demetci et al., 2021; Moriel et al., 2021; Schiebinger et al., 2019; Yang et al., 2020).

Formally, we call $T_0 : \mathbb{R}^d \rightarrow \mathbb{R}^d$ an OT map, or a Monge map, when it is the minimizer to the Monge problem (1781):

$$\min_{T \in \mathcal{T}(P, Q)} \int \frac{1}{2} \|x - T(x)\|_2^2 dP(x), \quad (1)$$

where P and Q are two probability measures on $\Omega \subseteq \mathbb{R}^d$, and $\mathcal{T}(P, Q)$ is the set of *pushforward* maps from P to Q , namely $\mathcal{T}(P, Q) := \{T : \Omega \rightarrow \Omega \mid T_{\#}P := P \circ T^{-1} = Q\}$. The existence of such a minimizer is guaranteed under mild regularity conditions on the two measures (Santambrogio, 2015). The task of map estimation is to provide an estimator \hat{T}_n on the basis of samples P_n from P and Q_n from Q whose expected $L^2(P)$ distance to T_0 is small.

When $T_0 \in \mathcal{C}^\alpha$ (the space of $\lfloor \alpha \rfloor$ -times differentiable functions whose $\lfloor \alpha \rfloor$ th derivative is $\alpha - \lfloor \alpha \rfloor$ Hölder smooth) and P and Q satisfy additional technical assumptions, Hütter & Rigollet (2021) provided the first estimator that achieved the following estimation rate,

$$\mathbb{E} \|\hat{T}_n - T_0\|_{L^2(P)}^2 \lesssim n^{-\frac{2\alpha}{2\alpha-2+d}} \log^3(n), \quad (2)$$

which they showed to be minimax optimal up to logarithmic factors. Though statistically optimal, their estimator requires gridding the space of observations, with computational complexity scaling exponentially in d . This work left open the question of finding a computationally tractable estimator that has good statistical properties. While Manole et al. (2021) and Deb et al. (2021) proposed minimax optimal estimators, the fastest estimator between these two works has a runtime of $\tilde{O}(n^3)$. Similarly, Muzellec et al. (2021) devised a nearly minimax map estimator in the high smoothness regime based on kernel sum-of-squares, however the underlying constant depends exponentially in d as their regularization parameter goes to zero.

¹Center for Data Science, New York University, New York, USA ²Google Research, currently at Apple ³Courant Institute of Mathematical Sciences, New York University, New York, USA. Correspondence to: Aram-Alexandre Pooladian <aram-alexandre.pooladian@nyu.edu>.

Leveraging the computational benefits of the Sinkhorn algorithm, Pooladian & Niles-Weed (2021) analyzed the computational and statistical properties of the *Entropic map*, which to any input x associates the conditional expectation of Y given $X = x$ under the entropy-regularized optimal coupling. This estimator can be computed in $\tilde{O}(n^2)$ time and is near optimal when the regularity of T_0^{-1} is low. Together, these results indicate that entropic regularization is a promising methodology for the estimation of Monge maps.

However, entropic regularization is known to introduce some form of bias when solving OT problems, notably when the regularization parameter is large (Cuturi & Peyré, 2018; Schmitzer, 2019). A popular remedy for this phenomenon has been to *debias* said estimates by introducing auxiliary OT problems (Chizat et al., 2020; Feydy et al., 2019; Genevay et al., 2018b). This debiasing procedure has been shown to offer both theoretical and practical benefits for estimating transport distances, suggesting that debiasing should be applied systematically when considering regularized OT problems, notably with the entropic estimation of Monge maps. Surprisingly, we show that the situation is more nuanced: we prove that debiasing can offer benefits when the regularization parameter is small and the number of samples is large but can substantially *degrade* statistical performance when these conditions are not met.

Main Contributions: In this paper, we explore whether debiasing is a fruitful approach for estimating Monge maps. Specifically, our contributions are the following:

1. We define the debiased map estimator, and prove a rate of convergence to the Monge map in the small regularization limit (§ 3, Theorem 2);
2. We prove quantitative differences between the Entropic map and its debiased variant in the large regularization regime, and present examples where debiasing provably degrades performance (§ 3, Theorem 3 and Theorem 4);
3. We derive closed-form expressions in the Gaussian-to-Gaussian case, and show that debiasing yields a better approximation to the Monge map when the regularization parameter is small (§ 4);
4. We illustrate advantages and pitfalls of debiasing on several numerical experiments (§ 5). Notably, we show that finite-sample effects can substantially worsen the performance of the debiased map.

We use the symbol C to denote a positive constant whose value may change from line to line, and write $a \lesssim b$ if there exists a constant $C > 0$ such that $a \leq Cb$.

2. Background on Optimal transport

2.1. Optimal transport without regularization

For $\Omega \subseteq \mathbb{R}^d$ compact, we define $\mathcal{P}(\Omega)$ to be the space of (Borel) probability measures with support contained in Ω , and $\mathcal{P}_{ac}(\Omega)$ be those with densities. We first recall Brenier’s fundamental result on the existence of an optimal map.

Theorem 1 (Brenier’s Theorem (1991)). *Let $P \in \mathcal{P}_{ac}(\Omega)$ and $Q \in \mathcal{P}(\Omega)$. Then there exists a solution T_0 to Equation (1), with $T_0 = Id - \nabla f_0$, where f_0 is a 1-strongly concave function solving*

$$\sup_{(f,g) \in \Phi} \int f \, dP + \int g \, dQ, \quad (3)$$

where $\Phi = \{(f, g) \in L^1(P) \times L^1(Q) : f(x) + g(y) \leq \frac{1}{2}\|x - y\|_2^2 (x, y \in \Omega)\}$.

In other words, Brenier’s theorem asserts that when P has a density function, the optimal map is the gradient of a convex function, $x \mapsto \frac{1}{2}\|x\|_2^2 - f_0(x)$. We call the maximizers (f_0, g_0) to Equation (3) *optimal (Kantorovich) potentials*.

If P does not have a density, an optimal transport map may not exist between P and Q . To remedy this, Kantorovich (1942) devised a convex relaxation of the optimal transport problem where one is optimizing over *plans* instead of maps

$$\text{OT}_0(P, Q) := \inf_{\pi \in \Pi(P, Q)} \iint \frac{1}{2}\|x - y\|_2^2 \, d\pi(x, y), \quad (4)$$

where $\Pi(P, Q)$ is the set of joint probability measures with marginals P and Q , called the set of *couplings*. When an optimal map exists, it also gives rise to a solution to Eq. (4); however, unlike the Monge problem, Eq. (4) always admits a minimizer when P and Q have finite second moments. We call such a minimizer an *optimal plan*, denoted π_0 . Note that OT_0 is also the definition of $\frac{1}{2}W_2^2(P, Q)$, the squared 2-Wasserstein distance between P and Q .

2.2. Entropic optimal transport

We define *Entropic optimal transport* to be the objective function that arises when we add the *Kullback–Liebler (KL) divergence*, denoted $D_{\text{KL}}(\cdot \| \cdot)$, as a regularizer to Eq. (4). Recall that $D_{\text{KL}}(\mu \| \nu) = \int \log(\frac{d\mu}{d\nu}) \, d\mu$. Thus, for $P, Q \in \mathcal{P}_{ac}(\Omega)$, we write the ε -regularized OT problem as

$$\begin{aligned} \text{OT}_\varepsilon(P, Q) &= \inf_{\pi \in \Gamma(P, Q)} \iint c \, d\pi + \varepsilon D_{\text{KL}}(\pi \| P \otimes Q) \quad (5) \\ &= \inf_{\pi \in \Gamma(P, Q)} \iint c \, d\pi + \varepsilon \iint \log \pi \, d\pi \\ &\quad - \varepsilon (\text{Ent}(P) + \text{Ent}(Q)), \end{aligned}$$

where we abbreviate $c(x, y) = \frac{1}{2}\|x - y\|_2^2$, and the *entropy* of an absolutely continuous probability measure is denoted

by $\text{Ent}(P) = \int p(x) \log(p(x)) dx$; note that the first line also holds in the absence of densities. Since OT_ε is now strongly convex, we call its unique minimizer the *optimal entropic plan*, denoted π_ε , which exists under a finite-second moment condition (Peyré & Cuturi, 2019).

