

On the Complexity of Approximating Multimarginal Optimal Transport

Tianyi Lin*

*Department of Electrical Engineering and Computer Science
University of California
Berkeley, CA 94720-1776, USA*

DARREN_LIN@BERKELEY.EDU

Nhat Ho*

*Department of Statistics and Data Sciences
University of Texas, Austin
Texas, 78712-1823, USA*

MINHNHAT@UTEXAS.EDU

Marco Cuturi

*Google Brain
Department of Statistics, CREST - ENSAE*

CUTURI@GOOGLE.COM

Michael I. Jordan

*Department of Electrical Engineering and Computer Science and Department of Statistics
University of California
Berkeley, CA 94720-1776, USA*

JORDAN@CS.BERKELEY.EDU

Editor: Lorenzo Rosasco

Abstract

We study the complexity of approximating the multimarginal optimal transport (MOT) distance, a generalization of the classical optimal transport distance, considered here between m discrete probability distributions supported each on n support points. First, we show that the standard linear programming (LP) representation of the MOT problem is not a minimum-cost flow problem when $m \geq 3$. This negative result implies that some combinatorial algorithms, e.g., network simplex method, are not suitable for approximating the MOT problem, while the worst-case complexity bound for the deterministic interior-point algorithm remains a quantity of $\tilde{O}(n^{3m})$. We then propose two simple and *deterministic* algorithms for approximating the MOT problem. The first algorithm, which we refer to as *multimarginal Sinkhorn* algorithm, is a provably efficient multimarginal generalization of the Sinkhorn algorithm. We show that it achieves a complexity bound of $\tilde{O}(m^3 n^m \varepsilon^{-2})$ for a tolerance $\varepsilon \in (0, 1)$. This provides a first *near-linear time* complexity bound guarantee for approximating the MOT problem and matches the best known complexity bound for the Sinkhorn algorithm in the classical OT setting when $m = 2$. The second algorithm, which we refer to as *accelerated multimarginal Sinkhorn* algorithm, achieves the acceleration by incorporating an estimate sequence and the complexity bound is $\tilde{O}(m^3 n^{m+1/3} \varepsilon^{-4/3})$. This bound is better than that of the first algorithm in terms of $1/\varepsilon$, and accelerated alternating minimization algorithm (Tupitsa et al., 2020) in terms of n . Finally, we compare our new algorithms with the commercial LP solver GUROBI. Preliminary results on synthetic data and real images demonstrate the effectiveness and efficiency of our algorithms.

* Tianyi Lin and Nhat Ho contributed equally to this work.

Keywords: Multimarginal optimal transport, computational complexity, network flow, Multimarginal Sinkhorn algorithm, Accelerated multimarginal Sinkhorn algorithm.

1. Introduction

The multimarginal optimal transport (MOT) (Gangbo and Swiech, 1998; Pass, 2015), the general problem of aligning or correlating $m \geq 2$ probability measures so as to maximize efficiency (with respect to a given cost function), is a generalization of the optimal transport (OT) problem (Villani, 2003). From the Kantorovich formulation (Kantorovich, 1942), we seek to solve the following optimization problem,

$$\min_{\gamma \in \Pi(\mu_1, \mu_2, \dots, \mu_m)} \int_{M_1 \times M_2 \times \dots \times M_m} c(x_1, x_2, \dots, x_m) d\gamma(x_1, x_2, \dots, x_m), \quad (1)$$

where $\Pi(\mu_1, \mu_2, \dots, \mu_m)$ denotes the set of positive joint measures on the product space $M_1 \times M_2 \times \dots \times M_m$ whose marginals are $\{\mu_i\}_{i \in [m]}$, and $c(\cdot)$ is a given cost function. In the discrete setting where each of μ_i is supported on n support points, the MOT problem is equivalent to a linear programming (LP) problem with mn constraints and n^m variables, which means that any algorithm requires at least n^m arithmetic operations in general.

The MOT problem has been recognized as the backbone of numerous important applications, such as matching in economics (Ekeland, 2005; Carlier and Ekeland, 2010a,b), density functional theory in physics (Seidl et al., 2007; Buttazzo et al., 2012; Cotar et al., 2013; Mendl and Lin, 2013), generalized Euler flow in fluid dynamics (Brenier, 1989, 1999, 2008) and financial mathematics (Dolinsky and Soner, 2014; Galichon et al., 2014). Over the past five years, the MOT problem has begun to attract considerable attention, due in part to a wide variety of emerging applications in machine learning, including generative adversarial networks (GANs) (Choi et al., 2018; Cao et al., 2019), clustering (Mi and Bento, 2020), domain adaptation (Hui et al., 2018; He et al., 2019) and Wasserstein barycenters (Agueh and Carlier, 2011; Cuturi and Doucet, 2014; Benamou et al., 2015; Carlier et al., 2015; Srivastava et al., 2018). Due to the space limit, we refer the interested readers to Pass (2015) for other applications of the MOT problem and Peyré and Cuturi (2019) for more details of the MOT problem from a computational point of view.

In order to further motivate the MOT problem, we briefly describe two representative application problems arising from machine learning.

Example 1. *The multimarginal Wasserstein GANs (Cao et al., 2019) are proposed to optimize a feasible MOT distance among different domains. This approach is based on a new dual formulation of the MOT distance and overcomes the limitations of existing methods by alleviating the distribution mismatching issue and exploiting cross-domain correlations.*

We consider $m \geq 2$ target domains $\{\mathcal{D}_k\}_{k \in [m]}$ and the associated generative models g_k parameterized by θ_k for all $k \in [m]$. Let $\mathcal{F} = \{f : \mathbb{R}^d \rightarrow \mathbb{R}\}$ be the class of discriminators parameterized by w , we define the MOT distance in the dual form as follows,

$$W(\hat{\mathbb{P}}_s, \hat{\mathbb{P}}_{\theta_1}, \dots, \hat{\mathbb{P}}_{\theta_m}) = \max_f \mathbb{E}_{x \sim \hat{\mathbb{P}}_s} [f(x)] - \sum_{k=1}^m \lambda_k^+ \mathbb{E}_{x \sim \hat{\mathbb{P}}_{\theta_k}} [f(x)], \quad \text{s.t. } \hat{\mathbb{P}}_{\theta_k} \in \mathcal{D}_k, f \in \Omega,$$

where $\hat{\mathbb{P}}_s$ is the real source distribution, $\hat{\mathbb{P}}_{\theta_k}$ is the distribution generated by g_k for all $k \in [m]$, and $\Omega = \{f \in \mathcal{F} \mid f(x) - \sum_{k=1}^m \lambda_k^+ f(\hat{x}^{(k)}) \leq c(x, \hat{x}^{(1)}, \dots, \hat{x}^{(m)})\}$ where $x \in \hat{\mathbb{P}}_s$ and $\hat{x}^{(k)} \in \hat{\mathbb{P}}_{\theta_k}$

for all $k \in [m]$ are samples. Note that λ_k^+ reflects the importance of the k -th target domain and is set as $1/m$ in practice when no prior knowledge is available.

Example 2. The free-support Wasserstein barycenter (Agueh and Carlier, 2011) is defined as a weighted barycenter of input measures $\{\mu_k\}_{k \in [m]}$ defined on \mathbb{R}^d according to the OT distance. As shown by Agueh and Carlier (2011), the computation of barycenters of measures can be computed by solving a MOT problem.

We consider the discrete setting where input measures are $\mu_k = \sum_{i=1}^n p_{k,i} \delta_{x_i}$ with weights $p_k = (p_{k,1}, \dots, p_{k,n}) \in \Delta^n$, the support points $\{x_i\}_{i \in [n]} \subseteq \mathbb{R}^d$ and the Dirac measure δ . It is shown in (Agueh and Carlier, 2011) that the Wasserstein barycenter of $\{\mu_k\}_{k \in [m]}$ with weights $\lambda = (\lambda_1, \dots, \lambda_m) \in \Delta^m$ according to the OT distance with the quadratic Euclidean distance ground cost function $c = \|\cdot\|^2$ is

$$\mu_\lambda := \sum_{1 \leq i_k \leq n, \forall k \in [m]} \gamma_{i_1, \dots, i_m} \delta_{A_{i_1, \dots, i_m}(x)},$$

where $A_{i_1, \dots, i_m}(x) = \sum_{k=1}^m \lambda_k x_{i_k}$ is the Euclidean barycenter and $\gamma \in \mathbb{R}^{n \times \dots \times n}$ is an optimal multimarginal transportation plan that solves the MOT problem in the LP form of

$$\min_{X \in \mathbb{R}^{n \times \dots \times n}} \langle C, X \rangle, \quad \text{s.t.} \quad \sum_{1 \leq i_l \leq n, l \neq k, \forall l \in [m]} X_{i_1, \dots, i_{k-1}, j, i_{k+1}, \dots, i_m} = p_{kj} \text{ for all } (k, j) \in [m] \times [n],$$

where C is defined as $C_{i_1, \dots, i_m} = \sum_{k=1}^m (\lambda_k/2) \|x_{i_k} - A_{i_1, \dots, i_m}(x)\|^2$ for $(i_1, \dots, i_m) \in [n] \times \dots \times [n]$. In practice, we set $\lambda_k = 1/m$ for all $k \in [m]$ when no prior knowledge is available.

It is worthy noting that the barycenter μ_λ is in general composed of more than m Diracs, and that these Diracs are not constrained to be on the support points $\{x_i\}_{i \in [n]}$. This is different from the fixed-support Wasserstein barycenter that must be on the same support points $\{x_i\}_{i \in [n]}$ of the input measures. To be specific, the free-support Wasserstein barycenter is the “true” barycenter of measures, while the fixed-support Wasserstein barycenter is an approximation on the fixed support points. But, on the flip side of the coin, the fixed-support Wasserstein barycenter can be computed without solving any MOT problem and the complexity bound is polynomial in m , n and $1/\varepsilon$ (Kroshnin et al., 2019; Lin et al., 2020) where ε is the desired accuracy.

Algorithms for the OT problem. The OT problem is a special instance of the MOT problem with $m = 2$ and has been studied thoroughly during the past decade. To the best of our knowledge, there are mainly two group of algorithms for solving the OT problem.

The first line of algorithms are combinatorial graph algorithms (Klein, 1967; Edmonds and Karp, 1972; Hassin, 1983; Tardos, 1985; Galil and Tardos, 1988; Goldberg and Tarjan, 1990; Hassin, 1992; Ervolina and McCormick, 1993b,a; Orlin, 1993, 1997; Goldberg and Rao, 1998). Indeed, the OT problem is a minimum-cost flow problem (Schrijver, 2003), which has graph structure and leads to efficient combinatorial algorithms mentioned before. Examples include the primal-dual cost scaling algorithm (Goldberg and Tarjan, 1990) and the network simplex algorithm (Orlin, 1997); see also Daitch and Spielman (2008) and Lee and Sidford (2014) for some recent progresses.

The second line of algorithms, initialized with the Sinkhorn algorithm (Cuturi, 2013), are developed for solving the OT problem through either entropy regularization or continuous

optimization algorithmic frameworks. Examples include Greenkhorn algorithm (Altschuler et al., 2017; Lin et al., 2019a), accelerated first-order primal-dual algorithms (Dvurechensky et al., 2018b), accelerated Sinkhorn algorithms (Lin et al., 2019b; Guminov et al., 2019), and some other algorithms (Blanchet et al., 2018; Jambulapati et al., 2019; Lahn et al., 2019; Xie et al., 2020). Even though these algorithms are very efficient, with easy to implement routines in practice, the Sinkhorn algorithm and its accelerated variants remain as the state-of-the-art approach for the OT problem and serve as the default solver in the celebrated POT package (Flamary and Courty, 2017).

Algorithms for the MOT problem. While the theory for computing the OT distance has received ample attention, the theory for computing the MOT distance is still nascent. Since the MOT problem has the LP representation with mn constraints and n^m variables, it can be solved by many LP algorithms, e.g., the interior-point algorithm, whose complexity bounds are however not near-linear. That is to say, the dependence of n is linear in n^m up to the logarithmic factors.

Two specialized algorithms are available for solving the MOT problem: multimarginal Sinkhorn algorithm and accelerated alternating minimization algorithm. The former one generalizes the Sinkhorn algorithm to the MOT setting but only has the asymptotic convergence analysis (Benamou et al., 2015, 2019; Peyré and Cuturi, 2019); the latter one is proposed by the concurrent work (Tupitsa et al., 2020) for solving the same dual entropic regularized MOT problem as ours and achieves the complexity bound of $\tilde{O}(m^3 n^{m+1/2} \varepsilon^{-1})$ when applied to solve the MOT problem along with our rounding scheme. However, their algorithm is not a near-linear time approximation algorithm and the dependence of n can be potentially improved.

