

An explicit analysis of the entropic penalty in linear programming

Jonathan Weed

JWEED@MIT.EDU

Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Editors: Sébastien Bubeck, Vianney Perchet and Philippe Rigollet

Abstract

Solving linear programs by using entropic penalization has recently attracted new interest in the optimization community, since this strategy forms the basis for the fastest-known algorithms for the optimal transport problem, with many applications in modern large-scale machine learning. Crucial to these applications has been an analysis of how quickly solutions to the penalized program approach true optima to the original linear program. More than 20 years ago, Cominetti and San Martín showed that this convergence is exponentially fast; however, their proof is asymptotic and does not give any indication of how accurately the entropic program approximates the original program for any particular choice of the penalization parameter. We close this long-standing gap in the literature regarding entropic penalization by giving a new proof of the exponential convergence, valid for any linear program. Our proof is non-asymptotic, yields explicit constants, and has the virtue of being extremely simple. We provide matching lower bounds and show that the entropic approach does not lead to a near-linear time approximation scheme for the linear assignment problem.

Keywords: Entropic penalization, optimal transport, assignment problem, Sinkhorn algorithm.

1. Introduction

In 1992, Fang initiated the study of the *entropic penalty* for linear programs. Given a basic linear program of the form

$$\begin{aligned} \min \quad & c^\top x \\ \text{subject to} \quad & Ax = b \\ & x \geq 0, \end{aligned} \tag{LP}$$

he proposed to solve instead the penalized program

$$\begin{aligned} \min \quad & c^\top x - \eta^{-1}H(x) \\ \text{subject to} \quad & Ax = b, \end{aligned} \tag{Pen}$$

where $H(x) := \sum_i x_i \log \frac{1}{x_i}$ is the *Shannon entropy* of x viewed as a probability vector and $\eta > 0$ is a penalization parameter. The term $-\eta^{-1}H(x)$ plays the role of a strongly convex regularizer, which also enforces the constraint $x \geq 0$. As $\eta \rightarrow \infty$, we recover (LP); however, one hopes that solving (Pen) is significantly easier.

Solving linear programs via entropic penalization does not initially seem like an especially attractive choice. Unlike the well known logarithmic penalty, the entropic penalty is not self-concordant, which makes it a poor barrier function for interior point methods (Boyd and Vandenberghe, 2004). Nevertheless, the entropic penalty has been applied to various linear and nonlinear problems with empirical success (Fang et al., 1997). It is also notable for its connection to other

fields. In statistics, it has been used as a tool for model selection and aggregation (Juditsky et al., 2008; Rigollet and Tsybakov, 2011) and is intimately related to the maximum entropy principle for statistical inference (Jaynes, 1982) and to maximum likelihood estimation (Chrétien and Hero, 2000). The entropic penalty is also closely connected to first-order optimization methods such as mirror descent (Bubeck, 2015) and to online learning algorithms for combinatorially structured problems (Freund and Schapire, 1997; Cesa-Bianchi and Lugosi, 2006; Helmbold and Warmuth, 2009; Koolen et al., 2010; Audibert et al., 2013).

The recent resurgence of interest in the entropic penalty in the machine-learning community has been driven by the fact that it can be used to obtain state-of-the-art methods for the *optimal transport problem* (Cuturi, 2013; Cuturi and Doucet, 2014; Solomon et al., 2015; Genevay et al., 2016; Benamou et al., 2016; Altschuler et al., 2017). The use of the entropic penalty for such problems dates back to Schrödinger (1931) (see Léonard, 2014) and to Brègman (1967), who noted the connection between the entropic penalty and the computation of a projection onto the feasible set $\{x : Ax = b, x \geq 0\}$ with respect to the generalized Kullbeck-Leibler divergence.

What makes this penalty especially useful for transport problems is that the solution to (Pen) can be computed quickly by a simple iterative algorithm known as the *Sinkhorn* or *RAS* algorithm (Sinkhorn, 1967). This fact was popularized by Cuturi (2013), and his work led to widespread adoption of the entropic penalty for computing optimal transport. The introduction of the entropic penalty makes an enormous difference in practice: since optimal transport can be formulated as a linear program, it can, of course, be solved in polynomial time, but the numerical experiments conducted by Cuturi (2013) indicate that the same linear program with an entropic penalty can be solved up to 10,000 times faster, as long as η is not too large. On the other hand, those experiments also showed that solving the penalized program becomes costly as η increases.

This same phenomenon is present in theory as well as in practice. A recent theoretical analysis (Altschuler et al., 2017) of the method of Cuturi (2013) suggests that the time required to solve the optimal transport problem via entropic regularization scales linearly with η . Even the guarantees of the most recent algorithms for solving (Pen) decay as η grows. For example, when $\{Ax = b, x \geq 0\}$ defines the Birkhoff polytope (Brualdi, 2006) of $n \times n$ doubly stochastic matrices, an approximate solution to (Pen) can be found in time¹ $\tilde{O}(n^2 \|c\|_\infty \eta)$ (Cohen et al., 2017; Allen-Zhu et al., 2017). In particular, when $\eta = \tilde{O}(1)$, the algorithms of Cohen et al. (2017) and Allen-Zhu et al. (2017) run in time which is *nearly linear* in the size of the input.

Nevertheless, if we wish to obtain a good approximation to the solution of (LP), η cannot be taken too small. In the $\eta \rightarrow 0$ limit, the solution to (Pen) converges to the maximum-entropy point in the feasible set, which may be far from the optimum. If the goal is to approximately solve (LP), then η must be large enough that the solution of (Pen) is still close to an optimum of (LP).

To summarize: the choice of η is essential. Too large, and the computational benefits of using the regularizer disappear; too small, and the entropic term induces significant bias and (Pen) is a poor approximation to the original problem. The chief aim of this work is to quantify this trade-off.