Similar in spirit to Eq. (3), the Entropic OT problem admits the following dual formulation under mild conditions,

$$\begin{aligned} \text{OT}_\varepsilon(P, Q) = & \sup_{f \in L^1(P), g \in L^1(Q)} \int f dP + \int g dQ \\ & - \varepsilon \iint e^{(f(x)+g(y)-c(x,y))/\varepsilon} dP(x) dQ(y) + \varepsilon, \end{aligned} \quad (6)$$

where the maximizers, denoted $(f_\varepsilon, g_\varepsilon)$ are called *optimal entropic potentials*. Primal-dual relationships provide that the optimal entropic coupling is closely related to $(f_\varepsilon, g_\varepsilon)$ through the following representation (Csiszár, 1975):

$$d\pi_\varepsilon(x, y) = \exp\left(\frac{f_\varepsilon(x) + g_\varepsilon(y) - c(x, y)}{\varepsilon}\right) dP(x) dQ(y).$$

As a result, and since π_ε integrates to 1, it holds for the maximizers that Equation (6) simplifies to the expression (see also other facts in Appendix A):

$$\text{OT}_\varepsilon(P, Q) = \int f_\varepsilon dP + \int g_\varepsilon dQ.$$

2.3. The Sinkhorn divergence

While $\text{OT}_\varepsilon(P, Q)$ is a reasonable approximation of $W_2^2(P, Q)$, we have that $\text{OT}_\varepsilon(P, P) \neq 0$ but $W_2^2(P, P) = 0$, which leads to an inherent bias. In view of this phenomenon, several works have suggested *centering* the Entropic OT objective in the following manner

$$S_\varepsilon(P, Q) := \text{OT}_\varepsilon(P, Q) - \frac{1}{2}\text{OT}_\varepsilon(P, P) - \frac{1}{2}\text{OT}_\varepsilon(Q, Q), \quad (7)$$

resulting in the so-called *Sinkhorn divergence*. Centering the Entropic OT objective is akin to *debiasing* the objective, as now $S_\varepsilon(P, Q) = 0 \iff P = Q$ (Feydy et al., 2019).

The corrective terms, $\text{OT}_\varepsilon(P, P)$ and $\text{OT}_\varepsilon(Q, Q)$, each admit their own dual formulation that is similar to Equation (6), with the key difference that they each possess only one optimal potential (Feydy et al., 2019). For example, we denote by α_ε the optimal entropic *self-potential*:

$$\begin{aligned} \text{OT}_\varepsilon(P, P) = & \sup_{\alpha \in L^1(P)} \int 2\alpha dP \\ & - \varepsilon \iint \left(e^{(\alpha(x)+\alpha(y)-c(x,y))/\varepsilon} - 1 \right) dP(x) dP(y) \\ & = 2 \int \alpha_\varepsilon dP, \end{aligned} \quad (8)$$

and similarly by β_ε the optimal self-potential corresponding to $\text{OT}_\varepsilon(Q, Q)$. Thus, at optimality, one has the relationship

$$S_\varepsilon(P, Q) = \int (f_\varepsilon - \alpha_\varepsilon) dP + \int (g_\varepsilon - \beta_\varepsilon) dQ \quad (9)$$

$$=: \int f_\varepsilon^D dP + \int g_\varepsilon^D dQ, \quad (10)$$

where we define $(f_\varepsilon^D, g_\varepsilon^D)$ to be the *debiased optimal entropic potentials*, or *Sinkhorn potentials*.

The Sinkhorn divergence has found many applications in the machine learning community as a way of eliminating the “bias” in the entropic OT objective (Chizat et al., 2020; Feydy et al., 2019; Genevay et al., 2018b). For example, in the case of estimating the Wasserstein distance, this debiasing effect can be made rigorous: Pal (2019) showed that OT_ε gives an additive $O(\varepsilon \log(\varepsilon^{-1}))$ approximation to W_2^2 . By contrast, Chizat et al. (2020) showed that $S_\varepsilon = W_2^2 + O(\varepsilon^2)$.

2.4. The Entropic map

Pooladian & Niles-Weed (2021) studied the *Entropic map*, defined to be the *barycentric projection* of the optimal entropic plan, defined as

$$T_\varepsilon(x) := \mathbb{E}_{\pi_\varepsilon}[Y|X = x], \quad (11)$$

as an estimator for T_0 between two measures P and Q . This object is partially motivated by the following observation, which acts as an entropic analogue to Brenier’s Theorem (see Pooladian & Niles-Weed, 2021, Prop. 2):

$$T_\varepsilon = \text{Id} - \nabla f_\varepsilon. \quad (12)$$

Moreover, in the finite-sample regime, this estimator can be computed in $\tilde{O}(n^2)$ operations and can be efficiently parallelized on GPUs, which is in contrast to other estimators in the literature that require either at least $\tilde{O}(n^3)$ complexity and are not easily parallelizable (Deb et al., 2021; Hütter & Rigollet, 2021; Manole et al., 2021).

3. Entropic map vs. Sinkhorn map

In this work, we introduce the *Sinkhorn map*

$$T_\varepsilon^D(x) := x - \nabla f_\varepsilon^D(x) = T_\varepsilon(x) + \nabla \alpha_\varepsilon(x) \quad (13)$$

which can be viewed as the natural, and perhaps desirable, extension of the Entropic map from § 2.4 following the centering or *debiasing* principle advocated in § 2.3.

3.1. Biased and Debaised MSE

As discussed in § 2.3, debiasing yields crucial improvements in approximating $W_2^2(P, Q)$. Our primary focus is to determine whether or not such an improvement can be

made in the context of map estimation. Assuming we have full access to the source and target measures P and Q , we are interested in comparing the limiting behaviors of their Mean-Squared Error (MSE), $\mathcal{R}(T_\varepsilon)$ and $\mathcal{R}(T_\varepsilon^D)$, where for a given T_0 and any estimator $S : \mathbb{R}^d \rightarrow \mathbb{R}^d$,

$$\mathcal{R}(S) := \|S - T_0\|_{L^2(P)}^2. \quad (14)$$

Small ε limit As a by-product of their analysis, Pooladian & Niles-Weed (2021) proved the following estimate:

Proposition 1. *Let $P, Q \in \mathcal{P}_{ac}(\Omega)$ with bounded densities with compact support, with the density of Q bounded from below. Let $T_0 := \nabla\varphi_0$ be the optimal map between P and Q , where $\mu I \preceq \nabla^2\varphi_0 \preceq LI$, and $T_0^{-1} \in C^\alpha$ for $\alpha > 1$. Then for ε sufficiently small,*

$$\mathcal{R}(T_\varepsilon) \lesssim \varepsilon^2 I_0(P, Q) + \varepsilon^{(\bar{\alpha}+1)/2},$$

where $\bar{\alpha} = 3 \wedge \alpha$ and $I_0(P, Q)$ is the integrated Fisher information along the $W_2^2(P, Q)$ geodesic.

We notice the following phenomena: in the high smoothness regime, we obtain a quadratic rate of convergence. On the other hand, as $\alpha \rightarrow 1$, the convergence rate is linear in the regularization parameter.

Our first contribution, which is potentially of independent interest, is a gradient estimate of the self-potential α_ε : we show that the effect of adding $\nabla\alpha_\varepsilon$ to the Entropic map as defined in Eq. (13) is small. Proofs of the results found in this section are deferred to Appendix B.

Lemma 1. *Let α_ε be the solution to $OT_\varepsilon(P, P)$ with P having a density function with compact support. Then*

$$\|\nabla\alpha_\varepsilon\|_{L^2(P)}^2 \leq \frac{\varepsilon^2}{4} I_0(P, P), \quad (15)$$

where $I_0(P, P) := \int \|\nabla \log p(x)\|^2 p(x) dx$ is the Fisher information of P .

Combining Proposition 1 with Lemma 1, we obtain the following result, which shows that in the $\varepsilon \rightarrow 0$ regime, we are no worse when we debias our map estimator.

Theorem 2. *Under the assumptions of Proposition 1, for ε small enough, there exists $C_1, C_2 > 0$ such that*

$$\mathcal{R}(T_\varepsilon^D) \leq \mathcal{R}(T_\varepsilon) + C_1\varepsilon^2 + C_2\varepsilon^{(\bar{\alpha}+5)/4},$$

where C_1 and C_2 are constants depending on $I_0(P, P)$ and $I_0(P, Q)$. In particular, if $I_0(P, P)$ and $I_0(P, Q)$ are finite, then as $\varepsilon \rightarrow 0$,

$$T_\varepsilon \rightarrow T_0, \quad T_\varepsilon^D \rightarrow T_0 \quad \text{in } L^2(P).$$

Large ε limit We now turn our attention to the large ε regime. Proposition 2 characterizes the limiting behavior of the maps, whereas Theorem 3 concerns their MSE.

Proposition 2. *Let $P, Q \in \mathcal{P}(\Omega)$, and let $\mu_P, \mu_Q \in \mathbb{R}^d$ denote their means. As $\varepsilon \rightarrow \infty$, and in $L^2(P)$ sense, one has*

$$T_\varepsilon \rightarrow \mu_Q, \quad T_\varepsilon^D \rightarrow Id + (\mu_Q - \mu_P). \quad (16)$$




Figure 1. Visual behavior of the Entropic T_ε and Sinkhorn (debiased) T_ε^D maps between two Gaussians using closed-form expressions derived in § 4, evaluated on arbitrary samples from the input Gaussian. The estimators approach T_0 as $\varepsilon \approx 0$. When ε is relatively large, we see that T_ε is biased towards the mean of the target, which in this case is zero. In contrast, T_ε^D is still maintaining the overall shape of the target.