Contribution: In this paper, we study the complexity of approximating the MOT problem between m discrete probability distributions with at most n points in their respective supports. Our contributions can be summarized as follows:

1. We show that the standard LP representation of the MOT problem is not a minimum-cost flow problem when $m \geq 3$. This implies the inefficiency of many combinatorial algorithms, including network simplex method, as well as the worst-case complexity bound of $\tilde{O}(n^{3m})$ for the standard *deterministic* interior-point algorithms.
2. We propose two simple and *deterministic* algorithms for solving the entropic regularized MOT problem. The first one is named as *multimarginal Sinkhorn algorithm* which can be also used to solve the MOT problem along with a new rounding scheme. The achieved complexity bound is $\tilde{O}(m^3 n^m \varepsilon^{-2})$, which is near-linear in terms of n^m , demonstrating that our algorithm is unimprovable in terms of n in general setting. To the best of our knowledge, this is a first *near-linear time* approximation algorithm for solving the MOT problem while the existing ones are either only shown convergent (Benamou et al., 2015, 2019; Peyré and Cuturi, 2019) or not near-linear time (Tupitsa et al., 2020). The second one is named as *accelerated multimarginal Sinkhorn algorithm* and achieves the complexity bound of $\tilde{O}(m^3 n^{m+1/3} \varepsilon^{-4/3})$ when applied to solve the MOT problem. This complexity bound is better than that of the first algorithm in terms of $1/\varepsilon$, and the accelerated alternating minimization algorithm (Tupitsa et al., 2020) in terms of n .

3. We compare our algorithms with the commercial LP solver GUROBI. Preliminary results on both synthetic data and real images demonstrate the effectiveness and efficiency of our algorithms in practice.

Organization. The remainder of the paper is organized as follows. In Section 2, we present the background materials on the MOT problem and derive some important properties of the objective function in the dual entropic regularized MOT problem. In Section 3, we show that the standard LP representation of the MOT problem is not a minimum-cost flow problem when $m \geq 3$. In Sections 4 and 5, we propose the multimarginal Sinkhorn and accelerated multimarginal Sinkhorn algorithms for solving the entropic regularized MOT problem. We also demonstrate that these algorithms can solve the MOT problem efficiently along with our new rounding scheme. In Section 6, we present some numerical results which validate the efficiency of our algorithms. We finally conclude this paper in Section 7.

Notation. We let $[n]$ be the set $\{1, 2, \dots, n\}$ and \mathbb{R}_+^n be the set of all vectors in \mathbb{R}^n with non-negative components. $\mathbf{1}_n \in \mathbb{R}^n$ refers to a vector with all of its components are 1 and Δ^n is denoted as the probability simplex in \mathbb{R}_+^n : $\Delta^n = \{u \in \mathbb{R}_+^n : \mathbf{1}_n^\top u = 1\}$. For a set S , we denote $|S|$ as its cardinality. For a differentiable function f , we denote ∇f and $\nabla_\beta f$ as the full gradient of f and the gradient of f with respect to β . For a vector $x \in \mathbb{R}^n$ and $1 \leq p \leq \infty$, we denote $\|x\|_p$ as its ℓ_p -norm and $\|x\|$ as its ℓ_2 -norm for simplicity. Let x and y be two vectors of same dimension, we denote $\min\{x, y\}$ as the component-wise minimum of x and y . For a tensor $A = (A_{i_1, \dots, i_m}) \in \mathbb{R}^{n_1 \times \dots \times n_m}$, we write $\|A\|_\infty = \max_{1 \leq i_k \leq n_k, \forall k \in [m]} |A_{i_1, \dots, i_m}|$ and $\|A\|_1 = \sum_{1 \leq i_k \leq n_k, \forall k \in [m]} |A_{i_1, \dots, i_m}|$, and denote $r_k(A) \in \mathbb{R}^{n_k}$ as its k -th marginal for $k \in [m]$ and each component is defined by

$$[r_k(A)]_j := \sum_{1 \leq i_l \leq n_l, \forall l \neq k} A_{i_1, \dots, i_{k-1}, j, i_{k+1}, \dots, i_m}.$$

Let A and B be two tensors of same dimension, we denote their Frobenius inner product as

$$\langle A, B \rangle := \sum_{1 \leq i_k \leq n_k, \forall k \in [m]} A_{i_1, \dots, i_m} B_{i_1, \dots, i_m}.$$

Given the dimension n and accuracy ε , the notation $a = O(b(n, \varepsilon))$ stands for the upper bound $a \leq C \cdot b(n, \varepsilon)$ where $C > 0$ is independent of n and ε , and the notation $a = \tilde{O}(b(n, \varepsilon))$ indicates the previous inequality where C depends on the logarithmic function of n and ε .

2. Preliminaries

In this section, we first present the linear programming (LP) representation of the multimarginal optimal transport (MOT) problem as well as a formal specification of an approximate multimarginal transportation plan. Then, we describe the entropic regularized MOT problem and derive the dual entropic regularized MOT problem where the objective function is in the form of the logarithm of sum of exponents. Finally, we provide several properties of this function which are useful for the subsequent analysis.

2.1 Linear programming representation

The linear programming representation of the OT problem between two discrete probability distributions with n supports dates back to the seminar work by Kantorovich (1942), and can be written as

$$\min_{X \in \mathbb{R}^{n \times n}} \langle C, X \rangle \quad \text{s.t. } X \mathbf{1}_n = r, \quad X^\top \mathbf{1}_n = c, \quad X \geq 0.$$

In the above formulation, $X \in \mathbb{R}_+^{n \times n}$ denotes a *transportation plan*, $C \in \mathbb{R}_+^{n \times n}$ denotes a *nonnegative cost matrix*, and r and c stand for two probability distributions lying in the simplex Δ^n . Approximately solving the OT problem amounts to finding an ε -approximate transportation plan \hat{X} such that $\hat{X} \mathbf{1}_n = r$, $\hat{X}^\top \mathbf{1}_n = c$ and the following inequality holds true,

$$\langle C, \hat{X} \rangle \leq \langle C, X^* \rangle + \varepsilon.$$

where X^* is defined as an optimal transportation plan of the OT problem.

As a straightforward generalization of the OT problem, the MOT problem is also a LP. Indeed, the problem of computing the MOT distance between $m \geq 2$ discrete probability distributions with n supports is in the following form of

$$\min_{X \in \mathbb{R}^{n \times \dots \times n}} \langle C, X \rangle, \quad \text{s.t. } r_k(X) = r_k \text{ for any } k \in [m], \quad X \geq 0. \quad (2)$$

In the above formulation, X denotes the *multimarginal transportation plan*, $C \in \mathbb{R}_+^{n \times \dots \times n}$ denotes a *nonnegative cost tensor*, and $\{r_k\}_{k \in [m]}$ stand for a set of probability distributions all lying in Δ^n .

We see from Eq. (2), that the MOT problem is a linear programming with mn equality constraints and n^m variables. The solution we hope to achieve is an ε -approximate multimarginal transportation plan which generalizes the notion of an ε -approximate transportation plan of the OT problem. More specifically, we have the following definition of ε -approximate multimarginal transportation plan.

Definition 1. *The nonnegative tensor $\hat{X} \in \mathbb{R}_+^{n \times \dots \times n}$ is called an ε -approximate multimarginal transportation plan if $r_k(\hat{X}) = r_k$ for any $k \in [m]$ and the following inequality holds true,*

$$\langle C, \hat{X} \rangle \leq \langle C, X^* \rangle + \varepsilon,$$

where X^* is defined as an optimal multimarginal transportation plan of the MOT problem.

With this definition in mind, one of the goals of this paper is to develop *near-linear time approximation* algorithms for solving the MOT problem. In particular, we seek the algorithms whose running time required to obtain an ε -approximate multimarginal transportation plan is nearly linear in the number of unknown variables n^m . These algorithms are favorable in modern machine learning applications since they are unimprovable up to the logarithmic factors in general. Indeed, for the general MOT problem, the tensor $X \in \mathbb{R}_+^{n \times \dots \times n}$ has n^m unknown entries. In order to solve the MOT problem, the number of arithmetic operations required by any algorithms is at least n^m .

In the classical OT setting, Altschuler et al. (2017) has shown that the Sinkhorn algorithm is near-linear time approximation algorithm. Benamou et al. (2015, 2019) generalized

the Sinkhorn algorithm to the MOT setting but did not provide any complexity bound guarantee for their algorithms. Thus, it is still unclear whether there exists a near-linear time approximation algorithm for the general MOT problem.

2.2 Entropic regularized MOT and its dual form

Building on Cuturi’s entropic approach to the classical OT problem (Cuturi, 2013), we consider a regularized version of the MOT problem in which we add an entropic penalty function to the objective in Eq. (2). The resulting problem is in the following form:

$$\begin{aligned} \min_{X \in \mathbb{R}^{n \times \dots \times n}} \quad & \langle C, X \rangle - \eta H(X) \\ \text{s.t.} \quad & r_k(X) = r_k \text{ for any } k \in [m], X \geq 0, \end{aligned} \tag{3}$$

where $\eta > 0$ denotes the regularization parameter and $H(X)$ denotes the entropic regularization term, which is given by:

$$H(X) := -\langle X, \log(X) - \mathbf{1}_{n \times \dots \times n} \rangle.$$

It is important to note that if η is large, the resulting optimal value of the entropic regularized MOT problem (cf. Eq (3)) yields a poor approximation to the unregularized MOT problem. Moreover, another issue of entropic regularization is that the sparsity of the solution is lost. Even though an ε -approximate transportation plan can be found efficiently, it is not clear how different the resulting sparsity pattern of the obtained solution is with respect to the solution of the actual OT problem. In contrast, as a special instance of the MOT distance, the actual OT distance suffers from the curse of dimensionality (Dudley, 1969; Fournier and Guillin, 2015; Weed and Bach, 2019; Lei, 2020) and is significantly worse than its entropic regularized version in terms of the sample complexity (Genevay et al., 2019; Mena and Niles-Weed, 2019). This statistical drawback also holds true for the unregularized MOT distance in general.

While there is an ongoing debate in the literature on the merits of solving the actual OT problem *versus* its entropic regularized version, we adopt here the viewpoint that reaching an additive approximation of the actual MOT cost matters and therefore propose to scale η as a function of the desired accuracy of the approximation.

Then we proceed to derive the dual form of the entropic regularized MOT problem in Eq. (3). As in the usual 2-marginals OT case (Cuturi and Peyré, 2018), the dual form of the MOT problem with $m \geq 3$ remains an unconstrained smooth optimization problem.

By introducing the dual variables $\{\lambda_1, \dots, \lambda_m\} \subseteq \mathbb{R}^n$ and $\tau \in \mathbb{R}$, we can define the Lagrangian function of the entropic regularized MOT problem in Eq. (3) as follows:

$$\mathcal{L}(X, \lambda_1, \dots, \lambda_m) = \langle C, X \rangle - \eta H(X) - \sum_{k=1}^m \lambda_k^\top (r_k(X) - r_k). \tag{4}$$

Note that the entropy function $H(X)$ is not well defined for any negative matrix X . Thus, we can neglect the non-negative constraint $X \geq 0$ and define the above function \mathcal{L} whose domain is $\mathbb{R}_+^{n \times \dots \times n} \times \mathbb{R}^{nm}$. In order to derive the smooth dual objective function, we consider

the following minimization problem:

$$\min_{X: \|X\|_1=1} \langle C, X \rangle - \eta H(X) - \sum_{k=1}^m \lambda_k^\top (r_k(X) - r_k).$$

In the above problem, the objective function is strongly convex. Thus, the optimal solution is unique. After the simple calculations, the optimal solution $\bar{X} = X(\lambda_1, \dots, \lambda_m)$ has the following form:

$$\bar{X}_{i_1 \dots i_m} = \frac{e^{\eta^{-1}(\sum_{k=1}^m \lambda_{ki_k} - C_{i_1 i_2 \dots i_m})}}{\sum_{1 \leq i_k \leq n, \forall k \in [m]} e^{\eta^{-1}(\sum_{k=1}^m \lambda_{ki_k} - C_{i_1 i_2 \dots i_m})}}. \quad (5)$$

Plugging Eq. (5) into Eq. (4) yields that the dual form is:

$$\max_{\lambda_1, \dots, \lambda_m} \left\{ -\eta \log \left(\sum_{1 \leq i_1, \dots, i_m \leq n} e^{\eta^{-1}(\sum_{k=1}^m \lambda_{ki_k} - C_{i_1 i_2 \dots i_m})} \right) + \sum_{k=1}^m \lambda_k^\top r_k \right\}.$$