1. The notation $\tilde{O}(\cdot)$ hides polylogarithmic factors.

1.1. Prior work

The question of how well (Pen) approximates (LP) as a function of η was studied by [Cominetti and San Martín \(1994\)](#).² They showed that, under mild conditions, the optimal solution x^η to (Pen) approaches an optimal solution x^* to (LP) exponentially fast, in the sense that

$$\lim_{\eta \rightarrow \infty} \frac{\|x^\eta - x^*\|}{\exp(-M\eta)} = 0 \tag{1}$$

for some $M > 0$. However, their proof does not make it easy to determine the order of magnitude of M —and, in particular, its dependence on problem-specific quantities such as the dimension and size of the feasible set $\{x : Ax = b, x \geq 0\}$. This result is nevertheless tantalizing, insofar as it suggests that x^η will be close to x^* even for relatively small values of η . Of course, knowing the size of M is crucial to making this idea precise. Prior to this work, theirs was the most general analysis of (Pen) available.

1.2. Our contribution

In this work, motivated by the recent popularity of entropic penalization for optimal transport, we prove a version of (1) with easy-to-understand constants. Our analysis applies to any linear program of the form (LP). We show (Section 2) that the quality of the penalized solution x^η satisfies

$$c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x \leq \Delta \exp\left(-\eta \frac{\Delta}{R_1} + \frac{R_1 + R_H}{R_1}\right),$$

where Δ is the gap in objective value between an optimal vertex and any suboptimal vertex, and R_1 and R_H are the radius of the feasible set with respect to the ℓ_1 norm and the entropy, respectively. As a corollary, we obtain that the result (1) obtained by [Cominetti and San Martín \(1994\)](#) holds for any $M < \Delta/R_1$. In addition to making explicit their result, our proof has the virtue of being very simple, requiring only elementary facts about entropy. Moreover, we show (Section 3) that no general improvement in the dependence on Δ , R_1 , or R_H is possible, even for the simplest possible example, where the feasible set is the probability simplex.

Finally, specializing to the Birkhoff polytope (Section 4), we obtain nearly matching upper and lower bounds on the quality of the solution as a function of η . In particular, these imply that η cannot be taken to be $\tilde{O}(1)$, so that the entropic penalty is not a magic bullet for the assignment problem.

1.3. Assumptions

We assume throughout that $\mathcal{P} := \{x : Ax = b, x \geq 0\}$ is bounded. To ensure that (LP) is nontrivial, we assume that \mathcal{P} is nonempty and that $c^\top x$ is not constant over \mathcal{P} .

2. In fact, the main object of study of [Cominetti and San Martín \(1994\)](#) is the program $\max_y \left\{ b^\top y - \eta^{-1} \sum_i e^{-\eta(c_i - (A^\top y)_i)} \right\}$, which is the dual of (Pen) with the entropy replaced by the similar function $\bar{H}(x) := \sum_i x_i \left(\log \frac{1}{x_i} + 1 \right)$. They refer to the penalty appearing in this dual program as the *exponential penalty*. Our analysis applies equally well to their setting, but we focus on the vanilla entropic penalty for clarity.

1.4. Quantities of interest

For convenience, we collect here definitions of the three quantities Δ , R_1 , and R_H appearing in our bounds.

Definition 1 Let V be the set of vertices of \mathcal{P} . The suboptimality gap Δ is

$$\min_{v \in \mathcal{S}} c^\top (v - v^*),$$

where $v^* \in \operatorname{argmin}_{v \in V} c^\top v$ and $\mathcal{S} := \{v \in V : c^\top v > c^\top v^*\}$.

Definition 2 The ℓ_1 radius R_1 of \mathcal{P} is $\max_{x \in \mathcal{P}} \|x\|_1$.

Definition 3 The entropic radius R_H of \mathcal{P} is $\max_{x, y \in \mathcal{P}} H(x) - H(y)$.

2. Upper bound

In this section, we prove our main bound on the quality of the solution of (Pen). Our proof recovers the result of [Cominetti and San Martín \(1994\)](#) that the penalized solution approaches an optimal solution exponentially fast. Before doing so, however, we first prove a much simpler and weaker bound, which we call the *slow rate* (see [Fang and Tsao, 1993](#); [Altschuler et al., 2017](#), where this analysis also appeared):

Proposition 4 (Slow rate) For all $\eta > 0$,

$$c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x \leq \eta^{-1} (H(x^\eta) - H(x^*)) \leq \eta^{-1} R_H.$$

Proof Denote by x^* an optimal solution to (LP). Since x^* and x^η both lie in \mathcal{P} and x^η is an optimal solution to (Pen),

$$c^\top x^\eta - \eta^{-1} H(x^\eta) \leq c^\top x^* - \eta^{-1} H(x^*).$$

The claim follows. ■

Note that the slow rate is much worse than the fast rate we hope to prove; however, [Rigollet \(2017\)](#) noted that the slow rate is actually tight for an infinite-dimensional analogue of (LP). This indicates that the reason that a fast rate obtains for (LP) is that the finite-dimensional problem exhibits a *suboptimality gap* (known as an energy gap in the statistical physics literature; see [Mézard and Montanari, 2009](#)). Intuitively, the slow rate dominates the convergence until η is large enough that x^η is concentrated near enough to the optimal solution; after this point, convergence occurs exponentially fast. We will return to this point in Section 3.

We now turn to the main result.

Theorem 5 (Fast rate) If $\eta \geq \frac{R_1 + R_H}{\Delta}$, then the optimal solution x^η of (Pen) satisfies

$$c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x \leq \Delta \exp\left(-\eta \frac{\Delta}{R_1} + \frac{R_1 + R_H}{R_1}\right).$$

Theorem 5 implies a bound on the size of η required to obtain a solution of desired accuracy: to obtain a solution x^η satisfying $c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x \leq \varepsilon$, it suffices to take $\eta = \left(\frac{R_1}{\Delta} \log \frac{\Delta}{\varepsilon}\right)_+ + \frac{R_1 + R_H}{\Delta}$, where $(x)_+ := \max\{x, 0\}$.