Proposition 2 shows that as $\varepsilon \rightarrow \infty$, the Entropic map approaches the constant map $x \mapsto \mu_Q$, whereas the Sinkhorn map approaches the linear map $x \mapsto x + (\mu_Q - \mu_P)$. This contrasts with the $\varepsilon \rightarrow 0$ limit, where both maps tend to T_0 ; these two behaviors are illustrated in Figure 1.

Proposition 2 implies the following expressions for the limiting values of the MSE.

Theorem 3. *For $P, Q \in \mathcal{P}(\Omega)$, as $\varepsilon \rightarrow \infty$,*

$$\mathcal{R}(T_\varepsilon) \rightarrow \text{Var}(Q), \quad \mathcal{R}(T_\varepsilon^D) \rightarrow W_2^2(\bar{P}, \bar{Q}),$$

where $\bar{P} := P - \mu_P$ and $\bar{Q} := Q - \mu_Q$.

Crucially, neither $\text{Var}(Q)$ nor $W_2^2(\bar{P}, \bar{Q})$ dominates the other in general, and whether $\mathcal{R}(T_\varepsilon)$ or $\mathcal{R}(T_\varepsilon^D)$ is larger




Figure 2. Validation of theoretical results from Section 3 and Section 4, when $P = \mathcal{N}(0, I_2)$ and $Q_1 = \mathcal{N}(0, \text{diag}(2, 1))$ (left figure) and $Q_2 = \mathcal{N}(0, \text{diag}(0.1, 0.1))$ (right figure). Making use of closed-form expressions, Monte Carlo integration was used to make the figures.

depends on the properties of P and Q . Theorem 3 exposes one potential pitfall of debiasing: it can either help or harm the estimate of the Monge map, depending on the size of ε and the properties of P and Q .

Figure 2 illustrates this effect in the Gaussian-to-Gaussian setting. When ε is small, debiasing yields better estimates, but when ε is large, the approximation error of the maps converge to the proven quantities in Theorem 3, and debiasing does not necessarily dominate the biased estimator.

Debiasing can be harmful even for small ε Theorem 3 shows that debiasing is not uniformly helpful when ε is large. Nevertheless, one may conjecture that this phenomenon only manifests when $\varepsilon \rightarrow \infty$, and that there exists ε_0 such that debiasing is benign for $\varepsilon \leq \varepsilon_0$. Our next result shows that this conjecture is emphatically false: it is possible for debiasing to yield *arbitrarily worse* estimates of the Monge map, no matter how small ε is taken.

Theorem 4. *For any $\varepsilon < 1$ and any $M > 0$, there exist a pair of densities $(P_\varepsilon, Q_\varepsilon)$ for which $\mathcal{R}(T_\varepsilon^D) \geq MR(T_\varepsilon)$.*

4. Case Study: Transport between Gaussians

We now explicitly analyze the case $P = \mathcal{N}(0, A)$ and $Q = \mathcal{N}(b, B)$ with $A, B \succ 0$ (the case where the mean of the first measure is non-zero can be recovered with a simple translation). The optimal transport map between two such Gaussians has the following closed-form solution (Gelbrich, 1990)

$$\begin{aligned} T_0(x) &= C_0^{AB}x + b \\ &:= A^{-1/2}(A^{1/2}BA^{1/2})^{1/2}A^{-1/2}x + b. \end{aligned}$$

Since the Entropic map is defined as the conditional mean of the optimal entropic plan, this also has a closed form:

Proposition 3. (Janati et al., 2020; Mallasto et al., 2021) *For $P = \mathcal{N}(0, A)$ and $Q = \mathcal{N}(b, B)$ and $\varepsilon > 0$, the Entropic map is given by $T_\varepsilon(x) = C_\varepsilon^{AB}x + b$, where*

$$\begin{aligned} C_\varepsilon^{AB} &:= (A^{-1/2}[A^{1/2}BA^{1/2} + (\varepsilon^2/4)I]^{1/2}A^{-1/2} \\ &\quad - (\varepsilon/2)A^{-1}). \end{aligned} \quad (17)$$

We now characterize the Sinkhorn map, and prove that in the small ε regime, it is a better estimator of T_0 than the Entropic map; see Figure 2. Proofs pertaining to this section are found in Appendix C.

Proposition 4. *For $P = \mathcal{N}(0, A)$, $\alpha_\varepsilon(x) = \frac{1}{2}x^\top(I - C_\varepsilon^{AA})x$, where*

$$C_\varepsilon^{AA} = A^{-1/2}(A^2 + (\varepsilon^2/4)I)^{1/2}A^{-1/2} - (\varepsilon/2)A^{-1}.$$

Corollary 1. *For $P = \mathcal{N}(0, A)$ and $Q = \mathcal{N}(b, B)$, the debiased map estimator is $T_\varepsilon^D(x) := \tilde{C}_\varepsilon^{AB}x + b$, where*

$$\begin{aligned} \tilde{C}_\varepsilon^{AB} &= (A^{-1/2}(A^{1/2}BA^{1/2} + (\varepsilon^2/4)I)^{1/2}A^{-1/2} \\ &\quad + I - A^{-1/2}(A^2 + (\varepsilon^2/4)I)^{1/2}A^{-1/2}). \end{aligned} \quad (18)$$

Proof. This follows from our definition of $T_\varepsilon^D(x) = T_\varepsilon(x) + \nabla\alpha_\varepsilon(x)$, and Propositions 3 and 4, noting that $\tilde{C}_\varepsilon^{AB} = C_\varepsilon^{AB} + I - C_\varepsilon^{AA}$. \square

With these closed-form expressions in hand, we can proceed with the following theorems.

Theorem 5. *For $P = \mathcal{N}(0, A)$ and $Q = \mathcal{N}(b, B)$,*

$$\|T_\varepsilon - T_0\|_{L^2(P)}^2 \leq \frac{\varepsilon^2}{4}I_0(P, P) + O_{A,B}(\varepsilon^4)$$

where $I_0(P, P) = \text{Tr}(A^{-1})$.

A key difference between Theorem 5 and Proposition 1 is that the leading-order error term here scales like $\varepsilon^2 I_0(P, P)$ instead of $\varepsilon^2 I_0(P, Q)$. This is made possible by having a more precise control on T_ε . While at a glance this seems like a minor difference, it allows us to improve the convergence rate for the debiased Entropic map. We provide a short proof sketch that outlines why this true; the full proofs for Theorems 5 and 6 can be found in Appendix C. Essentially, Theorem 6 is possible because the leading-order contributions of T_ε are exactly canceled out by those from $\nabla\alpha_\varepsilon$, resulting in improved estimates for the debiased map.

Theorem 6. For $P = \mathcal{N}(0, A)$ and $Q = \mathcal{N}(b, B)$,

$$\|T_\varepsilon^D - T_0\|_{L^2(P)}^2 \lesssim \varepsilon^4 + O_{A,B}(\varepsilon^6).$$

Thus, for ε small enough, $\mathcal{R}(T_\varepsilon^D) \leq \mathcal{R}(T_\varepsilon)$.

Proof sketch. Using a Taylor expansion of the explicit expressions of the maps for the Gaussian-to-Gaussian case, one can check that

$$\begin{aligned} T_\varepsilon^D(x) &= T_\varepsilon(x) + \nabla\alpha_\varepsilon(x) \\ &= \left(T_0(x) - \frac{\varepsilon}{2} \nabla \log(p(x)) + O(\varepsilon^2) \right) \\ &\quad + \left(0 + \frac{\varepsilon}{2} \nabla \log(p(x)) + O(\varepsilon^2) \right) \\ &= T_0(x) + O(\varepsilon^2). \end{aligned}$$

Thus, one expects $\|T_\varepsilon^D - T_0\|_{L^2(P)}^2 = O(\varepsilon^4)$. \square

A visual representation of these rates appears in Figure 2: for ε small, we observe convergence rates with order $O(\varepsilon^2)$ and $O(\varepsilon^4)$ for the biased and debiased MSE, respectively.

5. Experiments

5.1. Finite-sample map estimation

The results of §§ 3-4 illustrate the benefits and pitfalls of debiasing to obtain good approximations of the Monge map between two *known* distributions P and Q . This next section shows the extent to which these findings extend to the finite-sample regime. Our findings are double-edged: when the number of samples is sufficiently large and the dimension is moderate, debiasing can still have benefits. However, for smaller sample sizes and larger dimensions, the statistical error begins to dominate, and debiasing yields *worse* performance. We illustrate both phenomena on a suite of benchmark examples.