In order to streamline our subsequent presentation, we perform a change of variables, $\beta_k = \eta^{-1} \lambda_k$, and reformulate the above problem as

$$\min_{\beta_1, \dots, \beta_m} \varphi(\beta_1, \dots, \beta_m) := \log \left(\sum_{1 \leq i_1, i_2, \dots, i_m \leq n} e^{\sum_{k=1}^m \beta_{ki_k} - \frac{C_{i_1 i_2 \dots i_m}}{\eta}} \right) - \sum_{k=1}^m \beta_k^\top r_k.$$

To further simplify the notation, we define $B(\beta) := (B_{i_1 \dots i_m})_{i_1, i_2, \dots, i_m \in [n]} \in \mathbb{R}^{n \times \dots \times n}$ where $\beta = (\beta_1, \dots, \beta_m)$ by

$$B_{i_1 \dots i_m} = e^{\sum_{k=1}^m \beta_{ki_k} - \frac{C_{i_1 i_2 \dots i_m}}{\eta}}.$$

To this end, we obtain the *dual entropic regularized MOT problem* defined by

$$\min_{\beta_1, \dots, \beta_m} \varphi(\beta_1, \dots, \beta_m) := \log(\|B(\beta_1, \dots, \beta_m)\|_1) - \sum_{k=1}^m \beta_k^\top r_k. \quad (6)$$

Remark 2. *The first part of the objective function φ is in the form of the logarithm of sum of exponents while the second part is a linear function. This is different from the objective function used in previous dual entropic regularized OT problem (Cuturi, 2013; Altschuler et al., 2017; Dvurechensky et al., 2018b; Lin et al., 2019a). We also note that Eq. (6) is a special instance of a softmax minimization problem, and the objective function φ is known to be smooth (Nesterov, 2005). Finally, we point out that the same problem was derived in the later work by Tupitsa et al. (2020) and used for analyzing the accelerated alternating minimization algorithm.*

In the remainder of the paper, we also denote $\beta^* = (\beta_1^*, \dots, \beta_m^*) \in \mathbb{R}^{mn}$ as an optimal solution of the dual entropic regularized MOT problem in Eq. (6).

2.3 Properties of dual entropic regularized multimarginal OT

In this section, we present several useful properties of the dual entropic regularized MOT in Eq. (6). In particular, we show that there exists an optimal solution β^* such that it has an upper bound in terms of the ℓ_∞ -norm.

Lemma 3. *For the dual entropic regularized MOT problem in Eq. (6), there exists an optimal solution $\beta^* = (\beta_1^*, \dots, \beta_m^*)$ such that*

$$\|\beta^*\|_\infty := \max_{1 \leq i \leq m} \|\beta_i^*\|_\infty \leq R, \quad (7)$$

where $R > 0$ is defined as

$$R := \frac{\|C\|_\infty}{\eta} - \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij} \right).$$

Proof First, we claim that there exists an optimal solution $\beta^* = (\beta_1^*, \dots, \beta_m^*)$ such that

$$\min_{1 \leq j \leq n} \beta_{ij}^* \leq 0 \leq \max_{1 \leq j \leq n} \beta_{ij}^* \text{ for any } i \in [m]. \quad (8)$$

Indeed, letting $\widehat{\beta}^* = (\widehat{\beta}_1^*, \dots, \widehat{\beta}_m^*)$ be an optimal solution to Eq. (6), the claim holds true if $\widehat{\beta}^*$ satisfies Eq. (8). Otherwise, we let m shift terms be

$$\Delta \widehat{\beta}_i = \frac{\max_{1 \leq j \leq n} \widehat{\beta}_{ij}^* + \min_{1 \leq j \leq n} \widehat{\beta}_{ij}^*}{2} \in \mathbb{R} \text{ for any } i \in [m].$$

and define $\beta^* = (\beta_1^*, \dots, \beta_m^*)$ by

$$\beta_i^* = \widehat{\beta}_i^* - \Delta \widehat{\beta}_i \mathbf{1}_n \text{ for any } i \in [m].$$

By the definition of β^* , it is clear that β^* satisfies Eq. (8). Since $\mathbf{1}_n^\top r_i = 1$ for all $i \in [m]$, we have $(\beta_i^*)^\top r_i = (\widehat{\beta}_i^*)^\top r_i - \Delta \widehat{\beta}_i$ for all $i \in [m]$. In addition, we have $\log(\|B(\beta_1^*, \dots, \beta_m^*)\|_1) = \log(\|B(\widehat{\beta}_1^*, \dots, \widehat{\beta}_m^*)\|_1) + \sum_{i=1}^m \Delta \widehat{\beta}_i$. Putting these pieces together yields $\varphi(\beta^*) = \varphi(\widehat{\beta}^*)$. Therefore, β^* is an optimal solution that satisfies Eq. (8).

Then, we show that

$$\max_{1 \leq j \leq n} \beta_{ij}^* - \min_{1 \leq j \leq n} \beta_{ij}^* \leq \frac{\|C\|_\infty}{\eta} - \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij} \right) \text{ for all } i \in [m]. \quad (9)$$

Indeed, for any $(j, l) \in [m] \times [n]$, we derive from the optimality condition of β^* that

$$\frac{e^{\beta_{jl}^*} \sum_{1 \leq i_k \leq n, \forall k \neq j} e^{\sum_{k \neq j} \beta_{ki_k}^* - \eta^{-1} C_{i_1 \dots l \dots i_m}}}{\|B(\beta_1^*, \dots, \beta_m^*)\|_1} = r_{jl} \geq \min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij}.$$

Since C is a nonnegative cost tensor, we have

$$\beta_{jl}^* \geq \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij} \right) - \log \left(\sum_{1 \leq i_k \leq n, \forall k \neq j} e^{\sum_{k \neq j} \beta_{ki_k}^*} \right) + \log(\|B(\beta_1^*, \dots, \beta_m^*)\|_1). \quad (10)$$

Since $r_{jl} \in [0, 1]$ and $C_{i_1 \dots i_m} \leq \|C\|_\infty$, we have

$$\beta_{jl}^* \leq \frac{\|C\|_\infty}{\eta} - \log \left(\sum_{1 \leq i_k \leq n, \forall k \neq j} e^{\sum_{k \neq j} \beta_{ki}^*} \right) + \log(\|B(\beta_1^*, \dots, \beta_m^*)\|_1). \quad (11)$$

Combining the bounds in Eq. (10) and Eq. (11) implies the desired Eq. (9).

Finally, we prove that Eq. (7) holds true. Indeed, Eq. (8) and Eq. (9) imply that

$$-\frac{\|C\|_\infty}{\eta} + \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij} \right) \leq \min_{1 \leq j \leq n} \beta_{ij}^* \leq 0 \text{ for any } i \in [m], \quad (12)$$

and

$$0 \leq \max_{1 \leq j \leq n} \beta_{ij}^* \leq \frac{\|C\|_\infty}{\eta} - \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij} \right) \text{ for any } i \in [m]. \quad (13)$$

Combining Eq. (12) and Eq. (13) with the definition of R implies that $\max_{1 \leq i \leq m} \|\beta_i^*\|_\infty \leq R$ and hence the desired Eq. (7). \blacksquare

The upper bound for the ℓ_∞ -norm of an optimal solution of dual entropic-regularized multimarginal OT in Lemma 3 directly leads to the following direct bound for the ℓ_2 -norm.

Corollary 4. *For the dual entropic regularized MOT problem in Eq. (6), there exists an optimal solution $\beta^* = (\beta_1^*, \dots, \beta_m^*)$ such that*

$$\|\beta^*\| \leq \sqrt{mn}R,$$

where $R > 0$ is defined in Lemma 3.

Since the function $-H(X)$ is strongly convex with respect to the ℓ_1 -norm on the probability simplex $Q \subseteq \mathbb{R}^{nm}$, the entropic regularized MOT problem in Eq. (3) is a special case of the following linearly constrained convex optimization problem:

$$\min_{x \in Q} f(x), \quad \text{s.t. } Ax = b,$$

where f is strongly convex with respect to the ℓ_1 -norm on the set Q :

$$f(x') - f(x) - (x' - x)^\top \nabla f(x) \geq \frac{\eta}{2} \|x' - x\|_1^2 \text{ for any } x', x \in Q.$$

We use the ℓ_2 -norm for the dual space of the Lagrange multipliers. By Nesterov (2005, Theorem 1), the dual objective function $\tilde{\varphi}$ satisfies the following inequality:

$$\tilde{\varphi}(\lambda') - \tilde{\varphi}(\lambda) - (\lambda' - \lambda)^\top \nabla \tilde{\varphi}(\lambda) \leq \frac{\|A\|_{1 \rightarrow 2}^2}{2\eta} \|\lambda' - \lambda\|^2 \text{ for any } \lambda', \lambda \in \mathbb{R}^{mn}.$$

Recall that the function $\tilde{\varphi}$ is given by

$$\tilde{\varphi}(\lambda) = -\eta \log \left(\sum_{1 \leq i_1, \dots, i_m \leq n} e^{\eta^{-1}(\sum_{k=1}^m \lambda_{ki_k} - C_{i_1 i_2 \dots i_m})} \right) + \sum_{k=1}^m \lambda_k^\top r_k.$$

We notice that the function φ in Eq. (6) is defined by

$$\varphi(\beta) = -\eta^{-1}\tilde{\varphi}(\eta(\beta + (1/m)\mathbf{1}_{mn})).$$

After some calculations, we have

$$\varphi(\beta') - \varphi(\beta) - (\beta' - \beta)^\top \nabla \varphi(\beta) \leq \left(\frac{\|A\|_{1 \rightarrow 2}^2}{2} \right) \|\beta' - \beta\|^2. \quad (14)$$

By definition, each column of the matrix A contains no more than m nonzero elements which are equal to one. Since $\|A\|_{1 \rightarrow 2}$ is equal to maximum ℓ_2 -norm of the column of this matrix, we have $\|A\|_{1 \rightarrow 2} = \sqrt{m}$. Thus, the dual objective function φ is m -gradient Lipschitz with respect to the ℓ_2 -norm. This implies that the squared norm of the gradient is bounded by the dual objective gap (Nesterov, 2018). We present this result in the following lemma and provide the proof for the sake of completeness.

Lemma 5. *For any given vector $\beta \in \mathbb{R}^{nm}$, we have*

$$\sum_{i=1}^m \left(\varphi(\beta) - \operatorname{argmin}_{\gamma \in \mathbb{R}^n} \varphi(\beta_1, \dots, \beta_{i-1}, \gamma, \beta_{i+1}, \dots, \beta_m) \right) \geq \left(\frac{1}{2m} \right) \|\nabla \varphi(\beta)\|^2.$$

Proof We derive from Eq. (14) with $\bar{\beta}_i = \beta_i - \frac{1}{m} \nabla_{\beta_i} \varphi(\beta)$ and $\bar{\beta}_k = \beta_k$ for $k \neq i$ that

$$\varphi(\beta) - \operatorname{argmin}_{\gamma \in \mathbb{R}^n} \varphi(\beta_1, \dots, \beta_{i-1}, \gamma, \beta_{i+1}, \dots, \beta_m) \leq \varphi(\beta) - \varphi(\bar{\beta}) \leq \varphi(\beta) - \left(\frac{1}{2m} \right) \|\nabla_i \varphi(\beta)\|^2.$$

Summing up the above inequality over $i \in [m]$ yields the desired inequality. \blacksquare

3. Computational Hardness

In this section, we show that the multimarginal optimal transport (MOT) problem in the form of Eq. (2) is not a minimum-cost flow problem when $m \geq 3$. The proof idea is based on a simple reduction with m -dimensional matching problem.

3.1 Unimodularity, minimum-cost flow and matching

We present some definitions and classical results in combinatorial optimization and graph theory, including unimodularity, minimum-cost flow and matching.

Definition 6. *A totally unimodular (TU) matrix is one for which every square submatrix has determinant -1 , 0 or 1 .*

A direct way to determine whether a matrix is totally unimodular or not is by computing the determinants of every square submatrix of this matrix. However, it is clearly intractable in general. The following proposition provides an alternative way to check whether a matrix is TU or not.

Proposition 7. *Let A be a $\{-1, 0, 1\}$ -valued matrix. A is TU if each column contains at most two nonzero entries and all rows are partitioned into two sets I_1 and I_2 such that: If two nonzero entries of a column have the same sign, they are in different sets. If these two entries have different signs, they are in the same set.*

In what follows, we present the definition of minimum-cost flow problem and prove that the constraint matrix of LP representation of a minimum-cost flow problem is TU. Such result is well known and can be derived from Berge (2001, Theorem 1, Chapter 15) which shows that the incidence matrices of every directed graphs are TU. For the sake of completeness, we provide the detailed proof based on Proposition 7.