Note that Theorem 5 only holds for η sufficiently large. The requirement that $\eta \geq \frac{R_1 + R_H}{\Delta}$ corresponds exactly to the requirement that the exponent appearing on the right side of the above equation is nonpositive. In Section 3, we show that this restriction is necessary, in the sense that there are penalized linear programs for which x^η does not make appreciable progress towards the minimizer until $\eta = \Omega\left(\frac{R_1 + R_H}{\Delta}\right)$.

The proof of Theorem 5 is elementary and relies on three simple lemmas about the entropy function, which we now state. These lemmas are easy to verify; proofs appear in Section 6. Recall the definition of the binary entropy function:

$$h(\lambda) = \lambda \log \frac{1}{\lambda} + (1 - \lambda) \log \frac{1}{1 - \lambda} \quad \forall \lambda \in [0, 1].$$

Lemma 6 *If x and y are nonnegative vectors and $\lambda \in [0, 1]$, then*

$$H(\lambda x + (1 - \lambda)y) \leq \lambda H(x) + (1 - \lambda)H(y) + \max\{\|x\|_1, \|y\|_1\} \cdot h(\lambda).$$

Lemma 7 *The function $f(\lambda) := \alpha h(\lambda) + \beta \lambda$ is increasing on the interval $[0, \frac{\beta}{\alpha + \beta}]$.*

Lemma 8 *If $0 \leq \rho \leq 1$, then*

$$\frac{h(\rho)}{\rho} \leq \log \frac{1}{\rho} + 1.$$

Proof [Proof of Theorem 5] Let V be the vertices of \mathcal{P} . Write $\mathcal{O} := \{v^* \in V : c^\top v^* = \min_{v \in V} c^\top v\}$ for the set of optimal vertex solutions for (LP), and let $\mathcal{S} := V \setminus \mathcal{O}$ be the set of suboptimal vertices. Since $x^\eta \in \mathcal{P} = \text{conv}(V)$, we can write

$$x^\eta = \sum_{v \in \mathcal{O}} \lambda_v v + \sum_{w \in \mathcal{S}} \lambda_w w$$

for some nonnegative vector λ satisfying $\sum_{v \in V} \lambda_v = 1$. If we let $\gamma = \sum_{w \in \mathcal{S}} \lambda_w$, then $x^\eta = (1 - \gamma)x^* + \gamma \tilde{x}$, where $x^* \in \text{conv}(\mathcal{O})$ and $\tilde{x} \in \text{conv}(\mathcal{S})$. Since x^* is a convex combination of elements of \mathcal{O} , it lies on the optimal face of \mathcal{P} and is an optimal solution to (LP). On the other hand, since \tilde{x} is a convex combination of suboptimal vertices, $c^\top(\tilde{x} - x^*) \geq \Delta$.

Let $g(\eta) := c^\top(x^\eta - x^*)$. We first prove two simple bounds on this quantity. First, we have a trivial lower bound:

$$g(\eta) = c^\top(x^\eta - x^*) = \gamma c^\top(\tilde{x} - x^*) \geq \Delta \gamma. \quad (2)$$

On the other hand, Proposition 4 implies

$$g(\eta) \leq \eta^{-1}(H(x^\eta) - H(x^*)). \quad (3)$$

By Lemma 6,

$$\begin{aligned} H(x^\eta) - H(x^*) &\leq ((1 - \gamma)H(x^*) + \gamma H(\tilde{x}) + R_1 h(\gamma)) - H(x^*) \\ &= R_1 h(\gamma) + \gamma(H(\tilde{x}) - H(x^*)) \\ &\leq R_1 h(\gamma) + R_H \gamma. \end{aligned} \quad (4)$$

By assumption, $\eta \geq \frac{R_1 + R_H}{\Delta}$, so using (2) and (3) yields

$$\gamma \leq \frac{g(\eta)}{\Delta} \leq \frac{\eta^{-1} R_H}{\Delta} \leq \frac{R_H}{R_1 + R_H} < 1.$$

Lemma 7 then implies that $R_1 h(\gamma) + R_H \gamma \leq R_1 h\left(\frac{g(\eta)}{\Delta}\right) + R_H \frac{g(\eta)}{\Delta}$. We can combine this observation with (4) to obtain

$$\eta g(\eta) \leq H(x^\eta) - H(x^*) \leq R_1 h\left(\frac{g(\eta)}{\Delta}\right) + R_H \frac{g(\eta)}{\Delta}.$$

Writing $\rho = \frac{g(\eta)}{\Delta}$ yields the fixed-point equation

$$(\eta\Delta - R_H)\rho \leq R_1 h(\rho).$$

By Lemma 8, $\frac{h(\rho)}{\rho} \leq \log \frac{1}{\rho} + 1$ for all $\rho \leq 1$. We obtain

$$\frac{\eta\Delta - R_H}{R_1} - 1 \leq \log \frac{1}{\rho},$$

and hence

$$g(\eta) = \Delta\rho \leq \Delta \exp\left(-\eta \frac{\Delta}{R_1} + \frac{R_1 + R_H}{R_1}\right),$$

as desired. ■

We also obtain a corollary which establishes the distance of x^η to the optimal face, which reproduces the result of [Cominetti and San Martín \(1994\)](#). Let $F := \text{conv}\{\mathcal{O}\}$ be the optimal face of \mathcal{P} with respect to the objective $c^\top x$, and denote by $d_1(x, F)$ the ℓ_1 distance of the point x to F .

Corollary 9 *If $\eta \geq \frac{R_1 + R_H}{\Delta}$, then*

$$d_1(x^\eta, F) \leq 2R_1 \exp\left(-\eta \frac{\Delta}{R_1} + \frac{R_1 + R_H}{R_1}\right).$$

In particular,

$$\lim_{\eta \rightarrow \infty} \frac{d_1(x^\eta, F)}{\exp(-M\eta)} = 0$$

for any $M < \Delta/R_1$.