In the finite-sample regime, the Entropic map can be written

$$\hat{T}_\varepsilon(x) = \frac{\sum_{i=1}^n Y_i e^{(\hat{g}_\varepsilon(Y_i) - c(x, Y_i))/\varepsilon}}{\sum_{i=1}^n e^{(\hat{g}_\varepsilon(Y_i) - c(x, Y_i))/\varepsilon}},$$

where we write $(\hat{f}_\varepsilon, \hat{g}_\varepsilon)$ as the optimal entropic potentials for $\text{OT}_\varepsilon(P_n, Q_n)$. Through the celebrated Sinkhorn's algorithm, it is well-known that these potentials can be computed in $\tilde{O}(n^2)$ operations (Altschuler et al., 2017; Peyré & Cuturi, 2019), and evaluating the map is done in linear time. A similar procedure can be performed for the self-potential $\hat{\alpha}_\varepsilon$ though it empirically converges significantly faster (Feydy et al., 2019). Together, this results in the finite-sample Sinkhorn map estimator

$$\begin{aligned} \hat{T}_\varepsilon^D(x) &= \hat{T}_\varepsilon(x) + \nabla\hat{\alpha}_\varepsilon(x) \\ &= \hat{T}_\varepsilon(x) + x - \frac{\sum_{i=1}^n X_i e^{(\hat{\alpha}_\varepsilon(X_i) - c(x, X_i))/\varepsilon}}{\sum_{i=1}^n e^{(\hat{\alpha}_\varepsilon(X_i) - c(x, X_i))/\varepsilon}}. \end{aligned}$$

Our numerical experiments were performed using Google Colab Pro, where our code is adapted from (Chizat et al., 2020) and is publicly available. Across all plots, we compute the MSE via Monte-Carlo integration, where we always use $5 \cdot 10^5$ points. Unless otherwise specified, we perform our experiments across 20 random trials in order to generate error bars. More details are included in Appendix E.

5.1.1. SMOOTH MAPS

Restricting our attention to P admitting a density with compact support, we define our target distribution as $Q := T_\# P$ for T smooth. For a given dimension, we fix our choice of ε , compute the Entropic and Sinkhorn maps as n varies, and then compute their MSE.

A first example is when $P = \text{Unif}([-1, 1]^d)$ with the map

$$\text{(E1)} \quad T(x) = \Omega_d x$$

where we use the diagonal matrix example from (Paty et al., 2020), defined as $(\Omega_d)_{ii} := 0.8 - \frac{0.4}{d-1}(i-1)$ ($i \in [d]$). Figure 3 shows the MSE of the two estimators as a function of the number of samples, where the Sinkhorn map is better by over an order of magnitude. Due to space constraints, the remainder of our synthetic experiments in this setting are deferred to Appendix E. Though, in all cases, the Sinkhorn map always has a lower MSE than the Entropic map.

5.1.2. INVESTIGATING THE IMPACT OF SMOOTHNESS

We briefly investigate the impact of debiasing when the optimal map is itself non-smooth. We show that this lack of smoothness only moderately affects the performance of our estimators. Taking $P = \text{Unif}([-1, 1]^d)$ again, an example of a non-smooth Monge map is

$$\text{(E2)} \quad T(x) = \partial \left(\frac{1}{2} \|x\|^2 + 2|x_1| \right) = x + 2e_1 \text{sign}(x_1).$$

Despite being the subgradient of a 1-strongly convex potential, (E2) is a discontinuous map that perturbs the input



 (a) Example (E1) in $d = 5$ with $\varepsilon = 0.5$

Figure 3. In moderate dimension for relatively large ε , \hat{T}_ε^D better approximates (E1) compared to \hat{T}_ε .

along the first coordinate. By approximating $\text{sign}(\cdot)$ with a smooth function, denoted $\text{sign}_\beta(\cdot)$ ¹, we can compare map estimation of (E2) and its smooth analogue:

$$(E2') \quad T(x) = x + 2e_1 \text{sign}_\beta(x_1).$$

In Figure 4, we see that the MSEs of both estimators are somewhat lower when the optimal map is smooth, with the Sinkhorn map still better.

5.2. Potential pitfalls in finite-sample estimation

We return to the example of estimating the transport map between two Gaussian distributions. Recall the behavior predicted in Theorem 5 and Theorem 6 and verified in Fig. 2: for ε sufficiently small, the debiased map provably gives a better approximation to the Monge map. However, Figure 5 shows that the *statistical* performance of the debiased map estimator can be substantially worse.

We fix the source distribution as $P = \mathcal{N}(0, I_d)$ where d is the dimension, and we randomly generate covariance matrices as outlined in Appendix E.3 to create the target distribution $Q = \mathcal{N}(0, \Sigma)$. We generate Σ such that its eigenvalues are smaller than 1, causing it to be more concentrated than the source. Across 15 trials, we learn the map using N points labeled in the figures. The lines labeled $N = \infty$ denote the error in the infinite-sample limit, obtained using the closed-form expressions given in § 4.

In $d = 2$, we see the performance of \hat{T}_ε^D plateau at small values of ε even when using $N = 10^4$ points to learn the map, whereas \hat{T}_ε does not seem to experience this effect. For $d = 15$, the performance of both estimators degrades, but the effects are worse for the Sinkhorn map, which is worse than the Entropic map for all $\varepsilon \leq 1$. We conjecture that the Sinkhorn map suffers more dramatically from finite-

¹ $\text{sign}_\beta(\cdot) = 2(1 + e^{-\beta x})^{-1} - 1$ with $\beta \gg 1$




 (a) Example (E2) in $d = 5$ with $\varepsilon = 0.5$

 (b) Example (E2') in $d = 5$ with $\varepsilon = 0.5$ and $\beta = 50$

Figure 4. Smoothing (E2) lowers the MSE for both \hat{T}_ε and \hat{T}_ε^D .

sample effects because the term $\nabla \hat{\alpha}_\varepsilon(x)$ injects additional statistical noise into the map estimate.

5.3. Application: predicting trajectories of genomes

We turn our attention to an application of map estimation using real-world data, where practitioners may not have a priori knowledge of a map even existing between the source and target measures. Such an example arises in (Demetci et al., 2021; Moriel et al., 2021; Schiebinger et al., 2019), where the task is to infer cellular evolution from population measurements. The original data is temporal, where cell measurements are taken across 18 days, and each sampled data point consists of over 1000 gene expressions. Following the setup of Schiebinger et al. (2019), we project the genes onto \mathbb{R}^{30} using PCA. Finally, we normalize the data so that each datapoint lies in a ball of unit radius. We denote the source and target distributions by X and Y , corresponding to the sampled data at day0 and day1.

Across a range of ε values, we perform the following experiment across 20 trials where the train/test split is 50/50. For a fixed value of ε , we use $(X_{\text{train}}, Y_{\text{train}})$ to learn the mappings \hat{T}_ε and \hat{T}_ε^D . Since there is no a priori notion of an optimal map, checking the MSE is not possible. Instead, we compute the (discrete) W_2 distance between a predicted map-




 (a) \hat{T}_ε vs. \hat{T}_ε^D with Σ concentrated in $d = 2$

 (b) \hat{T}_ε vs. \hat{T}_ε^D with Σ concentrated in $d = 15$

Figure 5. We observe strong finite-sample effects for the Sinkhorn map when the regularization parameter is small. While curse-of-dimensionality effects are pervasive in OT, Figure 5(a) shows that these effects occur even in $d = 2$ for the Sinkhorn map.

ping, such as $\hat{T}_\varepsilon(\mathbf{X}_{\text{test}})$ (similarly for the Sinkhorn map), and \mathbf{Y}_{test} , newly sampled points from the target distribution. We compute the (unregularized) W_2 distance using the Python OT (POT) package (Flamary et al., 2021).

Figure 6 shows the computed W_2 distance between the predicted mapping and test data. We notice that as ε get smaller, \hat{T}_ε begins to dominate. This is expected given our observations from Figure 5 in the high-dimensional regime. As in Figure 5, the covariance of the target distribution day1 has smaller eigenvalues than that of day0 .

Figure 7 is a visualization of the first two components of the maps and the target distribution, which parallels Figure 1: for large values of ε , the Entropic map is heavily concentrated towards the mean of the target data. On the other hand, as $\varepsilon \rightarrow 0$, the performance of \hat{T}_ε continues to improve, whereas that of \hat{T}_ε^D does not.




Figure 6. Comparing $W_2(\hat{T}(\mathbf{X}_{\text{test}}), \mathbf{Y}_{\text{test}})$ for different levels of ε . As ε gets smaller, we see the same finite-sample effects from Section 5.2, and \hat{T}_ε starts to perform better than \hat{T}_ε^D .