Definition 8. *The minimum-cost flow problem finds the cheapest possible way of sending a certain amount of flow through a flow network. Formally,*

$$\begin{aligned} \min \quad & \sum_{(u,v) \in E} f(u,v) \cdot a(u,v) \\ \text{s.t.} \quad & f(u,v) \geq 0, \quad \text{for all } (u,v) \in E, \\ & f(u,v) \leq c(u,v) \quad \text{for all } (u,v) \in E, \\ & f(u,v) = -f(v,u) \quad \text{for all } (u,v) \in E, \\ & \sum_{(u,w) \in E \text{ or } (w,u) \in E} f(u,w) = 0, \\ & \sum_{w \in V} f(s,w) = d \quad \text{and} \quad \sum_{w \in V} f(w,t) = d. \end{aligned}$$

The flow network $G = (V, E)$ is a directed graph $G = (V, E)$ with a source vertex $s \in V$ and a sink vertex $t \in V$, where each edge $(u, v) \in E$ has capacity $c(u, v) > 0$, flow $f(u, v) \geq 0$ and cost $a(u, v)$, with most minimum-cost flow algorithms supporting edges with negative costs. The cost of sending this flow along an edge (u, v) is $f(u, v) \cdot a(u, v)$. The problem requires an amount of flow d to be sent from source s to sink t . The definition of the problem is to minimize the total cost of the flow over all edges.

Proposition 9. *The constraint matrix arising from a minimum-cost flow problem is TU.*

Proof The standard LP representation of the minimum-cost flow problem is

$$\min_{x \in \mathbb{R}^{|E|}} c^\top x, \quad \text{s.t. } Ax = b, \quad l \leq x \leq u.$$

where $x \in \mathbb{R}^{|E|}$ with x_j being the flow through arc j , $b \in \mathbb{R}^{|V|}$ with b_i being external supply at node i and $\mathbf{1}^\top b = 0$, c_j is unit cost of flow through arc j , l_j and u_j are lower and upper bounds on flow through arc j and $A \in \mathbb{R}^{|V| \times |E|}$ is the arc-node incidence matrix with entries

$$A_{ij} = \begin{cases} -1 & \text{if arc } j \text{ starts at node } i \\ 1 & \text{if arc } j \text{ ends at node } i \\ 0 & \text{otherwise} \end{cases}.$$

Since each arc has two endpoints, the constraint matrix A is a $\{-1, 0, 1\}$ -valued matrix in which each column contains two nonzero entries 1 and -1 . Using Proposition 7, we obtain that A is TU and the rows of A are categorized into a single set. ■

We proceed to the definition of m -dimensional matching which generalizes 2-dimensional matching. We present it in graph-theoretic sense as follows.

Definition 10. Let S_1, S_2, \dots, S_m be finite and disjoint sets, and let T be a subset of $S_1 \times \dots \times S_m$. That is, T consists of vectors (z_1, \dots, z_m) such that $z_i \in S_i$ for all $i \in [m]$. Now $M \subseteq T$ is a m -dimensional matching if the following holds: for any two distinct vectors $(z_1, \dots, z_m) \in M$ and $(z'_1, \dots, z'_m) \in M$, we have $z_i \neq z'_i$ for all $i \in [m]$.

In computational complexity theory, m -dimensional matching refers to the following decision problem: given a set T and an integer k , decide whether there exists a m -dimensional matching $M \subseteq T$ with $|M| \geq k$. This problem is NP-complete even when $m = 3$ and $k = |S_1| = |S_2| = |S_3|$ (Karp, 1972; Garey and Johnson, 2002). A m -dimensional matching is also an exact cover since the set M covers each element of S_1, S_2, \dots, S_m exactly once.

3.2 Main result

The problem of computing the MOT distance between $m \geq 2$ discrete probability distributions with at most n supports is equivalent to solving the following LP (cf. Eq. (2)):

$$\min_{X \in \mathbb{R}^{n \times \dots \times n}} \langle C, X \rangle, \quad \text{s.t. } r_k(X) = r_k \text{ for any } k \in [m], X \geq 0.$$

In other words, the MOT problem is a LP with mn equality constraints and n^m variables. When $m = 2$, the MOT problem is the classical OT problem (Villani, 2003) which is known to be a minimum-cost flow problem. Such problem structure is computationally favorable and permits the development of provably efficient algorithms, including the network simplex algorithms (Orlin, 1997; Tarjan, 1997) and specialized interior-point algorithms (Lee and Sidford, 2014). However, it remains unknown if the MOT problem in the above LP form admits such a structural decomposition when $m \geq 3$.

We present a negative answer to this question for $m \geq 3$. Before proceeding to the main theorem, we provide a simple yet intuitive counterexample.

Example 3. We consider arguably the simplest MOT problem, with $m = 3$ distributions supported on $n = 2$ elements each. We consider the $n^m = 8$ entries of a multimarginal tensor transportation plan, and number them slice by slice. A naive enumeration of all the marginal constraints results in nm linear equalities, but some of them are redundant since they involve several times the constraints that the sum of the elements of that tensor sum to 1. The number of required constraints is $m(n - 1) + 1$, namely only 4 mass conservation constraints are effective in this case. We therefore obtain the following matrix,

$$A = \begin{pmatrix} 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

We form the sub-matrix by only considering the first, fourth, sixth, and seventh columns of A , and can then check that the resulting matrix has determinant equal to 2, namely,

$$\det(A_{1,4,6,7}) = \det \left(\begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 \end{pmatrix} \right) = 2.$$

Therefore, the marginal constraint matrix is not totally unimodular, illustrating that the MOT with $(m, n) = (3, 2)$ is not a minimum-cost flow problem. More generally, one can numerically check that the constraint matrix corresponding to m marginals with n points each has size $(mn - m + 1) \times n^m$, and that it is not totally unimodular by selecting a subset of $(mn - m + 1)$ columns (out of n^m) that form a determinant that is neither $-1, 0, 1$. The constraint matrix itself can be obtained recursively, by defining first $L_{n,1} = I_n$, to apply next that for $t \geq 2$,

$$L_{n,t} = \begin{bmatrix} \mathbf{1}_n \otimes L_{n,t-1} & I_n \otimes \mathbf{1}_{n^{t-1}} \end{bmatrix} \in \mathbb{R}^{n^t \times nt},$$

where \otimes is Kronecker's product. In that case, $L_{n,m}$ corresponds to the matrix constraint of the dual multimarginal OT problem, which involves constraints of the type $(\alpha_1)_{i_1} + (\alpha_2)_{i_1} + \dots + (\alpha_m)_{i_m} \leq C_{i_1 i_2 \dots i_m}$ as mentioned in the next section. The constraint matrix in the primal, specified over the entries of transportation tensors, is $A_{n,m} = \tilde{L}_{n,m}^\top$, where $\tilde{L}_{n,m}$ is equal to $L_{n,m}$ stripped of $m - 1$ columns (one for each marginal but for the first), indexed for instance at $2n, 3n, \dots, nm$.

Example 3 provides some intuitions why the MOT problem is not a minimum-cost flow problem when $m \geq 3$. However, it is not easy to extend this approach to the general setting. Indeed, the constraint matrix in Eq. (2) becomes complicated when m and n are considerably large. Thus, it is challenging to compute the determinants of even a small fraction of sub-matrices, which is necessary to determine whether the constraint matrix is totally unimodular or not. While the direct calculation is intractable, some combinatorial optimization toolbox, e.g., Ghouila-Houri's theorem (Ghouila-Houri, 1962), might be helpful. However, we do not have concrete idea now and leave this topic to the future work.

Despite the above discussion, we can prove that the MOT problem in Eq. (2) is not a minimum-cost flow problem when $m \geq 3$ by using a simple reduction with m -dimensional matching problem. Roughly speaking, if the MOT problem is a minimum-cost flow problem when $m \geq 3$, its integer programming counterpart with specific choice of the cost tensor C and marginals $\{r_k\}$ must not be NP-hard. However, due to such specific choice, we can prove that this integer programming counterpart is equivalent to m -dimensional matching problem which is known as NP-complete when $m \geq 3$. This leads to the contradiction.

We present our theorem with the proof details as follows.

Theorem 11. *The MOT problem in the form of Eq. (2) is not a minimum-cost flow problem when $m \geq 3$.*

Proof We prove the result by contradiction. Indeed, we assume that the MOT problem in Eq. (2) is a minimum-cost flow problem when $m \geq 3$. Let $r_k = \frac{\mathbf{1}_n}{n}$ for all $k \in [m]$ in Eq. (2), the resulting LP is equivalent to the following problem

$$\min_X \langle C, X \rangle, \quad \text{s.t. } r_k(X) = \mathbf{1}_n \text{ for any } k \in [m], \quad X \geq 0. \quad (15)$$

We see from Eq. (15) that this is a minimum-cost flow problem where the constraint matrix and the right-hand side vector are both integer-valued. Then we consider the integer programming counterpart of Eq. (15) which is defined by

$$\begin{aligned} \min_X \quad & \langle C, X \rangle, \\ \text{s.t.} \quad & r_k(X) = \mathbf{1}_n, \text{ for any } k \in [m], \\ & X_{i_1 i_2 \dots i_m} \in \{0, 1\} \text{ for any } (i_1, \dots, i_m) \in [n] \times \dots \times [n]. \end{aligned} \quad (16)$$

It is well known in the combinatorial optimization literature (Schrijver, 2003) that Eq. (16) is not NP-hard when $m \geq 3$.

On the other hand, we claim that Eq. (16) is NP-complete when $m \geq 3$ since it reduces to an m -dimensional matching problem. Indeed, we let $S_i = [n]$ for all $i \in [m]$ and $T = [n] \times [n] \times \dots \times [n]$ as well as the cost tensor C is defined by

$$C_{i_1 i_2 \dots i_m} = \begin{cases} 1 & (i_1, i_2, \dots, i_m) \in T \\ 0 & (i_1, i_2, \dots, i_m) \notin T \end{cases} .$$

Then the objective function of any feasible solution is n so any feasible solution is an optimal solution. Furthermore, finding any optimal solution X is equivalent to finding an m -dimensional matching M . Indeed, we can define an one-to-one mapping as follows,

$$X_{i_1 i_2 \dots i_m} = \begin{cases} 1 & (i_1, i_2, \dots, i_m) \in M \\ 0 & (i_1, i_2, \dots, i_m) \notin M \end{cases} .$$

It is clear that X is a feasible solution of Eq. (16) if and only if M is an m -dimensional matching M . Thus, Eq. (16) is NP-complete when $m \geq 3$, which leads to a clear contradiction. This completes the proof of Theorem 11. ■

3.3 Discussion

We make a few comments on our main result for the MOT problem in Eq. (2), which help strengthen the understanding of the MOT problem.

First, Theorem 11 only holds true for the general MOT problem in Eq. (2). To be more specific, we show that there exists the cost tensor C and marginals $\{r_k\}$ such that the MOT problem in Eq. (2) is not a minimum-cost flow problem. Nonetheless, the MOT problem is commonly referred to as the LP in Eq. (2) and thus it is important to understand the structure of this LP when $m \geq 3$. Further, we remark that there is no other reformulation of the general MOT problems which can be solved efficiently; in particular, Altschuler and Boix-Adsera (2021a) showed that the MOT problems with repulsive costs are computationally intractable: several such problems of interest are NP-hard to solve – even approximately.

Second, Theorem 11 does not rule out the possibility that a few instances of the MOT problem are minimum-cost flow problems when $m \geq 3$. However, such examples need to admit special structure and are thus rare in real applications given that Example 3 is one of the simplest MOT problems. Many common MOT problems, e.g., Wasserstein barycenters, are not minimum-cost flow problems. Indeed, Lin et al. (2020) proved this result for even the simplest Wasserstein barycenter problems in the standard LP form — the fixed-support Wasserstein barycenters when $m \geq 3$ and $n \geq 3$. Despite such negative result, the discrete Wasserstein barycenter problems have their own structure (Anderes et al., 2016) and can be efficiently solved in practice (Cuturi and Doucet, 2014; Benamou et al., 2015; Carlier et al., 2015; Staib et al., 2017; Clatici et al., 2018; Dvurechensky et al., 2018a; Kroshnin et al., 2019; Ge et al., 2019). In conclusion, the MOT problem with $m \geq 3$ is different from the OT problem and we believe that the minimum-cost flow is not sufficient for characterizing the structure of the MOT problem.

Finally, Theorem 11 affects the complexity bound of various algorithms for solving the MOT problem when $m \geq 3$. If we could write the MOT problem as a minimum-cost flow problem on the directed graph with n^m edges and mn vertices, the network simplex method achieves the complexity bound of $\tilde{O}(m^2 n^{m+2})$ (Orlin, 1997) or better bound of $\tilde{O}(mn^{m+1})$ (Tarjan, 1997), while the specialized interior-point algorithm achieves the complexity bound of $\tilde{O}(\sqrt{mnn^m})$ (Lee and Sidford, 2014). However, due to Theorem 11, these bounds of network simplex method and specialized interior-point algorithms are not valid and only the standard interior-point algorithms can achieve much worse complexity bound of $\tilde{O}(n^{3m})$ (Wright, 1997). We are also aware of a *stochastic central path method* (Cohen et al., 2019) which achieves better complexity bound of $\tilde{O}(n^{\omega m})$ with the coefficient of matrix multiplication $\omega \approx 2.38$. However, this algorithm is seemingly not implementable in practice and not comparable with the deterministic algorithms we propose in this paper.