Proof Using the notation of Theorem 5, we have that there exist points $x^*, \tilde{x} \in \mathcal{P}$ such that x^* is optimal and

$$x^\eta = (1 - \gamma)x^* + \gamma\tilde{x},$$

for $\gamma \leq \frac{g(\eta)}{\Delta} \leq \exp\left(-\eta \frac{\Delta}{R_1} + \frac{R_1 + R_H}{R_1}\right)$. We obtain

$$d_1(x^\eta, F) \leq \|x^\eta - x^*\|_1 \leq \gamma \|\tilde{x} - x^*\|_1 \leq 2\gamma R_1,$$

and the claim follows. ■

The quantity Δ is quite brittle, since it can be affected by the presence of even a single almost-optimal vertex whose objective value is very close to that of the optimal vertex. However, the definition of Δ can be relaxed slightly to account for this case, as the following corollary shows.

Corollary 10 For any $\tau > 0$, denote by \mathcal{O}_τ the set $\{v \in V : c^\top v - \min_{x \in \mathcal{P}} c^\top x \leq \tau\}$, and let $\Delta_\tau := \min_{w \in \mathcal{O}_\tau, v \in V \setminus \mathcal{O}_\tau} c^\top (v - w)$. If $\eta \geq \frac{R_1 + R_H}{\Delta_\tau}$, then the optimal solution x^η of (Pen) satisfies

$$c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x \leq \Delta_\tau \exp\left(-\eta \frac{\Delta_\tau}{R_1} + \frac{R_1 + R_H}{R_1}\right) + \tau.$$

Proof If $c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x \leq \tau$, the claim is vacuous, so assume that $c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x > \tau$. Let $\tau^* := \max_{v \in \mathcal{O}_\tau} c^\top v - \min_{x \in \mathcal{P}} c^\top x$, and note that $\tau^* \leq \tau$. Given an optimal solution x^* to (LP), let $\mathcal{P}_\tau := \mathcal{P} \cap \{x : c^\top (x - x^*) \geq \tau^*\}$. We have $\min_{x \in \mathcal{P}_\tau} c^\top x - \min_{x \in \mathcal{P}} c^\top x = \tau^* \leq \tau$. Moreover, the suboptimality gap of \mathcal{P}_τ is Δ_τ .

By assumption, $x^\eta \in \mathcal{P}_\tau$, and so x^η is the solution to (Pen) over the smaller polytope \mathcal{P}_τ . Applying Theorem 5 to \mathcal{P}_τ yields the claim. \blacksquare

While the quantities R_1 and R_H are easy to calculate, evaluating the suboptimality gap Δ is not easy in general. Nevertheless, as we noted above, intuition from statistical physics implies that some dependence on Δ is necessary to obtain exponential convergence, a point which we substantiate in Section 3. We note the obvious fact that this dependence can be removed for integral polytopes, which are a core object of study in combinatorial optimization (Schrijver, 2003).

Corollary 11 If \mathcal{P} is integral and the entries of c are integers, then

$$c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x \leq \exp\left(-\frac{\eta}{R_1} + \frac{R_1 + R_H}{R_1}\right)$$

for all $\eta \geq R_1 + R_H$.

Proof By definition, the vertices of \mathcal{P} have integer coordinates, so for any vertex v , if c is an integer vector then $c^\top v \in \mathbb{Z}$. Therefore if v^* is an optimal vertex and $c^\top v > c^\top v^*$, then $c^\top (v - v^*) \geq 1$, so $\Delta \geq 1$. \blacksquare

3. Lower bound

In this section, we present an explicit example of a simple family of linear programs for which our analysis is tight, up to constant factors. This example evinces the two phenomena present in Theorem 5: the convergence of $c^\top x^\eta$ to the optimum is slow until η is of order $\frac{R_1 + R_H}{\Delta}$, and once this threshold is reached convergence happens at precisely the speed indicated in the upper bound. This example also validates the intuition presented above about the necessary dependence on the suboptimality gap: exponentially fast convergence is obtained only when $c^\top x^\eta - \min_{x \in \mathcal{P}} c^\top x \leq \Delta$.

Fix positive constants α and β and a dimension $d \geq 2$. Let $c \in \mathbb{R}^d$ be given by

$$c_i = \begin{cases} 0 & \text{if } i = 0, \\ \alpha & \text{otherwise,} \end{cases}$$

and consider the linear program

$$\begin{aligned} \min \quad & c^\top x \\ \text{subject to} \quad & \sum_i x_i = \beta \\ & x \geq 0, \end{aligned} \tag{5}$$

Note that the polytope \mathcal{P} defined by the constraints of (5) is a rescaled version of the d -dimensional probability simplex. We focus on the following penalized program:

$$\begin{aligned} \min \quad & c^\top x - \eta^{-1} H(x) \\ \text{subject to} \quad & \sum_i x_i = \beta. \end{aligned} \tag{6}$$

We make the following simple observations about (5) and (6):

- The unique optimal solution to (5) is $x^* = e_1$, the first elementary basis vector, and $c^\top x^* = 0$.
- The maximum value of $c^\top x$ over \mathcal{P} is $\alpha\beta$, achieved at any vertex other than e_1 .
- For this polytope, $\Delta = \alpha\beta$, $R_1 = \beta$, and $R_H = \beta \log d$.

The penalized program (6) has an explicit solution, which is given by a rescaled version of the Gibbs distribution (Mézard and Montanari, 2009).

Proposition 12 *The optimal solution x^η to (6) is given by*

$$x_i^\eta = \frac{\beta e^{-\eta c_i}}{\sum_j e^{-\eta c_j}}.$$

The guarantee of Theorem 5 requires that $\eta \geq \frac{R_1 + R_H}{R_H} = \frac{1 + \log d}{\alpha}$. We now show that when η is significantly smaller than this quantity, the solution to the penalized program is far from the true optimum. Indeed, the following proposition establishes that we cannot even achieve a constant-factor improvement over the *maximum* value of $c^\top x$ over \mathcal{P} until η is of order $\frac{\log d}{\alpha}$.