Figure 7. Visualizing genome predictions of the Entropic and Sinkhorn maps for small and large values of ε through the first two principal components. Much like Figure 1, for large ε , the Entropic map is biased towards the mean of the target.

6. Conclusion

In this work, we investigate theoretical and empirical properties of two estimators of Monge maps: the Entropic map and its debiased counterpart, the Sinkhorn map. In the population regime, we characterize their convergence as the regularization strength ε approaches either 0 or ∞ , and illustrate these results experimentally. Our findings indicate that debiasing does not always help, contradicting a domi-

nant belief in the regularized OT literature. For example, in the Gaussian-to-Gaussian transportation problem, we prove (in the population setting) that the Sinkhorn map is a strictly better estimator as $\varepsilon \rightarrow 0$. However, we also show, both theoretically and empirically, that for a *fixed* ε , the Sinkhorn map can perform worse than the Entropic map, and that these effects are magnified with small sample sizes. We confirm these findings experimentally on real datasets. On the other hand, when empirically estimating smooth maps, we notice that for an *appropriately chosen* value of ε , the Sinkhorn map can yield better estimates than the Entropic map. These findings indicate that the benefits of debiasing are not robust to the choice of the regularization parameter, and that it should be applied selectively.

Acknowledgements

AAP and JNW acknowledge the support of a Google Research Collabs grant.

References

- Altschuler, J., Weed, J., and Rigollet, P. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA, 2017*.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. *Proceedings of the 34th International Conference on Machine Learning*, 70:214–223, 2017.
- Brenier, Y. Polar factorization and monotone rearrangement of vector-valued functions. *Comm. Pure Appl. Math.*, 44(4):375–417, 1991. ISSN 0010-3640. doi: 10.1002/cpa.3160440402. URL <https://doi.org/10.1002/cpa.3160440402>.
- Chizat, L., Roussillon, P., Léger, F., Vialard, F.-X., and Peyré, G. Faster Wasserstein distance estimation with the Sinkhorn divergence. *Advances in Neural Information Processing Systems*, 33, 2020.
- Conforti, G. and Tamanini, L. A formula for the time derivative of the entropic cost and applications. *Journal of Functional Analysis*, 280(11):108964, 2021.
- Csiszár, I. I -divergence geometry of probability distributions and minimization problems. *Ann. Probability*, 3: 146–158, 1975.
- Cuturi, M. and Peyré, G. Semidual regularized optimal transport. *SIAM Review*, 60(4):941–965, 2018.
- Deb, N., Ghosal, P., and Sen, B. Rates of estimation of optimal transport maps using plug-in estimators via barycentric projections. *arXiv preprint arXiv:2107.01718*, 2021.
- Demetci, P., Santorella, R., Sandstede, B., and Singh, R. Unsupervised integration of single-cell multi-omics datasets with disparities in cell-type representation. *bioRxiv*, 2021. doi: 10.1101/2021.11.09.467903. URL <https://www.biorxiv.org/content/early/2021/11/11/2021.11.09.467903>.
- Feydy, J., Séjourné, T., Vialard, F.-X., Amari, S.-i., Trounev, A., and Peyré, G. Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2681–2690. PMLR, 2019.
- Finlay, C., Gerolin, A., Oberman, A. M., and Pooladian, A.-A. Learning normalizing flows from entropy-kantorovich potentials. *arXiv preprint arXiv:2006.06033*, 2020.
- Flamary, R., Courty, N., Gramfort, A., Alaya, M. Z., Boissunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N. T., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D. J., Tavenard, R., Tong, A., and Vayer, T. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021. URL <http://jmlr.org/papers/v22/20-451.html>.
- Gelbrich, M. On a formula for the l^2 wasserstein metric between measures on euclidean and hilbert spaces. *Mathematische Nachrichten*, 147(1):185–203, 1990.
- Genevay, A., Peyré, G., and Cuturi, M. Learning generative models with Sinkhorn divergences. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, pp. 1608–1617, 2018a.
- Genevay, A., Peyré, G., and Cuturi, M. Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain, 2018b*.
- Grathwohl, W., Chen, R. T., Bettencourt, J., Sutskever, I., and Duvenaud, D. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.
- Huang, C.-W., Chen, R. T. Q., Tsirigotis, C., and Courville, A. Convex Potential Flows: Universal Probability Distributions with Optimal Transport and Convex Optimization. In *ICLR*, 2021.
- Hütter, J.-C. and Rigollet, P. Minimax estimation of smooth optimal transport maps. *The Annals of Statistics*, 49(2): 1166–1194, 2021.

- Janati, H., Muzellec, B., Peyré, G., and Cuturi, M. Entropic optimal transport between unbalanced gaussian measures has a closed form. *Advances in Neural Information Processing Systems*, 33, 2020.
- Kantorovitch, L. On the translocation of masses. *C. R. (Doklady) Acad. Sci. URSS (N.S.)*, 37:199–201, 1942.
- Mallasto, A., Gerolin, A., and Minh, H. Q. Entropy-regularized 2-wasserstein distance between gaussian measures. *Information Geometry*, pp. 1–35, 2021.
- Manole, T., Balakrishnan, S., Niles-Weed, J., and Wasserman, L. Plugin estimation of smooth optimal transport maps. *arXiv preprint arXiv:2107.12364*, 2021.
- Monge, G. Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences*, pp. 666–704, 1781.
- Moriel, N., Senel, E., Friedman, N., Rajewsky, N., Karaïskos, N., and Nitzan, M. Novosparc: flexible spatial reconstruction of single-cell gene expression with optimal transport. *Nature Protocols*, 16(9):4177–4200, 2021.
- Muzellec, B., Vacher, A., Bach, F., Vialard, F.-X., and Rudi, A. Near-optimal estimation of smooth transport maps with kernel sums-of-squares. *arXiv preprint arXiv:2112.01907*, 2021.
- Pal, S. On the difference between entropic cost and the optimal transport cost. *arXiv preprint arXiv:1905.12206*, 2019.
- Paty, F.-P., d’Aspremont, A., and Cuturi, M. Regularity as regularization: Smooth and strongly convex brenier potentials in optimal transport. In *International Conference on Artificial Intelligence and Statistics*, pp. 1222–1232. PMLR, 2020.
- Petersen, K. B., Pedersen, M. S., et al. The matrix cookbook. *Technical University of Denmark*, 7(15):510, 2008.
- Peyré, G. and Cuturi, M. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6): 355–607, 2019.
- Pooladian, A.-A. and Niles-Weed, J. Entropic estimation of optimal transport maps. *arXiv preprint arXiv:2109.12004*, 2021.
- Salimans, T., Zhang, H., Radford, A., and Metaxas, D. Improving GANs using optimal transport. In *International Conference on Learning Representations*, 2018.
- Santambrogio, F. Optimal transport for applied mathematicians. 2015.
- Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Gould, J., Liu, S., Lin, S., Berube, P., et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- Schmitzer, B. Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM Journal on Scientific Computing*, 41(3):A1443–A1481, 2019.
- Van Handel, R. Probability in high dimension. Technical report, PRINCETON UNIV NJ, 2014.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. In *Compressed sensing*, pp. 210–268. Cambridge Univ. Press, Cambridge, 2012.
- Yang, K. D., Damodaran, K., Venkatachalapathy, S., Soylemezoglu, A. C., Shivashankar, G., and Uhler, C. Predicting cell lineages using autoencoders and optimal transport. *PLoS computational biology*, 16(4):e1007828, 2020.

A. Remaining (Entropic) OT facts

Let $\Lambda_\varepsilon = (2\pi\varepsilon)^{d/2}$. For two measures with bounded densities and compact support, OT_ε admits the following *dynamical* formulation (Chizat et al., 2020; Conforti & Tamanini, 2021)

$$\begin{aligned} \text{OT}_\varepsilon(P, Q) + \varepsilon \log(\Lambda_\varepsilon) &= \inf_{\rho, v} \int_0^1 \int_{\mathbb{R}^d} \left(\frac{1}{2} \|v(t, x)\|_2^2 + \frac{\varepsilon^2}{8} \|\nabla_x \log(\rho(t, x))\|_2^2 \right) \rho(t, x) \, dx \, dt \\ &\quad - \frac{\varepsilon}{2} (\text{Ent}(P) + \text{Ent}(Q)), \end{aligned} \quad (19)$$

subject to $\partial_t \rho + \nabla \cdot (\rho v) = 0$, called the *continuity equation*, with $\rho(0, \cdot) = p(\cdot)$ and $\rho(1, \cdot) = q(\cdot)$.