4. Multimarginal Sinkhorn Algorithm

In this section, we propose and analyze a multimarginal Sinkhorn algorithm for solving the entropic regularized multimarginal optimal transport (MOT) problem. We also generalize the rounding scheme (Altschuler et al., 2017) to the MOT setting. Together with a new rounding scheme, our algorithm achieves a complexity bound of $\tilde{O}(m^3 n^m \varepsilon^{-2})$ when applied to solve the MOT problem. The proof techniques are heavily based on the smooth dual objective function in Eq. (6) and thus not a straightforward generalization of the analysis in the OT setting (Altschuler et al., 2017; Dvurechensky et al., 2018b; Lin et al., 2019a) where they use different form of dual objective function; see Remark 2 for the details.

4.1 Algorithmic procedure

We present the pseudocode of the multimarginal Sinkhorn algorithm in Algorithm 1. This algorithm is a new generalization of the classical Sinkhorn algorithm (Cuturi, 2013) and different from the existing multimarginal Sinkhorn algorithms (Benamou et al., 2015, 2019; Peyré and Cuturi, 2019). Indeed, the main difference lies in the greedy choice of the next marginal (cf. **Step 2**). This simple yet crucial modification makes our complexity bound analysis work while only the asymptotic convergence properties are proved for the existing multimarginal Sinkhorn algorithms.

Comments on algorithmic scheme. Algorithm 1 can be interpreted as a greedy block coordinate descent algorithm (Dhillon et al., 2011; Nutini et al., 2015) for solving the dual entropic regularized MOT problem in Eq. (6); see Tupitsa et al. (2020, Lemma 1) for the justification. However, the corresponding known complexity bounds for greedy block coordinate descent algorithms can not be applied to analyze Algorithm 1. Indeed, the per-iteration progress is quantified using the ℓ_2 -norm in the existing algorithmic scheme and convergence analysis. This will lead to the worse complexity bound than ours since it does not respect the structure of MOT problem. In contrast, Algorithm 1 employs KL divergence to quantify the per-iteration progress and link it to the ℓ_1 -norm via appeal to the Pinsker inequality (Cover and Thomas, 2012).

Algorithm 1 MULTISINKHORN($C, \eta, \{r_k\}_{k \in [m]}, \varepsilon'$)

Initialization: $t = 0$ and $\beta^0 \in \mathbb{R}^{mn}$ with $\beta^0 = \mathbf{0}_{mn}$.

while $E_t > \varepsilon'$ **do**

Step 1. Choose the greedy coordinate $K = \operatorname{argmax}_{1 \leq k \leq m} \rho(r_k, r_k(B(\beta^t)))$.

Step 2. Compute $\beta^{t+1} \in \mathbb{R}^{mn}$ by

$$\beta_k^{t+1} = \begin{cases} \beta_k^t + \log(r_k) - \log(r_k(B(\beta^t))), & k = K \\ \beta_k^t, & \text{otherwise} \end{cases} .$$

Step 3. Increment by $t = t + 1$.

end while

Output: $B(\beta^t)$.

More specifically, an exact coordinate update for the K -th variable is performed at each iteration while other variables are fixed. Here we choose K by using the greedy rule as follows,

$$K = \operatorname{argmax}_{1 \leq k \leq m} \rho(r_k, r_k(B(\beta^t))),$$

where $\rho : \mathbb{R}_+^n \times \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ is defined as

$$\rho(a, b) := \mathbf{1}_n^\top (b - a) + \sum_{i=1}^n a_i \log \left(\frac{a_i}{b_i} \right).$$

Following up the optimal transport literature (Cuturi, 2013; Altschuler et al., 2017), we set the stopping criterion as $E_t \leq \varepsilon'$ for some tolerance $\varepsilon' > 0$, where E_t is defined by

$$E_t := \sum_{k=1}^m \|r_k(B(\beta^t)) - r_k\|_1. \tag{17}$$

Comments on arithmetic operations per iteration. The most expensive step is to determine which coordinate is the greedy one. While the naive way requires $O(mn^m)$ arithmetic operations to compute all marginals $r_k(B(\beta^t))$, we can adopt some implementation tricks based on the observation that *one of $\rho(r_k, r_k(B(\beta^t))) = 0$ after the first step.*

Without loss of generality, we assume that $m \geq 3$ and $r_1(B(\beta^t)) = r_1$. The key step is to construct a small tensor A which has n^{m-1} entries: $A_{i_1, \dots, i_{m-1}} = \sum_{j=1}^n B_{j, i_1, \dots, i_{m-1}}$ for any $(i_1, \dots, i_{m-1}) \in [n] \times \dots \times [n]$. This requires $O(n^m)$ arithmetic operations. It is clear that $r_2(B(\beta^t)), \dots, r_m(B(\beta^t))$ exactly corresponds to the marginals of A and the computation only needs $O(mn^{m-1})$ arithmetic operations. Putting these pieces together yields that the arithmetic operations per iteration is $O(n^m)$ for the case of $m = O(n)$.

Rounding scheme. Algorithm 1 is developed for solving the entropic-regularized MOT problem and the output is not necessarily a feasible solution of *unregularized* MOT problem. To address this issue, we develop a new rounding scheme by extending Altschuler et al. (2017, Algorithm 2) to the MOT setting; see Algorithm 2. We can see that the difference between

. Meshi et al. (2012) showed that the greedy rule was implemented efficiently by using a max-heap structure for many structured problems.

Algorithm 2 ROUND($X, \{r_k\}_{k \in [m]}$)

Initialization: $X^{(0)} = X$.

for $k = 1$ to m **do**

 Compute $z_k = \min\{\mathbf{1}_n, r_k/r_k(X^{(k-1)})\} \in \mathbb{R}^n$.

for $j = 1$ to n **do**

$X_{i_1 i_2 \dots i_m}^{(k)} = z_{kj} X_{i_1 i_2 \dots i_m}^{(k-1)}$ in which $i_k = j$ is fixed.

end for

end for

 Compute $\text{err}_k \in \mathbb{R}^n$ such that $\text{err}_k = r_k - r_k(X^{(m)})$ for all $k \in [m]$.

 Compute $Y \in \mathbb{R}^{n \times n \times \dots \times n}$ by

$$Y_{i_1 i_2 \dots i_m} = X_{i_1 i_2 \dots i_m}^{(m)} + \frac{\prod_{k=1}^m \text{err}_k^{i_k}}{\|\text{err}_1\|_1^{m-1}} \text{ for any } (i_1, \dots, i_m) \in [n] \times \dots \times [n].$$

Output: Y .

the input and output of Algorithm 2 is simply a rank-one tensor. Using the approach presented by (Lacombe et al., 2018, Proposition 4), we can compute $\langle C, Y \rangle$ efficiently using $\langle C, X \rangle$. Finally, the total arithmetic operations required by Algorithm 2 is $O(mn^m)$.

Algorithm for the MOT problem. We present the pseudocode of our main algorithm in Algorithm 3, where Algorithm 1 and 2 are the subroutines. We notice that the regularization parameter η is scaled as a function of the desired accuracy $\varepsilon > 0$, and remark that **Step 1** is necessary since the multimarginal Sinkhorn algorithm is not well behaved if the marginal distributions do not have dense support.

4.2 Technical lemmas

In this section, we provide two technical lemmas which are important in the analysis of Algorithm 1. The first lemma shows that the dual objective gap at iteration t is bounded by the product between the residue term E_t and a constant depending on η , C and $\{r_i\}_{i \in [m]}$.

Lemma 12. *Let $\{\beta^t\}_{t \geq 0}$ be the iterates generated by Algorithm 1 and β^* be an optimal solution which is specified by Lemma 3. Then, the following inequality holds true:*

$$\varphi(\beta^t) - \varphi(\beta^*) \leq \bar{R} E_t, \quad \text{for all } t \geq 1.$$

where E_t is defined in Eq. (17) and $\bar{R} > 0$ is defined as

$$\bar{R} := \frac{\|C\|_\infty}{\eta} - \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij} \right).$$

Proof We first prove that the following inequality holds true,

$$\begin{aligned} \max_{1 \leq i \leq m} \left\{ \max_{1 \leq j \leq n} \beta_{ij}^t - \min_{1 \leq j \leq n} \beta_{ij}^t \right\} &\leq \bar{R}, \\ \max_{1 \leq i \leq m} \left\{ \max_{1 \leq j \leq n} \beta_{ij}^* - \min_{1 \leq j \leq n} \beta_{ij}^* \right\} &\leq \bar{R}. \end{aligned} \tag{18}$$

Algorithm 3 Approximating MOT by Algorithm 1 and 2

Input: $\eta = \frac{\varepsilon}{2m \log(n)}$ and $\varepsilon' = \frac{\varepsilon}{8\|C\|_\infty}$.

Step 1: Let $\tilde{r}_k \in \Delta_n$ for $\forall k \in [m]$ be defined as

$$(\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_m) = \left(1 - \frac{\varepsilon'}{4m}\right) (r_1, r_2, \dots, r_m) + \frac{\varepsilon'}{4mn} (\mathbf{1}_n, \mathbf{1}_n, \dots, \mathbf{1}_n).$$

Step 2: Compute $\tilde{X} = \text{MULTISINKHORN}(C, \eta, \{\tilde{r}_k\}_{k \in [m]}, \varepsilon'/2)$.

Step 3: Round $\hat{X} = \text{ROUND}(\tilde{X}, \{\tilde{r}_k\}_{k \in [m]})$.

Output: \hat{X} .

Note that the second inequality is a straightforward deduction of Eq. (9) in the proof of Lemma 3. Thus, it suffices to prove the first inequality.

We establish this by an induction argument. Indeed, this inequality holds trivially when $t = 0$. Assume that this inequality holds true for $t \leq T$. By the update for β in Algorithm 1, $\beta_k^{T+1} = \beta_k^T$ for all $k \neq K$, where $K = \arg\max_{1 \leq k \leq m} \rho(r_k, r_k(B(\beta^t)))$. This implies that

$$\max_{1 \leq j \leq n} \beta_{kj}^{T+1} - \min_{1 \leq j \leq n} \beta_{kj}^{T+1} \leq \bar{R} \text{ for all } k \neq K.$$

Now it remains to show $\max_{1 \leq j \leq n} \beta_{Kj}^{T+1} - \min_{1 \leq j \leq n} \beta_{Kj}^{T+1} \leq \bar{R}$. For any $l \in [n]$, we derive from the update formula of β_{Kl}^{T+1} that

$$e^{\beta_{Kl}^{T+1}} \sum_{1 \leq i_k \leq n, \forall k \neq K} e^{\sum_{k \neq K} \beta_{ki_k}^T - \eta^{-1} C_{i_1 \dots i_m}} = r_{Kl} \geq \min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij}.$$

Since C is a nonnegative cost tensor, we derive from the above inequality that

$$\beta_{Kl}^{T+1} \geq \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} r_{ij} \right) - \log \left(\sum_{1 \leq i_k \leq n, \forall k \neq K} e^{\sum_{k \neq K} \beta_{ki_k}^T} \right). \quad (19)$$

Since $r_{Kl} \leq 1$ and $C_{i_1 \dots i_m} \leq \|C\|_\infty$, we have

$$\beta_{Kl}^{T+1} \leq \frac{\|C\|_\infty}{\eta} - \log \left(\sum_{1 \leq i_k \leq n, \forall k \neq K} e^{\sum_{k \neq K} \beta_{ki_k}^T} \right). \quad (20)$$

Combining the bounds (19) and (20) implies the desired result.