Proposition 13 *For any $\varepsilon > 0$, if $\eta \leq \frac{\log \varepsilon d}{\alpha}$, then $c^\top x^\eta \geq (1 - \varepsilon)\alpha\beta$.*

Proof We prove the contrapositive. Note that $c^\top x^\eta = \alpha(\beta - x_0^\eta)$, so if $c^\top x^\eta < (1 - \varepsilon)\alpha\beta$ then $x_0^\eta > \varepsilon\beta$. By Proposition 12, we can write explicitly

$$x_0^\eta = \frac{\beta}{\sum_j e^{-\eta c_j}} = \frac{\beta}{1 + (d-1)e^{-\eta\alpha}}.$$

If $x_0^\eta > \varepsilon\beta$, then $de^{-\eta\alpha} < 1 + (d-1)e^{-\eta\alpha} < \frac{1}{\varepsilon}$, so $\eta > \frac{\log \varepsilon d}{\alpha}$, as claimed. ■

The next proposition shows that the Theorem 5 is tight up to a small constant factor.

Proposition 14 *If $\eta \geq \frac{1 + \log d}{\alpha} = \frac{R_1 + R_H}{\Delta}$, then*

$$c^\top x^\eta \geq \frac{1}{9}\alpha\beta e^{-\eta\alpha + 1 + \log d} = \frac{1}{9}\Delta \exp\left(-\eta \frac{\Delta}{R_1} + \frac{R_1 + R_H}{R_1}\right).$$

Proof If $\eta \geq \frac{\log d + 1}{\alpha}$, then $\sum_j e^{-\eta c_j} = 1 + (d-1)e^{-\eta\alpha} \leq 1 + \frac{d-1}{3d}$. Using Proposition 12,

$$c^\top x^\eta = (d-1)\alpha \frac{\beta e^{-\eta\alpha}}{\sum_j e^{-\eta c_j}} \geq \frac{1}{3}d\alpha\beta e^{-\eta\alpha} \geq \frac{1}{9}\alpha\beta e^{-\eta\alpha + 1 + \log d},$$

as claimed. ■

4. Entropic penalization for the assignment problem

In this section, we give an application of Theorem 5 to the *assignment problem*, a fundamental combinatorial optimization problem (Schrijver, 2003). Our motivation for analyzing this example explicitly is twofold. First, this is a case where entropic penalization has already been proposed as a good candidate algorithm (Kosowsky and Yuille, 1994; Sharify et al., 2011). Second, as noted in the introduction, new fast algorithms for the *matrix scaling problem* (Cohen et al., 2017; Allen-Zhu et al., 2017) show that a penalized version of the assignment problem with cost matrix C can be solved in time $\tilde{O}(n^2\|C\|_\infty\eta)$. These fast algorithms raise the prospect that entropic penalization could provide a near-linear time algorithm for the assignment problem, a major breakthrough (see, e.g., Mądry, 2013).

Whether this breakthrough is possible depends crucially on the size of η required to solve the problem accurately. The best bounds available from previous works on the problem (Kosowsky and Yuille, 1994; Sharify et al., 2011) require $\eta \gtrsim n \log n$ to achieve constant accuracy, which is just the guarantee given by the slow rate (Proposition 4). An open question implicit in these works is whether this is optimal, or whether $\eta = \tilde{O}(1)$ suffices. (In particular, this would imply a near-linear time algorithm for the assignment problem.) By applying Theorem 5 and exhibiting an almost-matching lower bound, we show exactly what rates are attainable for the Birkhoff polytope. In short, our hopes are dashed: η cannot be taken to be dimension free in general.

We first recall the problem. Given a bipartite graph with edge weights, the goal of the assignment problem is to find a minimum-cost perfect matching in the graph. This problem also has a well known linear programming formulation: given a matrix $C \in \mathbb{R}^{n \times n}$ of edge weights, the assignment problem is

$$\begin{aligned} \min \quad & \langle C, X \rangle \\ \text{subject to} \quad & X\mathbf{1} = \mathbf{1} \\ & X^\top \mathbf{1} = \mathbf{1} \\ & X \geq 0, \end{aligned} \tag{7}$$

The polytope given by the constraints $\{X\mathbf{1} = \mathbf{1}, X^\top \mathbf{1} = \mathbf{1}, X \geq 0\}$ is known as the *Birkhoff polytope*, and its vertices are the permutation matrices (Brauldi, 2006), a result known as the Birkhoff-von Neumann Theorem.

We first give an upper bound on the quality of X^η as a function of the regularization parameter η . We require a preliminary lemma, whose proof appears in Section 6.

Lemma 15 *The Birkhoff polytope \mathcal{B} has $R_1 = n$ and $R_H = n \log n$.*

Lemma 15 combined with Theorem 5 yields the following guarantee for the entropic penalty applied to the assignment problem. For normalization purposes, we assume that the entries of C are nonnegative integers, as is common in the combinatorial optimization literature.

Proposition 16 *An additive ε approximation to the assignment problem with cost matrix $C \in \mathbb{Z}_{\geq 0}^{n \times n}$ can be found by solving an entropy-penalized version of (7) with parameter $\eta = O(n \log \frac{n}{\varepsilon})$.*

Proof Since the entries of C are integers and the Birkhoff polytope is integral (Schrijver, 2003), Corollary 11 implies $\Delta \geq 1$. Theorem 5 implies that as long as $\varepsilon < 1$, we can obtain a solution \hat{X} such that

$$\langle C, \hat{X} \rangle - \min_{X^* \in \mathcal{B}} \langle C, X^* \rangle \leq \varepsilon$$

by solving the penalized program with $\eta = \frac{R_1}{\Delta} \log \frac{\Delta}{\varepsilon} + \frac{R_1+R_H}{\Delta} \leq n \log \frac{1}{\varepsilon} + n(1 + \log n)$, which is $O(n \log \frac{n}{\varepsilon})$. \blacksquare

Proposition 16 is disappointing: it guarantees exponential convergence of X^η to X^* only when $\eta \gtrsim n \log n$, a far cry from the hoped-for result that η could be taken $\tilde{O}(1)$. We now show that, up to logarithmic factors, this bound is tight. The following theorem implies that, even when $C \in \{0, 1\}^{n \times n}$, $\langle C, X^\eta \rangle$ can be bounded away from the optimal value if $\eta \ll n$.