When $Q = P$, we can upper bound the left-hand side of Equation (19) with the naive choices $v(t, x) \equiv 0$ and $\rho(t, \cdot) = p(\cdot)$, resulting in

$$\text{OT}_\varepsilon(P, P) + \varepsilon \log(\Lambda_\varepsilon) \leq \int_{\mathbb{R}^d} \frac{\varepsilon^2}{8} \|\nabla \log(p(x))\|^2 p(x) \, dx - \varepsilon \text{Ent}(P),$$

or equivalently,

$$\text{OT}_\varepsilon(P, P) + \varepsilon \log(\Lambda_\varepsilon) + \varepsilon \text{Ent}(P) \leq \frac{\varepsilon^2}{8} I_0(P, P). \quad (20)$$

Finally, we recall the modified dual formulation of $\text{OT}_\varepsilon(P, P)$ from Pooladian & Niles-Weed (2021)

Proposition 5. *Assume P has finite second moment and let γ_ε be the optimal entropic plan from P to itself. Then,*

$$\text{OT}_\varepsilon(P, P) = \sup_{\eta \in L^1(\gamma_\varepsilon)} \int \eta \, d\gamma_\varepsilon - \varepsilon \iint e^{(\eta(x, y) - \frac{1}{2} \|x - y\|^2) / \varepsilon} \, dP(x) \, dP(y) + \varepsilon. \quad (21)$$

B. Proofs from Section 3

For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we define its *convex conjugate* by $f^*(y) := \sup_x \{t^\top y - f(x)\}$. If $X \sim \mathcal{N}(\mu, \Sigma)$, recall that for any $t \in \mathbb{R}^d$,

$$\mathbb{E}[\exp\{t^\top X\}] = \exp\left\{t^\top \mu + \frac{1}{2} t^\top \Sigma t\right\}. \quad (22)$$

Proof of Lemma 1. Our proof technique is based on (Pal, 2019; Pooladian & Niles-Weed, 2021). Let $q_\varepsilon^x(y)$ denote the density of $\mathcal{N}(x, \varepsilon I_d)$. Using Proposition 5, writing γ_ε as the optimal entropic plan from P to itself, we plug in the test function

$$\eta(x, y) = \varepsilon(\chi(x, y) - \log(\Lambda_\varepsilon) - \log(p(y)))$$

with $\chi(x, y) = h(x)^\top (y - x) - (\varepsilon/2) \|h(x)\|^2$, where we omit the dependence on h for the time being. This results in

$$\text{OT}_\varepsilon(P, P) \geq \sup_\chi \iint \varepsilon(\chi(x, y) - \log(\Lambda_\varepsilon) - \log(p(y))) \, d\gamma_\varepsilon(x, y) - \varepsilon \iint e^{\chi(x, y)} q_\varepsilon^x(y) \, dy \, dP(x) + \varepsilon \quad (23)$$

$$= \sup_\chi \varepsilon \iint \chi(x, y) \, d\gamma_\varepsilon(x, y) - \varepsilon \log(\Lambda_\varepsilon) - \varepsilon \text{Ent}(P) - \varepsilon \iint e^{\chi(x, y)} q_\varepsilon^x(y) \, dy \, dP(x) + \varepsilon \quad (24)$$

Rearranging, this yields

$$\sup_\chi \iint \chi(x, y) \, d\gamma_\varepsilon(x, y) - \iint \left[e^{\chi(x, y)} - 1 \right] q_\varepsilon^x(y) \, dy \, dP(x) \leq \frac{1}{\varepsilon} (\text{OT}_\varepsilon(P, P) + \varepsilon \log(\Lambda_\varepsilon) + \varepsilon \text{Ent}(P)). \quad (25)$$

It follows by Equation (20) that the right-hand side is bounded above by $\frac{\varepsilon}{8} I_0(P, P)$. By well-known properties of moment generating functions of Gaussians (see Equation (22)), it also holds that

$$\mathbb{E}_{Y \sim q_\varepsilon^x} \left[e^{v^\top (Y - x) - (\varepsilon/2) \|v\|^2} - 1 \right] = 0 \quad (26)$$

for all $v \in \mathbb{R}^d$. Identifying $h(x)$ with v , we see that this implies

$$\iint \left[e^{\chi(x,y)} - 1 \right] q_\varepsilon^x(y) dy dP(x) = \int \left\{ \mathbb{E}_{Y \sim q_\varepsilon^x} \left[e^{h(x)^\top (Y-x) - (\varepsilon/2) \|h(x)\|^2} - 1 \right] \right\} dP(x) = 0. \quad (27)$$

Our bound now simplifies to

$$\sup_h \iint h(x)^\top (y-x) - \frac{\varepsilon}{2} \|h(x)\|^2 d\gamma_\varepsilon - 0 \leq \frac{\varepsilon}{8} I_0(P, P). \quad (28)$$

Taking $h(x) = -\varepsilon^{-1} \nabla \alpha_\varepsilon$, we use the fact that $\mathbb{E}_{\gamma_\varepsilon} [Y|X = x] = x - \nabla \alpha_\varepsilon(x)$ by definition, and that $\gamma_\varepsilon \in \Pi(P, P)$, resulting in

$$\iint (-\varepsilon^{-1} \nabla \alpha_\varepsilon(x))^\top (y-x) - \frac{1}{2\varepsilon} \|\nabla \alpha_\varepsilon(x)\|^2 d\gamma_\varepsilon = \frac{1}{2\varepsilon} \int \|\nabla \alpha_\varepsilon(x)\|^2 dP(x) \leq \frac{\varepsilon}{8} I_0(P, P). \quad (29)$$

Rearranging the remaining constants gives the desired bound. \square

Proof of Theorem 2. Expanding the square directly and applying Cauchy-Schwarz gives

$$\begin{aligned} \mathcal{R}(T_\varepsilon^D) &= \|T_\varepsilon^D - T_0\|_{L^2(P)}^2 \\ &= \|(T_\varepsilon - T_0) + \nabla \alpha_\varepsilon\|_{L^2(P)}^2 \\ &= \|T_\varepsilon - T_0\|_{L^2(P)}^2 + \|\nabla \alpha_\varepsilon\|_{L^2(P)}^2 + 2\langle T_\varepsilon - T_0, \nabla \alpha_\varepsilon \rangle_{L^2(P)} \\ &\leq \mathcal{R}(T_\varepsilon) + \|\nabla \alpha_\varepsilon\|_{L^2(P)}^2 + 2\sqrt{\mathcal{R}(T_\varepsilon)} \|\nabla \alpha_\varepsilon\|_{L^2(P)} \\ &\leq \mathcal{R}(T_\varepsilon) + (\varepsilon^2/4) I_0(P, P) + \varepsilon I_0(P, P) \sqrt{\mathcal{R}(T_\varepsilon)} \\ &\leq \mathcal{R}(T_\varepsilon) + \frac{3\varepsilon}{2} I_0(P, P) \sqrt{\mathcal{R}(T_\varepsilon)}, \end{aligned}$$

where the last inequality holds for ε small enough. The second claim follows by applying Proposition 1 and expanding the terms. \square

Proof of Proposition 2. Since π_ε and $P \otimes Q$ have first marginal P , we begin by applying the chain-rule for the KL divergence:

$$D_{\text{KL}}(\pi_\varepsilon \| P \otimes Q) = \mathbb{E}_{X \sim P} D_{\text{KL}}(\pi_\varepsilon^X \| Q),$$

where we denote π_ε^X as the (random) conditional density of π_ε given X . Since P and Q are compactly supported, we can apply the T_1 -transport inequality (Van Handel, 2014),

$$D_{\text{KL}}(\pi_\varepsilon \| P \otimes Q) = \mathbb{E}_{X \sim P} D_{\text{KL}}(\pi_\varepsilon^X \| Q) \gtrsim \mathbb{E}_{X \sim P} W_1^2(\pi_\varepsilon^X, Q).$$

We can further lower-bound this by the squared difference of the means of π_ε^X and Q ,

$$D_{\text{KL}}(\pi_\varepsilon \| P \otimes Q) \gtrsim \mathbb{E}_{X \sim P} W_1^2(\pi_\varepsilon^X, Q) \geq \mathbb{E}_{X \sim P} \|\mathbb{E}_{\pi_\varepsilon^X}[Y|X] - \mu_Q\|^2.$$

Taking the limit $\varepsilon \rightarrow \infty$, we invoke Lemma 3 and obtain the first claim:

$$\|T_\varepsilon - \mu_Q\|_{L^2(P)}^2 \lesssim D_{\text{KL}}(\pi_\varepsilon \| P \otimes Q) \rightarrow 0.$$

By an identical argument, it follows that

$$\|\text{Id} - \nabla \alpha_\varepsilon - \mu_P\|_{L^2(P)}^2 \rightarrow 0 \quad (\varepsilon \rightarrow \infty),$$

and thus the second claim follows by an application of the triangle inequality:

$$\begin{aligned} \|T_\varepsilon^D - (\mu_Q - \mu_P + \text{Id})\|_{L^2(P)}^2 &= \|(T_\varepsilon - \mu_Q) + (\nabla \alpha_\varepsilon - (-\mu_P + \text{Id}))\|_{L^2(P)}^2 \\ &\lesssim \|T_\varepsilon - \mu_Q\|_{L^2(P)}^2 + \|\nabla \alpha_\varepsilon + \mu_P - \text{Id}\|_{L^2(P)}^2 \\ &\rightarrow 0 \end{aligned}$$