Then, we proceed to the proof of Lemma 12. Since the function φ is convex and β^* is an optimal solution, Eq. (18) implies that

$$\varphi(\beta^t) - \varphi(\beta^*) \leq (\beta^t - \beta^*)^\top \nabla \varphi(\beta^t) = \sum_{k=1}^m (\beta_k^t - \beta_k^*)^\top \left(\frac{r_k(B(\beta^t))}{\|B(\beta^t)\|_1} - r_k \right).$$

Note that the initialization and the main update for the variable β in Algorithm 1 imply that $\|B(\beta^t)\|_1 = 1$ for all $t \geq 1$. Thus, we have

$$\varphi(\beta^t) - \varphi(\beta^*) \leq \sum_{k=1}^m (\beta_k^t - \beta_k^*)^\top (r_k(B(\beta^t)) - r_k). \quad (21)$$

Furthermore, we have $\mathbf{1}_n^\top r_k(B(\beta^t)) = \mathbf{1}_n^\top r_k = 1$ for all $k \in [m]$. This implies

$$\mathbf{1}_n^\top (r_k(B(\beta^t)) - r_k) = 0 \text{ for all } k \in [m]. \quad (22)$$

For all $k \in [m]$, we define $2m$ shift terms as follows,

$$\Delta\beta_k^t = \frac{\max_{1 \leq j \leq n} \beta_{kj}^t + \min_{1 \leq j \leq n} \beta_{kj}^t}{2}, \quad \Delta\beta_i^* = \frac{\max_{1 \leq j \leq n} \beta_{kj}^* + \min_{1 \leq j \leq n} \beta_{kj}^*}{2}.$$

Using these shift terms, we derive that

$$\begin{aligned} & (\beta_k^t - \beta_k^*)^\top (r_k(B(\beta^t)) - r_k) \\ & \stackrel{(22)}{=} (\beta_k^t - \Delta\beta_k^t \mathbf{1}_n)^\top (r_k(B(\beta^t)) - r_k) - (\beta_k^* - \Delta\beta_k^* \mathbf{1}_n)^\top (r_k(B(\beta^t)) - r_k) \\ & \leq (\|\beta_k^t - \Delta\beta_k^t \mathbf{1}_n\|_\infty + \|\beta_k^* - \Delta\beta_k^* \mathbf{1}_n\|_\infty) \|r_k(B(\beta^t)) - r_k\|_1 \\ & = \left(\max_{1 \leq j \leq n} \beta_{kj}^t - \min_{1 \leq j \leq n} \beta_{kj}^t + \max_{1 \leq j \leq n} \beta_{kj}^* - \min_{1 \leq j \leq n} \beta_{kj}^* \right) \frac{\|r_k(B(\beta^t)) - r_k\|_1}{2}. \end{aligned}$$

Plugging Eq. (18) into the above inequality yields

$$(\beta_k^t - \beta_k^*)^\top (r_k(B(\beta^t)) - r_k) \leq \bar{R} \|r_k(B(\beta^t)) - r_k\|_1. \quad (23)$$

Combining Eq. (23) and Eq. (21) yields

$$\varphi(\beta^t) - \varphi(\beta^*) \leq \bar{R} \left(\sum_{k=1}^m \|r_k(B(\beta^t)) - r_k\|_1 \right) = \bar{R} E_t.$$

As a consequence, we obtain the conclusion of the lemma. \blacksquare

The second lemma gives a descent inequality for the iterates generated by Algorithm 1 with a lower bound on the progress at each iteration.

Lemma 13. *Let $\{\beta^t\}_{t \geq 0}$ be the iterates generated by Algorithm 1. Then, the following inequality holds true:*

$$\varphi(\beta^t) - \varphi(\beta^{t+1}) \geq \frac{1}{2} \left(\frac{E_t}{m} \right)^2, \quad \text{for all } t \geq 1. \quad (24)$$

Proof We first show that

$$\varphi(\beta^t) - \varphi(\beta^{t+1}) \geq \frac{1}{m} \left(\sum_{k=1}^m \rho(r_k, r_k(B(\beta^t))) \right). \quad (25)$$

By the definition of φ , we have

$$\varphi(\beta^t) - \varphi(\beta^{t+1}) = \log(\|B(\beta^t)\|_1) - \log(\|B(\beta^{t+1})\|_1) - \sum_{k=1}^m (\beta_k^t - \beta_k^{t+1})^\top r_k. \quad (26)$$

From the update formula for β^{t+1} , it is clear that $\|B(\beta^t)\|_1 = \|B(\beta^{t+1})\|_1 = 1$ for all $t \geq 1$. Therefore, we have

$$\varphi(\beta^t) - \varphi(\beta^{t+1}) = -(\beta_K^t - \beta_K^{t+1})^\top r_K = (\log(r_K) - \log(r_K(B(\beta^t))))^\top r_K.$$

Since $\mathbf{1}_n^\top r_K = \mathbf{1}_n^\top r_K(B(\beta^t)) = 1$, we have $\varphi(\beta^t) - \varphi(\beta^{t+1}) = \rho(r_K, r_K(B(\beta^t)))$. Combining this equality with the fact that the K -th coordinate is the greedy one yields Eq. (25).

We proceed to prove Eq. (24). Indeed, by the Pinsker inequality, we have

$$\rho(r_k, r_k(B(\beta^t))) \geq \frac{1}{2} \|r_k(B(\beta^t)) - r_k\|_1^2 \text{ for any } k \in [m].$$

Plugging this inequality into Eq. (25) and using the Cauchy-Schwarz inequality yields

$$\begin{aligned} \varphi(\beta^t) - \varphi(\beta^{t+1}) &\geq \frac{1}{2m} \left(\sum_{k=1}^m \|r_k(B(\beta^t)) - r_k\|_1^2 \right) \\ &\geq \frac{1}{2m^2} \left(\sum_{k=1}^m \|r_k(B(\beta^t)) - r_k\|_1 \right)^2 = \frac{1}{2} \left(\frac{E_t}{m} \right)^2. \end{aligned}$$

This completes the proof. ■

4.3 Main results

We present an upper bound for the number of iterations required by Algorithm 1.

Theorem 14. *Let $\{\beta^t\}_{t \geq 0}$ be the iterates generated by Algorithm 1. The number of iterations required to reach the stopping criterion $E_t \leq \varepsilon'$ satisfies*

$$t \leq 2 + \frac{2m^2 \bar{R}}{\varepsilon'}, \quad (27)$$

where \bar{R} is defined in Lemma 12.

Proof Let β^* be an optimal solution of the dual entropic regularized MOT problem considered in Lemma 12. By letting the objective gap at each iteration be $\delta^t = \varphi(\beta^t) - \varphi(\beta^*)$, we derive from Lemma 12 and Lemma 13 that

$$\delta^t \leq \bar{R} E_t, \quad \delta^t - \delta^{t+1} \geq \frac{1}{2} \left(\frac{E_t}{m} \right)^2.$$

Putting these pieces together with the fact that $E_t \geq \varepsilon'$ as long as the stopping criterion is not fulfilled yields

$$\delta^t - \delta^{t+1} \geq \frac{1}{2} \left(\max \left\{ \left(\frac{\varepsilon'}{m} \right)^2, \left(\frac{\delta^t}{m \bar{R}} \right)^2 \right\} \right).$$

We now apply the switching strategy to obtain the desired upper bound in Eq. (27). Indeed, we have

$$\frac{\delta^{t+1}}{2m^2 \bar{R}^2} \leq \frac{\delta^t}{2m^2 \bar{R}^2} - \frac{(\delta^t)^2}{4m^4 \bar{R}^4}, \quad \delta^{t+1} \leq \delta^t - \frac{1}{2} \left(\frac{\varepsilon'}{m} \right)^2.$$

Fixing an integer $t_1 > 0$ and considering $t > t_1$, the first inequality further implies that

$$\frac{2m^2 \bar{R}^2}{\delta^{t+1}} - \frac{2m^2 \bar{R}^2}{\delta^t} \geq 1 \implies t_1 \leq 1 + \frac{2m^2 \bar{R}^2}{\delta^{t_1}},$$

and the second inequality further implies that

$$t - t_1 \leq 1 + 2(\delta^{t_1} - \delta^t) \left(\frac{m}{\varepsilon'}\right)^2 \implies t \leq 1 + t_1 + 2\delta^{t_1} \left(\frac{m}{\varepsilon'}\right)^2.$$

Let $s = \delta_{t_1} \leq \delta_1$, we obtain that the total number of iterations satisfies

$$t \leq \min_{0 \leq s \leq \delta_1} \left\{ 2 + \frac{2m^2 \bar{R}^2}{s} + 2s \left(\frac{m}{\varepsilon'}\right)^2 \right\} \leq 2 + \frac{2m^2 \bar{R}}{\varepsilon'}.$$

This completes the proof. \blacksquare

Before presenting the main result on the complexity bound of Algorithm 3, we provide the complexity bound of Algorithm 2 in the following theorem.

Theorem 15. *Let $X \in \mathbb{R}^{n \times \dots \times n}$ be a nonnegative tensor and $\{r_i\}_{i \in [m]} \subseteq \Delta^n$ be a sequence of probability vectors, Algorithm 2 returns a nonnegative tensor $Y \in \mathbb{R}^{n \times \dots \times n}$ satisfying that $r_k(Y) = r_k$ for all $k \in [m]$ and*

$$\|Y - X\|_1 \leq 2 \left(\sum_{k=1}^m \|r_k(X) - r_k\|_1 \right).$$

Proof By the definition of z_k and the update formula for $X^{(k)}$ for all $k \in [m]$, each entry of $X^{(m)}$ is nonnegative and

$$\text{err}_k = r_k - r_k(X^{(m)}) \geq 0 \text{ for all } k \in [m]. \quad (28)$$

This implies that $\|\text{err}_k\|_1 = 1 - \|X^{(m)}\|_1$ for all $k \in [m]$. Thus, we derive from Eq. (28) and the update formula for Y that each entry of Y is nonnegative.

Furthermore, we define A by $A_{i_1 \dots i_m} := \prod_{k=1}^m \text{err}_{k i_k}$ for all $(i_1, \dots, i_m) \in [n] \times \dots \times [n]$ and find that

$$[r_k(A)]_j = \text{err}_{kj} \left(\sum_{1 \leq i_l \leq n, \forall l \neq k} \prod_{l \neq k} \text{err}_{l i_l} \right) = \text{err}_{kj} \prod_{l \neq k} \|\text{err}_l\|_1. \quad (29)$$

Therefore, we conclude that

$$r_k(Y) = r_k(X^{(m)}) + \frac{r_k(A)}{\|\text{err}_1\|_1^{m-1}} \stackrel{(29)}{=} r_k(X^{(m)}) + \text{err}_k = r_k \text{ for all } k \in [m].$$

It remains to estimate the ℓ_1 bound between Y and X . Indeed, we have

$$\|X\|_1 - \|X^{(m)}\|_1 = \|X^{(0)}\|_1 - \|X^{(m)}\|_1 = \sum_{k=1}^m (\|X^{(k-1)}\|_1 - \|X^{(k)}\|_1). \quad (30)$$

Since $\|X^{(k-1)}\|_1 - \|X^{(k)}\|_1$ is the amount of mass removed from $X^{(k-1)}$ by rescaling the k th subtensor when $r_{kj}(X^{(k-1)}) \geq r_{kj}$, we have

$$\|X^{(k-1)}\|_1 - \|X^{(k)}\|_1 = \mathbf{1}_n^\top (\max\{0, r_k(X^{(k-1)}) - r_k\}) \text{ for all } k \in [m].$$

A simple calculation using the fact that $X^{(0)} = X$ shows that

$$\|X^{(0)}\|_1 - \|X^{(1)}\|_1 = \frac{1}{2}(\|r_1(X) - r_1\|_1 + \|X\|_1 - 1). \quad (31)$$

Moreover, $r_k(X^{(0)})$ is entrywise larger than $r_k(X^{(k-1)})$ for all $k \in [m]$. That is to say, $r_k(X^{(k-1)}) \leq r_k(X^{(k-2)}) \leq \dots \leq r_k(X^{(0)}) = r_k(X)$. This implies

$$\|X^{(k)}\|_1 - \|X^{(k+1)}\|_1 \leq \|r_{k+1}(X) - r_{k+1}\|_1 \text{ for all } k \in [m-1]. \quad (32)$$

Plugging Eq. (31) and Eq. (32) into Eq. (30) yields

$$\|X\|_1 - \|X^{(m)}\|_1 \leq \frac{1}{2}(\|r_1(X) - r_1\|_1 + \|X\|_1 - 1) + \sum_{k=2}^m \|r_k(X) - r_k\|_1. \quad (33)$$

By the definition of Y , we have

$$\|X - Y\|_1 \leq \|X - X^{(m)}\|_1 + \frac{\|A\|_1}{\|\text{err}_1\|_1^{m-1}} \stackrel{(29)}{=} \|X - X^{(m)}\|_1 + \|\text{err}_1\|_1.$$

Since X is entrywise larger than $X^{(m)}$ and $\|\text{err}_1\|_1 = 1 - \|X^{(m)}\|_1$, we have

$$\|X - Y\|_1 \leq \|X\|_1 - \|X^{(m)}\|_1 + 1 - \|X^{(m)}\|_1 = 2(\|X\|_1 - \|X^{(m)}\|_1) + 1 - \|X\|_1. \quad (34)$$

Plugging Eq. (33) into Eq. (34) yields the desired result. \blacksquare

We are ready to present the complexity bound of Algorithm 3 for solving the MOT problem in Eq. (2). Note that $\varepsilon' = \varepsilon/(8\|C\|_\infty)$ is defined using the desired accuracy $\varepsilon > 0$.