Theorem 17 *Let $C \in \{0, 1\}^{n \times n}$ be the matrix given by*

$$C_{ij} := \begin{cases} 0 & \text{if } j = i \text{ or } j = i + 1, \\ 1 & \text{otherwise.} \end{cases} \quad (8)$$

If $\eta \leq n \log \frac{1-\varepsilon}{\varepsilon}$, then

$$\langle C, X^\eta \rangle - \langle C, X^* \rangle \geq \varepsilon$$

Proof The matrix defined in (8) admits the unique optimum solution $X^* = I$, the identity matrix, and optimal value $\langle C, I \rangle = 0$. For any permutation $\Pi \neq I$, on the other hand, $\langle C, \Pi \rangle \geq 1$.

We prove the contrapositive. By the Birkhoff-von Neumann theorem, we can write X^η as a convex combination of permutation matrices:

$$X^\eta = \lambda_I I + \sum_{\Pi \neq I} \lambda_\Pi \Pi.$$

By assumption, $\langle C, X^\eta \rangle < \varepsilon$, so $\lambda_I > 1 - \varepsilon$. This implies that $X_{ii}^\eta > 1 - \varepsilon$ for $1 \leq i \leq n$, and therefore that $X_{i,i+1}^\eta < \varepsilon$ for $1 \leq i \leq n - 1$.

Sinkhorn's theorem (Sinkhorn, 1967) combined with first-order optimality conditions for the penalized program guarantee that $X^\eta = D_1 A D_2$ for positive diagonal matrices D_1 and D_2 and $A_{ij} := \exp(-\eta C_{ij})$. Write d for the vector of diagonal entries of D_2 . For $1 \leq i \leq n - 1$, we have

$$\frac{d_i}{d_{i+1}} = \frac{A_{i,i}}{A_{i,i+1}} \frac{d_i}{d_{i+1}} = \frac{X_{i,i}^\eta}{X_{i,i+1}^\eta} > \frac{1 - \varepsilon}{\varepsilon}. \quad (9)$$

Finally, we note that

$$\frac{X_{n,n}^\eta}{X_{n,1}^\eta} = \frac{A_{n,n}}{A_{n,1}} \frac{d_n}{d_1} = \exp(\eta) \frac{d_n}{d_1},$$

and since $X_{n,1}^\eta < \varepsilon$, we obtain

$$\frac{d_n}{d_1} > \exp(-\eta) \frac{1 - \varepsilon}{\varepsilon}. \quad (10)$$

Combining (9) for $1 \leq i \leq n - 1$ with (10) yields

$$0 > n \log \frac{1 - \varepsilon}{\varepsilon} - \eta,$$

which implies

$$\eta > n \log \frac{1 - \varepsilon}{\varepsilon},$$

as claimed. \blacksquare

5. Conclusions

Our focus in this work has been on making explicit the asymptotic analysis of [Cominetti and San Martín \(1994\)](#). Their paper has been cited consistently in the computational optimal transport community as giving the best account of the speed of convergence of the penalized program to the original linear program (see [Genevay et al., 2016](#); [Benamou et al., 2015, 2016](#); [Blondel et al., 2017](#); [Carlier et al., 2017](#); [Denitju et al., 2014](#); [Dessein et al., 2016](#); [Di Marino et al., 2017](#); [Díaz et al., 2015](#); [Genevay et al., 2016](#); [Schmitzer, 2016](#); [Peyré and Cuturi, 2017](#); [Luise et al., 2018](#)). We hope that the simple and explicit proof here will clarify the nature of the exponential rate proved in [Cominetti and San Martín \(1994\)](#), and provide a framework for a more refined analysis of the entropic penalty for linear programs of interest.

One puzzle that remains is to give theoretical justification to the observation of [Cuturi \(2013\)](#) that small values of η achieve good accuracy on real-world optimal transport data. It is clear that the analysis of [Theorem 5](#) could be improved via a more refined understanding of the “energy spectrum” of optimal transport (i.e., the size and structure of the set of nearly-optimal transports), but obtaining this understanding even in the case where the costs are i.i.d. random variables is a very deep question ([Aldous, 2001](#)). We leave obtaining a more sophisticated grasp on the behavior of this trajectory for future work.

6. Proofs of Lemmas

Proof [Proof of [Lemma 6](#)] Write $\bar{\lambda} = 1 - \lambda$, and let $R = \max\{\|x\|_1, \|y\|_1\}$.

$$\begin{aligned}
 H(\lambda x + \bar{\lambda} y) &= \sum_i (\lambda x_i + \bar{\lambda} y_i) \left(\log \frac{1}{\lambda x_i + \bar{\lambda} y_i} \right) \\
 &= \sum_i \lambda x_i \left(\log \frac{1}{\lambda x_i + \bar{\lambda} y_i} \right) + \sum_i \bar{\lambda} y_i \left(\log \frac{1}{\lambda x_i + \bar{\lambda} y_i} \right) \\
 &\leq \sum_i \lambda x_i \left(\log \frac{1}{\lambda x_i} \right) + \sum_i \bar{\lambda} y_i \left(\log \frac{1}{\bar{\lambda} y_i} \right) \\
 &= \lambda H(x) + \bar{\lambda} H(y) + \lambda \log \frac{1}{\lambda} \sum_i x_i + \bar{\lambda} \log \frac{1}{\bar{\lambda}} \sum_i y_i \\
 &\leq \lambda H(x) + \bar{\lambda} H(y) + Rh(\lambda).
 \end{aligned}$$

■

Proof [Proof of [Lemma 7](#)] The derivative of f satisfies

$$f'(\lambda) = \alpha \log \frac{1 - \lambda}{\lambda} + \beta.$$

When $\lambda \in (0, \frac{\beta}{\alpha + \beta}]$,

$$\log \frac{1 - \lambda}{\lambda} \geq -\log \frac{\beta}{\alpha} > -\frac{\beta}{\alpha},$$

so $f'(\lambda) > 0$. The claim follows.