\square

Proof of Theorem 3. By an application of the (reverse) triangle inequality,

$$\begin{aligned} \left| \|T_\varepsilon - T_0\|_{L^2(P)} - \sqrt{\text{Var}(Q)} \right| &= \left| \|T_\varepsilon - T_0\|_{L^2(P)} - \|\mu_Q - Y\|_{L^2(Q)} \right| \\ &= \left| \|T_\varepsilon - T_0\|_{L^2(P)} - \|\mu_Q - T_0\|_{L^2(P)} \right| \\ &\leq \|T_\varepsilon - \mu_Q\|_{L^2(P)}, \end{aligned}$$

which converges to zero by Proposition 2 as $\varepsilon \rightarrow \infty$. Taking squares, we have the first claim; the second follows by a similar argument. \square

Proof of Theorem 4. Fix $\varepsilon \in (0, 1)$ and take $P = \mathcal{N}(0, 1)$ and $Q = \mathcal{N}(0, \sigma^2)$. Set $\sigma^2 = \varepsilon^{2m}$ for $m > 0$ to be specified. Then

$$\begin{aligned} \mathcal{R}(T_\varepsilon) &= ((\varepsilon^{2m} + (\varepsilon^2/4))^{0.5} - (\varepsilon/2) + \varepsilon^m)^2, \\ \mathcal{R}(T_\varepsilon^D) &= ((\varepsilon^{2m} + \varepsilon^2/4)^{0.5} + 1 - \varepsilon^m - (1 + \varepsilon^2/4)^{0.5})^2. \end{aligned}$$

Taking $m \rightarrow \infty$, we see that $\mathcal{R}(T_\varepsilon) \rightarrow 0$ and $\mathcal{R}(T_\varepsilon^D) \rightarrow (\varepsilon/2) + 1 - (1 + \varepsilon^2/4)^{0.5} > 0$. Thus, for any $M > 0$, there exists an m large enough that $\mathcal{R}(T_\varepsilon^D) \geq M \cdot \mathcal{R}(T_\varepsilon)$. \square

C. Proofs from Section 4

Proof of Proposition 3. The optimal entropic plan between two Gaussians is known to be the following multivariate Gaussian distribution (Janati et al., 2020)

$$\pi_\varepsilon = N \left(\begin{pmatrix} 0 \\ b \end{pmatrix}, \begin{pmatrix} A & \Sigma_\varepsilon \\ \Sigma_\varepsilon^\top & B \end{pmatrix} \right), \quad (30)$$

where $\Sigma_\varepsilon = A^{1/2}(A^{1/2}BA^{1/2} + (\varepsilon^2/4)I)^{1/2}A^{-1/2} - (\varepsilon/2)I$. The conditional mean of this joint multivariate Gaussian is known to have the following closed form solution (Petersen et al., 2008), which completes the claim

$$\mathbb{E}_{\pi_\varepsilon}[Y|X = x] = \Sigma_\varepsilon^\top A^{-1}x + b = C_\varepsilon^{AB}x + b.$$

\square

Proof of Proposition 4. If γ_ε is the optimal entropic plan arising from $\text{OT}_\varepsilon(P, P)$, mimicking Equation (30), we have

$$\gamma_\varepsilon = N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} A & \tilde{\Sigma}_\varepsilon \\ \tilde{\Sigma}_\varepsilon^\top & A \end{pmatrix} \right),$$

where $\tilde{\Sigma}_\varepsilon = A^{1/2}(A^2 + (\varepsilon^2/4)I)^{1/2}A^{-1/2} + (\varepsilon/2)I$. The conditional expectation in this case simplifies to

$$\mathbb{E}_{\pi_\varepsilon^A}[Y|X = x] = \tilde{\Sigma}_\varepsilon^\top A^{-1}x = C_\varepsilon^{AA}x.$$

By combining this computation with Eq. (11) and Eq. (13), we have that

$$\nabla \alpha_\varepsilon(x) = x - (C_\varepsilon^{AA}x) = (I - C_\varepsilon^{AA})x.$$

\square

Proof of Theorem 5. We begin by creating short-hand notation for convenience. Writing $M = A^{1/2}BA^{1/2}$, then $C_0^{AB} = A^{-1/2}M^{1/2}A^{-1/2}$. We also denote $M_\varepsilon := (M + (\varepsilon^2/4)I)^{1/2} - M^{1/2}$, which is positive semidefinite. A direct computation yields (cf. (Petersen et al., 2008))

$$\|T_\varepsilon - T\|_{L^2(P)}^2 = \mathbb{E}_{X \sim P} \|DX\|^2 = \mathbb{E}_{X \sim P} X^\top D^2 X = \text{Tr}(D^2 A), \quad (31)$$

where $D := C_\varepsilon^{AB} - C_0^{AB} = A^{-1/2}M_\varepsilon A^{-1/2} - (\varepsilon/2)A^{-1}$. Computing the square and post-multiplying by A gives

$$\begin{aligned} D^2 A &= \frac{\varepsilon^2}{4} A^{-2} A + (A^{-1/2} M_\varepsilon A^{-1/2})^2 A - \frac{\varepsilon}{2} [A^{-1/2} M_\varepsilon A^{-1/2} A^{-1} + A^{-1} A^{-1/2} M_\varepsilon A^{-1/2}] A \\ &= \frac{\varepsilon^2}{4} A^{-1} + (A^{-1/2} M_\varepsilon A^{-1/2})^2 A - \frac{\varepsilon}{2} [A^{-1/2} M_\varepsilon A^{-1/2} + A^{-1} A^{-1/2} M_\varepsilon A^{-1/2} A]. \end{aligned}$$

Since the trace operator is linear and invariant under permutations, we arrive at

$$\begin{aligned} \text{Tr}(D^2 A) &= \frac{\varepsilon^2}{4} \text{Tr}(A^{-1}) + \text{Tr}(A^{-1/2} M_\varepsilon A^{-1/2} A^{-1/2} M_\varepsilon A^{-1/2} A) - \varepsilon \text{Tr}(A^{-1/2} M_\varepsilon A^{-1/2}) \\ &= \frac{\varepsilon^2}{4} \text{Tr}(A^{-1}) + \text{Tr}(A^{-1/2} (M_\varepsilon^2 - \varepsilon M_\varepsilon) A^{-1/2}) \\ &\leq \frac{\varepsilon^2}{4} \text{Tr}(A^{-1}) + \text{Tr}(M_\varepsilon^2 A^{-1}). \end{aligned}$$

Invoking Lemma 2, we arrive at

$$\|T_\varepsilon - T_0\|_{L^2(P)}^2 \leq \frac{\varepsilon^2}{4} \text{Tr}(A^{-1}) + \varepsilon^4 \text{Tr}(A^{-1}) C_{\lambda_{\min}(M)},$$

but a direct computation shows that

$$\begin{aligned} I_0(P, P) &= \int \|\nabla \log p(x)\|^2 p(x) \, dx \\ &= \int \|A^{-1} x\|^2 p(x) \, dx \\ &= \mathbb{E}_{X \sim \mathcal{N}(0, A)} [X^\top A^{-2} X] \\ &= \text{Tr}(A^{-2} A) \\ &= \text{Tr}(A^{-1}), \end{aligned}$$

which completes the proof. □

Proof of Theorem 6. We keep the same definitions of M and M_ε as in Theorem 5, and also introduce

$$A_\varepsilon := (A^2 + (\varepsilon^2/4)I)^{1/2} - A.$$

As last time, we have by direct computation that

$$\|T_\varepsilon^D - T_0\|_{L^2(P)}^2 = \text{Tr}(\tilde{D}^2 A)$$

where

$$\tilde{D} := \tilde{C}_\varepsilon^{AB} - C_0^{AB} = A^{-1/2} (M_\varepsilon - A_\varepsilon) A^{-1/2},$$

where both $A_\varepsilon, M_\varepsilon \succeq 0$. We begin by computing the squared term inside the trace:

$$(\tilde{C}_\varepsilon^{AB} - C_0^{AB})^2 = (A^{-1/2} M_\varepsilon A^{-1/2})^2 + (A^{-1/2} A_\varepsilon A^{-1/2})^2 - A^{-1/2} M_\varepsilon A^{-1} A_\varepsilon A^{-1/2} - A^{-1/2} A_\varepsilon A^{-1} M_\varepsilon A^{-1/2}.$$

By the permutation invariance of trace, we arrive at

$$\text{Tr}(\tilde{D}^2 A) = \text{Tr}((M_\varepsilon^2 + A_\varepsilon^2 - 2A_\varepsilon M_\varepsilon) A^{-1}).$$