Theorem 16. *Algorithm 3 returns an ε -approximate multimarginal transportation plan $\widehat{X} \in \mathbb{R}^{n \times \dots \times n}$ within*

$$O\left(\frac{m^3 n^m \|C\|_\infty^2 \log(n)}{\varepsilon^2}\right)$$

arithmetic operations.

Proof We first claim that

$$\langle C, \widehat{X} \rangle - \langle C, X^* \rangle \leq m\eta \log(n) + 4 \left(\sum_{k=1}^m \|r_k(\widetilde{X}) - r_k\|_1 \right) \|C\|_\infty. \quad (35)$$

where \widetilde{X} is defined in **Step 2** of Algorithm 3 and \widehat{X} is returned by Algorithm 3 and X^* is an optimal multimarginal transportation plan. By the definition of $\{\tilde{r}_k\}_{k \in [m]}$ and using $\sum_{k=1}^m \|r_k(\widetilde{X}) - \tilde{r}_k\|_1 \leq \varepsilon'/2$, we have

$$\sum_{k=1}^m \|r_k(\widetilde{X}) - r_k\|_1 \leq \sum_{k=1}^m (\|r_k(\widetilde{X}) - \tilde{r}_k\|_1 + \|\tilde{r}_k - r_k\|_1) \leq \frac{\varepsilon'}{2} + \sum_{k=1}^m \frac{\varepsilon'}{2m} = \varepsilon'.$$

Plugging the above inequality into Eq. (35) and using $\eta = \varepsilon/(2m \log(n))$ and $\varepsilon' = \varepsilon/(8\|C\|_\infty)$, we obtain that $\langle C, \widehat{X} \rangle - \langle C, X^* \rangle \leq \varepsilon$.

It remains to bound the number of iterations required by Algorithm 1 to reach $E_t \leq \varepsilon'/2$ (cf. **Step 2** of Algorithm 3). Using Theorem 14, we have

$$t \leq 2 + \frac{4m^2\bar{R}}{\varepsilon'}.$$

By the definition of \bar{R} (cf. Lemma 12), $\eta = \varepsilon/(2m \log(n))$ and $\varepsilon' = \varepsilon/(8\|C\|_\infty)$, we have

$$\begin{aligned} t &\leq 2 + \frac{32m^2\|C\|_\infty}{\varepsilon} \left(\frac{\|C\|_\infty}{\eta} - \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} \tilde{r}_{ij} \right) \right) \\ &\leq 2 + \frac{32m^2\|C\|_\infty}{\varepsilon} \left(\frac{2m \log(n)\|C\|_\infty}{\varepsilon} - \log \left(\frac{\varepsilon}{32mn\|C\|_\infty} \right) \right) \\ &= O \left(\frac{m^3\|C\|_\infty^2 \log(n)}{\varepsilon^2} \right). \end{aligned}$$

Since each iteration of Algorithm 1 requires $O(n^m)$ arithmetic operations, the total arithmetic operations required by **Step 2** of Algorithm 3 is $O(m^3n^m\|C\|_\infty^2 \log(n)\varepsilon^{-2})$. In addition, computing a set of vectors $\{\tilde{r}_k\}_{k \in [m]}$ requires $O(mn)$ arithmetic operations and Algorithm 2 requires $O(mn^m)$ arithmetic operations. Putting these pieces together yields that the complexity bound of Algorithm 3 is $O(m^3n^m\|C\|_\infty^2 \log(n)\varepsilon^{-2})$.

Proof of Eq. (35): Using Theorem 15, we obtain that \hat{X} is a feasible solution to the MOT problem in Eq. (2) and

$$\|\hat{X} - \tilde{X}\|_1 \leq 2 \left(\sum_{k=1}^m \|r_k(\tilde{X}) - r_k\|_1 \right).$$

This implies that

$$\langle C, \hat{X} \rangle - \langle C, \tilde{X} \rangle \leq 2\|C\|_\infty \left(\sum_{k=1}^m \|r_k(\tilde{X}) - r_k\|_1 \right). \quad (36)$$

Letting X^* be an optimal solution of the MOT problem and \tilde{Y} be the output returned by Algorithm 2 with an input X^* and $\{r_k(\tilde{X})\}_{k \in [m]}$, Theorem 15 implies

$$\|\tilde{Y} - X^*\|_1 \leq 2 \left(\sum_{k=1}^m \|r_k(X^*) - r_k(\tilde{X})\|_1 \right) = 2 \left(\sum_{k=1}^m \|r_k(\tilde{X}) - r_k\|_1 \right). \quad (37)$$

Since \tilde{X} is returned by Algorithm 1, we have $\|\tilde{X}\|_1 = 1$. By the optimality condition, there exists $\tilde{\beta} \in \mathbb{R}^{mn}$ such that $\tilde{X} = B(\tilde{\beta})$ and $\tilde{\beta}$ is an optimal solution of the following problem:

$$\min_{\beta_1, \dots, \beta_m \in \mathbb{R}^n} \log(\|B(\beta_1, \dots, \beta_m)\|_1) - \sum_{i=1}^m \beta_i^\top r_i(\tilde{X}).$$

This implies that \tilde{X} is an optimal solution of the following problem:

$$\min \langle C, X \rangle - \eta H(X), \quad \text{s.t. } r_k(X) = r_k(\tilde{X}) \text{ for all } k \in [m].$$

Since \tilde{Y} is feasible for the above problem, we have $\langle C, \tilde{X} \rangle - \eta H(\tilde{X}) \leq \langle C, \tilde{Y} \rangle - \eta H(\tilde{Y})$. Using the property of entropy regularization function (Cover and Thomas, 2012), we have $0 \leq H(\tilde{X}), H(\tilde{Y}) \leq m \log(n)$. Putting these pieces yields

$$\langle C, \tilde{X} \rangle - \langle C, \tilde{Y} \rangle \leq m\eta \log(n). \quad (38)$$

Combining Eq. (37) and Eq. (38) together with the Hölder inequality yields

$$\langle C, \tilde{X} \rangle - \langle C, X^* \rangle \leq m\eta \log(n) + 2\|C\|_\infty \left(\sum_{k=1}^m \|r_k(\tilde{X}) - r_k\|_1 \right). \quad (39)$$

Combining Eq. (36) and Eq. (39) yields

$$\langle C, \hat{X} \rangle - \langle C, X^* \rangle \leq m\eta \log(n) + 4\|C\|_\infty \left(\sum_{k=1}^m \|r_k - r_k(\tilde{X})\|_1 \right).$$

This completes the proof of Eq. (35). ■

Remark 17. *Theorem 16 demonstrates that the complexity bound of Algorithm 3 is near-linear in n^m , which is the number of unknown variable of the MOT problem in Eq. (2). This is the best possible dependence on n that we can hope for an optimization algorithm when applied to solve the general MOT problem. Further, the complexity bound has the dependence m^3 which seems unimprovable using the current techniques; indeed, the iteration number of Algorithm 1 is proportional to m^2 and the regularization parameter η is necessarily proportional to $1/m$ such that the output returned by Algorithm 1 can be rounded to an ε -approximate multimarginal transport plan.*

Remark 18. *Even though Benamou et al. (2015) has shown that the Sinkhorn-type algorithm can be more efficient in practice than LP solvers, the full theoretical analysis is not given. In contrast, our theoretical analysis provides the provably efficient way to solve the MOT problem, demonstrating the importance of the greedy update rule in the Sinkhorn-type algorithm. This leads to an algorithmic framework in which each iteration might take $O(n^m)$ number of arithmetic operations in the worst case. However, we can develop some efficient subroutines by exploiting the special structure of many MOT problems in practice and show that the required number of arithmetic operations is only polynomial in m and n (Altschuler and Boix-Adserà, 2021b,a; Altschuler and Boix-Adserà, 2022). Some of their results are based on both the algorithmic scheme and the theoretical analysis of multimarginal Sinkhorn, demonstrating the fundamental role that our analysis play in understanding the MOT problem.*

5. Accelerating Multimarginal Sinkhorn Algorithm

In this section, we present an *accelerated multimarginal Sinkhorn algorithm* for solving the entropic regularized MOT problem in Eq. (3). Together with a rounding scheme, our algorithm can be used for solving the MOT problem in Eq. (2) and achieves a complexity

Algorithm 4 ACCELERATED MULTISINKHORN($C, \eta, \{\tilde{r}_k\}_{k \in [m]}, \varepsilon'$)

Input: $t = 0, \theta_0 = 1, K = 1$ and $\tilde{\beta}^0 = \tilde{\beta}^0 = \mathbf{0}_n$.

while $E_t > \varepsilon'$ **do**

Step 1. Compute $\tilde{\beta}^t = (1 - \theta_t)\tilde{\beta}^t + \theta_t\tilde{\beta}^t$.

Step 2. Compute $\tilde{\beta}^{t+1} \in \mathbb{R}^{mn}$ by

$$\tilde{\beta}_k^{t+1} = \tilde{\beta}_k^t - \frac{1}{m\theta_t} \left(\frac{r_k(B(\tilde{\beta}^t))}{\|B(\tilde{\beta}^t)\|_1} - r_k \right) \text{ for all } k \in [m].$$

Step 3. Compute $\hat{\beta}^t = \tilde{\beta}^t + \theta_t(\tilde{\beta}^{t+1} - \tilde{\beta}^t)$.

Step 4. Compute $\hat{\beta}^t \in \mathbb{R}^{mn}$ by

$$\hat{\beta}_k^t = \begin{cases} \tilde{\beta}_k^t + \log(r_k) - \log(r_k(B(\hat{\beta}^t))), & k = K, \\ \tilde{\beta}_k^t, & \text{otherwise.} \end{cases}$$

Step 5. Compute $\beta^t = \operatorname{argmin}\{\varphi(\beta) \mid \beta \in \{\tilde{\beta}^t, \hat{\beta}^t\}\}$.

Step 6. Choose the greedy coordinate $K = \operatorname{argmax}_{1 \leq k \leq m} \rho(r_k, r_k(B(\beta^t)))$.

Step 7. Compute $\tilde{\beta}^{t+1} \in \mathbb{R}^{mn}$ by

$$\tilde{\beta}_k^{t+1} = \begin{cases} \beta_k^t + \log(r_k) - \log(r_k(B(\beta^t))), & k = K, \\ \beta_k^t, & \text{otherwise.} \end{cases}$$

Step 8. Compute $\theta_{t+1} = \theta_t(\sqrt{\theta_t^2 + 4} - \theta_t)/2$.

Step 9. Increment by $t = t + 1$.

end while

Output: $B(\beta^t)$.

bound of $\tilde{O}(m^3 n^{m+1/3} \varepsilon^{-4/3})$, which improves that of the multimarginal Sinkhorn algorithm in terms of $1/\varepsilon$ and accelerated alternating minimization algorithm (Tupitsa et al., 2020) in terms of n . The proof idea comes from a novel combination of Nesterov's estimated sequence and the techniques for analyzing the multimarginal Sinkhorn algorithm.

5.1 Algorithmic procedure

We present the pseudocode of accelerated multimarginal Sinkhorn algorithm in Algorithm 4. This algorithm achieves the acceleration by using Nesterov's estimate sequences (Nesterov, 2018). While our algorithm can be interpreted as an accelerated block coordinate descent algorithm, it is worthy noting that our algorithm is purely *deterministic* and thus differs from other accelerated randomized algorithms (Nesterov, 2012; Lin et al., 2015; Fercoq and Richtárik, 2015; Allen-Zhu et al., 2016; Lu et al., 2018; Diakonikolas and Orecchia, 2018) in the machine learning and optimization literature.

Comments on algorithmic scheme. Algorithm 4 is a novel combination of Nesterov's estimate sequences, a monotone search step, the choice of greedy coordinate and two coordinate updates. Nesterov's estimate sequences (**Step 1-3**) are crucial for optimizing a dual objective function φ faster than Algorithm 1. The coordinate update (**Step 4**) guarantees that $\varphi(\hat{\beta}^t) \leq \varphi(\tilde{\beta}^t)$ and $\|B(\hat{\beta}^t)\|_1 = 1$. The monotone search step (**Step 5**) guarantees that

Algorithm 5 Approximating MOT by Algorithms 2 and 4

Input: $\eta = \frac{\varepsilon}{2m \log(n)}$ and $\varepsilon' = \frac{\varepsilon}{8\|C\|_\infty}$.