■

Proof [Proof of Lemma 8] By definition $h(\rho) = \rho \log \frac{1}{\rho} + (1 - \rho) \log \frac{1}{1-\rho}$, so it suffices to show that

$$\frac{(1 - \rho)}{\rho} \log \frac{1}{1 - \rho} \leq 1 \quad \forall \rho \in [0, 1].$$

This inequality is easily verified by noting that the derivative of the left side is nonpositive on $(0, 1)$ and $\lim_{\rho \rightarrow 0^+} \frac{(1-\rho)}{\rho} \log \frac{1}{1-\rho} = 1$. ■

Proof [Proof of Lemma 15] It is trivial to see that all $X \in \mathcal{B}$ satisfy $\sum_{ij} X_{ij} = n$, so $R_1 = n$. For any $X \in \mathcal{B}$,

$$\begin{aligned} H(X) &= \sum_{ij} X_{ij} \log \frac{1}{X_{ij}} \\ &= \sum_i \sum_j X_{ij} \log \frac{1}{X_{ij}} \\ &= \sum_i H(X_i), \end{aligned}$$

where X_i denotes the i th row of X . Since each row of X is a nonnegative vector of dimension n whose entries sum to 1, for each $1 \leq i \leq n$ the bound $0 \leq H(X_i) \leq \log n$ holds. Therefore $0 \leq H(X) \leq n \log n$ for all $X \in \mathcal{B}$, which proves that $R_H \leq n \log n$. ■

Acknowledgments

This work was supported in part by NSF Graduate Research Fellowship DGE-1122374. The author would like to thank J. Altschuler and P. Rigollet for useful discussions, as well as the anonymous referees for their suggestions.

References

- David J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Random Structures Algorithms*, 18(4):381–418, 2001. ISSN 1042-9832. URL <https://doi.org/10.1002/rsa.1015>.
- Zeyuan Allen-Zhu, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson. Much faster algorithms for matrix scaling. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 890–901. IEEE Computer Society, 2017. ISBN 978-1-5386-3464-6. doi: 10.1109/FOCS.2017.87. URL <https://doi.org/10.1109/FOCS.2017.87>.
- Jason Altschuler, Jonathan Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 1961–1971, 2017. URL <http://papers.nips.cc/paper/6792-near-linear-time-approximation-algorithms-for-optimal-transport-via-sinkhorn-iteration>.

- Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2013.
- Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- Jean-David Benamou, Guillaume Carlier, and Luca Nenna. A numerical method to solve multi-marginal optimal transport problems with Coulomb cost. In *Splitting Methods in Communication, Imaging, Science, and Engineering*, Sci. Comput., pages 577–601. Springer, Cham, 2016.
- Mathieu Blondel, Vivien Seguy, and Antoine Rolet. Smooth and sparse optimal transport. *arXiv preprint arXiv:1710.06276*, 2017.
- Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004. ISBN 0-521-83378-7.
- L. M. Brègman. A relaxation method of finding a common point of convex sets and its application to the solution of problems in convex programming. *Ž. Vyčisl. Mat. i Mat. Fiz.*, 7:620–631, 1967. ISSN 0044-4669.
- Richard A. Brualdi. *Combinatorial Matrix Classes*, volume 108 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2006. ISBN 978-0-521-86565-4; 0-521-86565-4. doi: 10.1017/CBO9780511721182. URL <http://dx.doi.org/10.1017/CBO9780511721182>.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3-4):231–357, 2015. doi: 10.1561/22000000050. URL <https://doi.org/10.1561/22000000050>.
- Guillaume Carlier, Vincent Duval, Gabriel Peyré, and Bernhard Schmitzer. Convergence of entropic schemes for optimal transport and gradient flows. *SIAM Journal on Mathematical Analysis*, 49(2):1385–1418, 2017.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, 2006. ISBN 978-0-521-84108-5; 0-521-84108-9. doi: 10.1017/CBO9780511546921. URL <http://dx.doi.org/10.1017/CBO9780511546921>.
- Stéphane Chrétien and Alfred O. Hero, III. Kullback proximal algorithms for maximum-likelihood estimation. *IEEE Trans. Inform. Theory*, 46(5):1800–1810, 2000. ISSN 0018-9448. doi: 10.1109/18.857792. URL <http://dx.doi.org/10.1109/18.857792>.
- Michael B. Cohen, Aleksander Mądry, Dimitris Tsipras, and Adrian Vladu. Matrix scaling and balancing via box constrained newton’s method and interior point methods. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 902–913. IEEE Computer Society, 2017. ISBN 978-1-5386-3464-6. doi: 10.1109/FOCS.2017.88. URL <https://doi.org/10.1109/FOCS.2017.88>.