Again, by Lemma 2, we have

$$M_\varepsilon = \frac{\varepsilon^2}{8} M^{-1/2} + O_{A,B}(\varepsilon^4), \quad A_\varepsilon = \frac{\varepsilon^2}{8} A^{-1} + O_A(\varepsilon^4).$$

Taking squares and expanding within the trace operator yields the desired claim

$$\text{Tr}((M_\varepsilon^2 + A_\varepsilon^2 - 2A_\varepsilon M_\varepsilon) A^{-1}) = \frac{\varepsilon^4}{64} \text{Tr}(M^{-1} A^{-1} + A^{-3} - 2M^{-1/2} A^{-2}) + O_{A,B}(\varepsilon^6)$$

□

D. Lemmas

Lemma 2. For $C \succ 0$, the following expansion holds

$$(C + \eta I)^{1/2} = C^{1/2} + \sum_{k \geq 1} (-1)^{k-1} \eta^k C^{1/2-k} c_k,$$

with $c_k = 2^{-k} \cdot (2k - 1)!!$. In particular, taking $\eta = \varepsilon^2/4$, we have that the expansion is, to second order

$$(C + (\varepsilon^2/4)I)^{1/2} = C^{1/2} + \frac{\varepsilon^2}{8} C^{-1/2} - \frac{\varepsilon^4}{128} C^{-3/2} + R_\varepsilon, \quad (32)$$

where the remainder term R_ε satisfies $\|R_\varepsilon\| \lesssim \lambda_{\min}^{-5/2}(C) \varepsilon^6$.

Proof. Since I commutes with $C \succ 0$, the above claim follows by considering the Taylor expansion of the scalar function $f(x) = (x + \eta)^{1/2}$ and replacing the relevant quantities with their matrix counterparts. The second claim follows by noticing that the second order term is dominated by the smallest eigenvalue of C raised to the relevant power, which amounts to $\lambda_{\min}^{-5/2}(C)I$. \square

Lemma 3. Let π_ε be the minimizer for $OT_\varepsilon(P, Q)$ under with P and Q having compact support. Then $D_{\text{KL}}(\pi_\varepsilon \| P \otimes Q) \rightarrow 0$ as $\varepsilon \rightarrow \infty$.

Proof of Lemma 3. Note that $OT_\varepsilon(P, Q)$ is bounded above by $C^* := \iint \frac{1}{2} \|x - y\|^2 d(P \otimes Q) > 0$, so

$$\iint \frac{1}{2} \|x - y\|^2 d\pi_\varepsilon + \varepsilon D_{\text{KL}}(\pi_\varepsilon \| P \otimes Q) \leq C^*.$$

The first term on the left-hand side is also bounded above by some constant, as $\pi_\varepsilon \in \Gamma(P, Q)$, where P and Q have compact support. Taking the limit of $\varepsilon \rightarrow \infty$, it must be the case that $D_{\text{KL}}(\pi_\varepsilon \| P \otimes Q) \rightarrow 0$. \square

E. Numerics

E.1. Compact, elliptically contoured distributions

Here we recall some facts about how to sample from compact, elliptically contoured distributions, as outlined in (Chizat et al., 2020). To sample from such a (centered) distribution, one can use the following simple recipe: for a given covariance matrix $A \in \mathbb{R}^d$,

1. Sample $U \sim \text{Unif}(\mathbb{S}^{d-1})$ (Uniform distribution over the $d - 1$ sphere),
2. Sample $Z \sim \mathcal{N}(0, 1)$,
3. Set $R := \alpha |\arctan Z/\beta|^{1/d}$, where $\alpha > 0$ such that $\mathbb{E}[R^2] = d$,
4. $X = R \cdot A^{1/2}U$.

In Step 3, we use the choice of $\beta = 2$ and R is ultimately computed via Monte Carlo integration with 10^7 points. For simplicity, we choose the same covariance matrices as in (Chizat et al., 2020) as their code is publicly available.

E.2. Synthetic examples

We run a log-stabilized version of Sinkhorn’s algorithm to alleviate numerical instabilities that would otherwise arise for small choices of ε . For both $d = 5$ and $d = 10$, we have $N_S = 50000$ as the number of points to approximate the MSE $\|\hat{T} - T_0\|_{L^2(P)}^2$ via Monte-Carlo integration after having learned the maps. For fixed ε (we chose $\varepsilon = 0.5$ in both $d = 5$ and $d = 10$), we sample n points from the source (where n varies between 100 and 10000), and then re-sample n points from the source and map them according to T , effectively generating samples from the target distribution. We run this procedure 20 times to generate error bars on the plots.

Example (E3) is between two elliptically contoured distributions, where the optimal transport map is known to be a linear map. Specifically, for centered, compact, elliptically contoured distributions with covariance A for the source and B for the target, then the optimal transport map is written as

$$T(x) = A^{1/2}(A^{-1/2}BA^{-1/2})^{1/2}A^{1/2}x.$$

Example (E4) is the coordinate-wise exponential map, $T(x) = (\exp(x_i))_{i=1}^d$, with source distribution $P = \text{Unif}([-1, 1]^d)$.

REMAINING PLOTS IN 5D






Figure 8. Remaining examples in $d = 5$


EXAMPLES PLOTS IN 10D




(a) Example (E1)




(b) Example (E2)



(c) Example (E2')



(d) Example (E3)



(e) Example (E4)

Figure 9. Previous examples but with $d = 10$

E.3. Gaussian to Gaussian map estimation

E.3.1. RANDOMLY GENERATING COVARIANCE MATRICES

We follow a procedure similar to what is outlined in (Chizat et al., 2020), which we re-write here for completeness. We begin with a matrix $M \in \mathbb{R}^{d \times k}$ with $M_{ij} \sim N(0, 1)$ i.i.d. entries, and $k = d/\alpha$ for $\alpha \in (0, 1)$. Defining $\tilde{A} = MM^T \succeq 0$, it is known that the eigenvalues of \tilde{A} are contained (with high probability) a small enlargement of the interval $[(1 - \sqrt{\alpha})^2, (1 + \sqrt{\alpha})^2]$; this follows from Gordon’s inequality (Vershynin, 2012, Theorem 5.32). Choosing $\alpha = 1/3$ and writing $A = \gamma \tilde{A} / \text{Tr}(\tilde{A})$ for a positive constant γ , we can (randomly) control the spread of the eigenvalues of the covariance matrix $A \in \mathbb{R}^{d \times d}$ through the parameter $\gamma > 0$.

For the “concentrated” covariance matrices, we chose $\gamma = 0.2$ for $d = 2$ and $\gamma = 5$ in $d = 15$. For the “spread out” covariance matrices, we chose $\gamma = 5$ for $d = 2$ and $\gamma = 20$ in $d = 15$.

E.3.2. REMAINING PLOTS

In all examples here, $P = \mathcal{N}(0, I_d)$ and $Q = \mathcal{N}(0, \Sigma)$, where Σ is a randomly generated covariance matrix that has some eigenvalues greater than 1, and some less than 1, which we call “spread out” (see above). We present examples in $d = 2$ and $d = 15$, illustrating finite-sample effects in both low and high dimensions. We perform the MSE estimates using Monte-Carlo integration with $5 \cdot 10^5$ samples. For each figure, we perform 15 random trials of map estimation where we learn the map with the value N in the plots, and vary ε .

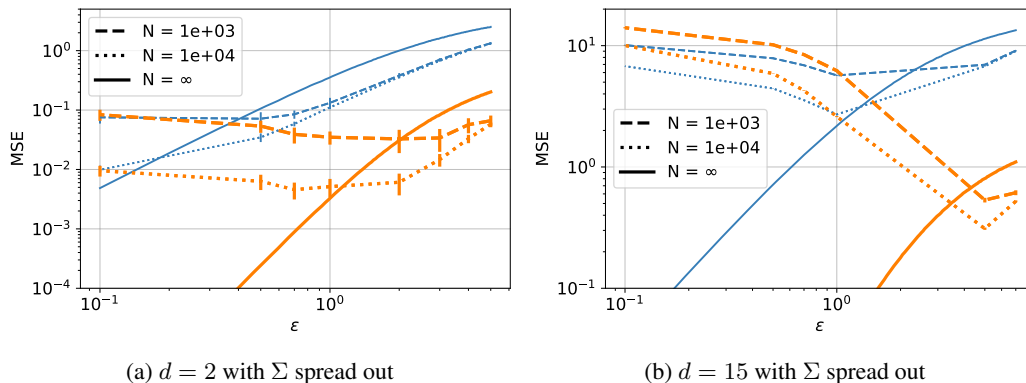


Figure 10. Finite-sample effect for \hat{T}_ε vs. \hat{T}_ε^D when Σ is spread out