- R. Cominetti and J. San Martín. Asymptotic analysis of the exponential penalty trajectory in linear programming. *Math. Programming*, 67(2, Ser. A):169–187, 1994. ISSN 0025-5610. doi: 10.1007/BF01582220. URL <http://dx.doi.org/10.1007/BF01582220>.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *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.*, pages 2292–2300, 2013. URL <http://papers.nips.cc/paper/4927-sinkhorn-distances-lightspeed-computation-of-optimal-transport>.
- Marco Cuturi and Arnaud Doucet. Fast computation of Wasserstein barycenters. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 685–693. JMLR.org, 2014. URL <http://jmlr.org/proceedings/papers/v32/cuturi14.html>.
- Andreea Denitui, Stefania Petra, Claudius Schnörr, and Christoph Schnörr. An entropic perturbation approach to TV-minimization for limited-data tomography. In Elena Barucci, Andrea Frosini, and Simone Rinaldi, editors, *Discrete Geometry for Computer Imagery - 18th IAPR International Conference, DGCI 2014, Siena, Italy, September 10-12, 2014. Proceedings*, volume 8668 of *Lecture Notes in Computer Science*, pages 262–274. Springer, 2014. ISBN 978-3-319-09954-5. doi: 10.1007/978-3-319-09955-2_22. URL https://doi.org/10.1007/978-3-319-09955-2_22.
- Arnaud Dessein, Nicolas Papadakis, and Jean-Luc Rouas. Regularized optimal transport and the ROT mover’s distance. *arXiv preprint arXiv:1610.06447*, 2016.
- Simone Di Marino, Augusto Gerolin, and Luca Nenna. Optimal transportation theory with repulsive costs. In *Topological optimization and optimal transport*, volume 17 of *Radon Ser. Comput. Appl. Math.*, pages 204–256. De Gruyter, Berlin, 2017.
- Juan Díaz, Tomás Rau, and Jorge Rivera. A matching estimator based on a bilevel optimization problem. *Review of Economics and Statistics*, 97(4):803–812, 2015.
- S. C. Fang. An unconstrained convex programming view of linear programming. *Z. Oper. Res.*, 36(2):149–161, 1992. ISSN 0340-9422. URL <https://doi.org/10.1007/BF01417214>.
- S.-C. Fang, J. R. Rajasekera, and H.-S. J. Tsao. *Entropy optimization and mathematical programming*, volume 8 of *International Series in Operations Research & Management Science*. Kluwer Academic Publishers, Boston, MA, 1997. ISBN 0-7923-9939-0. URL <https://doi.org/10.1007/978-1-4615-6131-6>.
- Shu-Cherng Fang and H-S Jacob Tsao. Linear programming with entropic perturbation. *Zeitschrift für Operations Research*, 37(2):171–186, 1993.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.

- Aude Genevay, Marco Cuturi, Gabriel Peyré, and Francis R. Bach. Stochastic optimization for large-scale optimal transport. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 3432–3440, 2016. URL <http://papers.nips.cc/paper/6566-stochastic-optimization-for-large-scale-optimal-transport>.
- David P Helmbold and Manfred K Warmuth. Learning permutations with exponential weights. *Journal of Machine Learning Research*, 10(Jul):1705–1736, 2009.
- Edwin T Jaynes. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*, 70(9):939–952, 1982.
- Anatoli Juditsky, Philippe Rigollet, and Alexandre Tsybakov. Learning by mirror averaging. *Ann. Statist.*, 36(5):2183–2206, 2008. ISSN 0090-5364.
- Wouter M. Koolen, Manfred K. Warmuth, and Jyrki Kivinen. Hedging structured concepts. In Adam Tauman Kalai and Mehryar Mohri, editors, *COLT 2010 - The 23rd Conference on Learning Theory, Haifa, Israel, June 27-29, 2010*, pages 93–105. Omnipress, 2010. ISBN 978-0-9822529-2-5. URL <http://colt2010.haifa.il.ibm.com/papers/COLT2010proceedings.pdf#page=101>.
- JJ Kosowsky and Alan L Yuille. The invisible hand algorithm: Solving the assignment problem with statistical physics. *Neural Networks*, 7(3):477–490, 1994.
- Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete Contin. Dyn. Syst.*, 34(4):1533–1574, 2014. ISSN 1078-0947. doi: 10.3934/dcds.2014.34.1533. URL <http://dx.doi.org/10.3934/dcds.2014.34.1533>.
- Giulia Luise, Alessandro Rudi, Massimiliano Pontil, and Carlo Ciliberto. Differential properties of Sinkhorn approximation for learning with Wasserstein distance. *arXiv preprint arXiv:1805.11897*, 2018.
- Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford Graduate Texts. Oxford University Press, Oxford, 2009. ISBN 978-0-19-857083-7. doi: 10.1093/acprof:oso/9780198570837.001.0001. URL <https://doi.org/10.1093/acprof:oso/9780198570837.001.0001>.
- Aleksander Mądry. Navigating central path with electrical flows: From flows to matchings, and back. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 253–262. IEEE Computer Society, 2013. ISBN 978-0-7695-5135-7. doi: 10.1109/FOCS.2013.35. URL <https://doi.org/10.1109/FOCS.2013.35>.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport. Book draft, 2017.
- Philippe Rigollet. Personal communication, 2017.

- Philippe Rigollet and Alexandre Tsybakov. Exponential screening and optimal rates of sparse estimation. *Ann. Statist.*, 39(2):731–771, 2011. ISSN 0090-5364. doi: 10.1214/10-AOS854. URL <http://dx.doi.org/10.1214/10-AOS854>.
- Bernhard Schmitzer. Stabilized sparse scaling algorithms for entropy regularized transport problems. *arXiv preprint arXiv:1610.06519*, 2016.
- Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. A*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. ISBN 3-540-44389-4.
- Erwin Schrödinger. Über die Umkehrung der Naturgesetze. *Angewandte Chemie*, 44(30):636–636, 1931.
- Meisam Sharify, Stéphane Gaubert, and Laura Grigori. Solution of the optimal assignment problem by diagonal scaling algorithms. *arXiv preprint arXiv:1104.3830*, 2011.
- Richard Sinkhorn. Diagonal equivalence to matrices with prescribed row and column sums. *The American Mathematical Monthly*, 74(4):402–405, 1967.
- Justin Solomon, Fernando De Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen, Tao Du, and Leonidas Guibas. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (TOG)*, 34(4):66, 2015.