Step 1: Let $\tilde{r}_k \in \Delta_n$ for $\forall k \in [m]$ be defined as

$$(\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_m) = \left(1 - \frac{\varepsilon'}{4m}\right) (r_1, r_2, \dots, r_m) + \frac{\varepsilon'}{4mn} (\mathbf{1}_n, \mathbf{1}_n, \dots, \mathbf{1}_n).$$

Step 2: Compute $\tilde{X} = \text{ACCELERATED MULTSINKHORN}(C, \eta, \{\tilde{r}_k\}_{k \in [m]}, \varepsilon'/2)$.

Step 3: Round $\hat{X} = \text{ROUND}(\tilde{X}, \{\tilde{r}_k\}_{k \in [m]})$.

Output: \hat{X} .

$\varphi(\beta^t) \leq \varphi(\hat{\beta}^t)$. The greedy coordinate update (**Step 6-7**) guarantees that $\varphi(\check{\beta}^{t+1}) \leq \varphi(\beta^t)$ with sufficiently large progress. Similar to Algorithm 1, the greedy rule is based on the function $\rho : \mathbb{R}_+^n \times \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ given by:

$$\rho(a, b) = \mathbf{1}_n^\top (b - a) + \sum_{i=1}^n a_i \log \left(\frac{a_i}{b_i} \right).$$

Furthermore, we also use the same quantity as that in the multimarginal Sinkhorn algorithm to measure the per-iteration residue of Algorithm 4:

$$E_t = \sum_{k=1}^m \|r_k(B(\beta^t) - r_k)\|_1. \quad (40)$$

Comments on arithmetic operations per iteration. The most expensive step is to compute $r_k(B(\bar{\beta}^t))/\|B(\bar{\beta}^t)\|_1$ for all $k \in [m]$. Since $B(\bar{\beta}^t)$ does not have any special property, it seems difficult to design some implementation trick to reduce the dependency on m . Thus, the arithmetic operations per iteration is still $O(mn^m)$. Note that, the accelerated alternating minimization algorithm in (Tupitsa et al., 2020) also requires $O(mn^m)$ arithmetic operations per iteration.

Algorithm for the MOT problem. We present the pseudocode of our main algorithm in Algorithm 5, where Algorithms 4 and 2 are the subroutines. The regularization parameter η is set as before, and **Step 1** is also necessary since the accelerated multimarginal Sinkhorn algorithm is not well behaved if the marginal distributions do not have dense support.

5.2 Technical lemmas

We first present two technical lemmas which are essential in the analysis of Algorithm 4. The first lemma provides an inductive relationship on the quantity

$$\delta_t = \varphi(\check{\beta}^t) - \varphi(\beta^*), \quad (41)$$

where β^* is an optimal solution of the dual entropic regularized MOT problem in Eq. (6). In order to facilitate the discussion, we recall Eq. (14) with $\|A\|_{1 \rightarrow 2} = \sqrt{m}$ as follows,

$$\varphi(\beta') - \varphi(\beta) - (\beta' - \beta)^\top \nabla \varphi(\beta) \leq \left(\frac{m}{2}\right) \|\beta' - \beta\|^2, \quad (42)$$

which will be used in the proof of the first lemma.

Lemma 19. *Let $\{\tilde{\beta}^t\}_{t \geq 0}$ be the iterates generated by Algorithm 4 and β^* be an optimal solution of the dual entropic regularized MOT problem. Then the quantity δ_t defined by Eq. (41) satisfies the following inequality,*

$$\delta_{t+1} \leq (1 - \theta_t)\delta_t + \frac{m\theta_t^2}{2} \left(\|\beta^* - \tilde{\beta}^t\|^2 - \|\beta^* - \tilde{\beta}^{t+1}\|^2 \right).$$

Proof Using Eq. (42) with $\beta' = \tilde{\beta}^t$ and $\beta = \tilde{\beta}^t$, we have

$$\varphi(\tilde{\beta}^t) \leq \varphi(\tilde{\beta}^t) + \theta_t(\tilde{\beta}^{t+1} - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t) + \left(\frac{m\theta_t^2}{2} \right) \|\tilde{\beta}^{t+1} - \tilde{\beta}^t\|^2.$$

By simple calculations, we find that

$$\begin{aligned} \varphi(\tilde{\beta}^t) &= (1 - \theta_t)\varphi(\tilde{\beta}^t) + \theta_t\varphi(\tilde{\beta}^t), \\ (\tilde{\beta}^{t+1} - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t) &= -(\tilde{\beta}^t - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t) + (\tilde{\beta}^{t+1} - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t). \end{aligned}$$

Putting these pieces together yields that

$$\begin{aligned} \varphi(\tilde{\beta}^t) &\leq \theta_t \underbrace{\left(\varphi(\tilde{\beta}^t) + (\tilde{\beta}^{t+1} - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t) + \left(\frac{m\theta_t}{2} \right) \|\tilde{\beta}^{t+1} - \tilde{\beta}^t\|^2 \right)}_{\text{I}} \\ &\quad + \underbrace{(1 - \theta_t)\varphi(\tilde{\beta}^t) - \theta_t(\tilde{\beta}^t - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t)}_{\text{II}}. \end{aligned} \tag{43}$$

We first estimate the term II. Indeed, it follows from the definition of $\tilde{\beta}^t$ that

$$-\theta_t(\tilde{\beta}^t - \tilde{\beta}^t) = \theta_t\tilde{\beta}^t + (1 - \theta_t)\tilde{\beta}^t - \tilde{\beta}^t = (1 - \theta_t)(\tilde{\beta}^t - \tilde{\beta}^t).$$

Using this equation and the convexity of φ , we have

$$\text{II} = (1 - \theta_t)(\varphi(\tilde{\beta}^t) + (\tilde{\beta}^t - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t)) \leq (1 - \theta_t)\varphi(\tilde{\beta}^t). \tag{44}$$

Then we proceed to estimate the term I. Indeed, by the update formula for $\tilde{\beta}^{t+1}$ and the definition of φ , we have

$$(\beta - \tilde{\beta}^{t+1})^\top (\nabla \varphi(\tilde{\beta}^t) + m\theta_t(\tilde{\beta}^{t+1} - \tilde{\beta}^t)) = 0 \text{ for all } \beta \in \mathbb{R}^{mn}.$$

Letting $\beta = \beta^*$ and rearranging the resulting equation yields that

$$(\tilde{\beta}^{t+1} - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t) = (\beta^* - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t) + \frac{m\theta_t}{2} \left(\|\beta^* - \tilde{\beta}^t\|^2 - \|\beta^* - \tilde{\beta}^{t+1}\|^2 - \|\tilde{\beta}^{t+1} - \tilde{\beta}^t\|^2 \right).$$

Using the convexity of φ again, we have $(\beta^* - \tilde{\beta}^t)^\top \nabla \varphi(\tilde{\beta}^t) \leq \varphi(\beta^*) - \varphi(\tilde{\beta}^t)$. Putting these pieces together yields that

$$\text{I} \leq \varphi(\beta^*) + \frac{m\theta_t}{2} \left(\|\beta^* - \tilde{\beta}^t\|^2 - \|\beta^* - \tilde{\beta}^{t+1}\|^2 \right). \tag{45}$$

Plugging Eq. (44) and Eq. (45) into Eq. (43) yields that

$$\varphi(\check{\beta}^t) \leq (1 - \theta_t)\varphi(\check{\beta}^t) + \theta_t\varphi(\beta^*) + \frac{m\theta_t^2}{2} \left(\|\beta^* - \tilde{\beta}^t\|^2 - \|\beta^* - \tilde{\beta}^{t+1}\|^2 \right).$$

Since $\check{\beta}^{t+1}$ is obtained by an coordinate update from β^t , we have $\varphi(\beta^t) \geq \varphi(\check{\beta}^{t+1})$. By the definition of β^t , we have $\varphi(\hat{\beta}^t) \geq \varphi(\beta^t)$. Since $\hat{\beta}^t$ is obtained by an coordinate update from $\check{\beta}^t$, we have $\varphi(\hat{\beta}^t) \geq \varphi(\check{\beta}^t)$. Putting these pieces together with yields that

$$\varphi(\check{\beta}^{t+1}) - \varphi(\beta^*) \leq (1 - \theta_t)(\varphi(\check{\beta}^t) - \varphi(\beta^*)) + \frac{m\theta_t^2}{2} \left(\|\beta^* - \tilde{\beta}^t\|^2 - \|\beta^* - \tilde{\beta}^{t+1}\|^2 \right).$$

This completes the proof. ■

The second lemma provides an upper bound for δ_t defined by Eq. (41) where $\{\check{\beta}^t\}_{t \geq 0}$ are generated by Algorithm 4 and β^* is an optimal solution defined by Corollary 4. Note that our lemma is a direct corollary of the analysis provided in Tseng (2008) and we provide the proof details for the sake of completeness.

Lemma 20. *Let $\{\check{\beta}^t\}_{t \geq 0}$ be the iterates generated by Algorithm 4 and β^* be an optimal solution of the dual entropic regularized MOT problem satisfying that $\|\beta^*\| \leq \sqrt{mn}R$ where R is defined in Corollary 4. Then the quantity δ_t defined by Eq. (41) satisfies the following inequality,*

$$\delta_t \leq \frac{2m^2nR^2}{(t+1)^2}.$$

Proof By simple calculations, we derive from the definition of θ_t that $(\theta_{t+1}/\theta_t)^2 = 1 - \theta_{t+1}$. Therefore, we conclude from Lemma 19 that

$$\left(\frac{1 - \theta_{t+1}}{\theta_{t+1}^2} \right) \delta_{t+1} - \left(\frac{1 - \theta_t}{\theta_t^2} \right) \delta_t \leq \frac{m}{2} \left(\|\beta^* - \tilde{\beta}^t\|^2 - \|\beta^* - \tilde{\beta}^{t+1}\|^2 \right).$$

Equivalently, we have

$$\left(\frac{1 - \theta_t}{\theta_t^2} \right) \delta_t + \left(\frac{m}{2} \right) \|\beta^* - \tilde{\beta}^t\|^2 \leq \left(\frac{1 - \theta_0}{\theta_0^2} \right) \delta_0 + \left(\frac{m}{2} \right) \|\beta^* - \tilde{\beta}^0\|^2.$$

Recall that $\theta_0 = 1$ and $\tilde{\beta}^0 = \mathbf{0}_{mn}$, we have $\delta_t \leq (m\theta_{t-1}^2/2)\|\beta^*\|^2 \leq (1/2)m^2nR^2\theta_{t-1}^2$. The remaining step is to show that $0 < \theta_t \leq 2/(t+2)$. Indeed, the claim holds when $t = 0$ as we have $\theta_0 = 1$. Assume that the claim holds for $t \leq t_0$, i.e., $\theta_{t_0} \leq 2/(t_0+2)$, we have

$$\theta_{t_0+1} = \frac{2}{1 + \sqrt{1 + 4/\theta_{t_0}^2}} \leq \frac{2}{t_0 + 3}.$$

Putting these pieces together yields the desired inequality for δ_t . ■

5.3 Main results

We present an upper bound for the number of iterations required by Algorithm 4.

Theorem 21. *Let $\{\beta^t\}_{t \geq 0}$ be the iterates generated by Algorithm 4. The number of iterations required to reach the stopping criterion $E_t \leq \varepsilon'$ satisfies*

$$t \leq 1 + 4 \left(\frac{\sqrt{nm}R}{\varepsilon'} \right)^{2/3},$$

where $R > 0$ is defined in Lemma 3.

Proof We first claim that

$$\varphi(\beta^t) - \varphi(\check{\beta}^{t+1}) \geq \frac{1}{2m} \left(\sum_{k=1}^m \|r_k(B(\beta^t)) - r_k\|_1^2 \right). \quad (46)$$

By the definition of φ , we have

$$\varphi(\beta^t) - \varphi(\check{\beta}^{t+1}) = \log(\|B(\beta^t)\|_1) - \log(\|B(\check{\beta}^{t+1})\|_1) - \sum_{k=1}^m (\beta_k^t - \check{\beta}_k^{t+1})^\top r_k. \quad (47)$$

From the update formula for $\hat{\beta}^t$ and $\check{\beta}^{t+1}$, it is clear that $\|B(\hat{\beta}^t)\|_1 = 1$ and $\|B(\check{\beta}^{t+1})\|_1 = 1$ for all $t \geq 0$. Then we derive from the monotone search step (cf. **Step 5**) that $\|B(\beta^t)\|_1 = 1$ for all $t \geq 1$. Therefore, we have

$$\varphi(\beta^t) - \varphi(\check{\beta}^{t+1}) = -(\beta_K^t - \check{\beta}_K^{t+1})^\top r_K = (\log(r_K) - \log(r_K(B(\beta^t))))^\top r_K.$$

Since $\mathbf{1}_n^\top r_K = \mathbf{1}_n^\top r_K(B(\beta^t)) = 1$, we have $\varphi(\beta^t) - \varphi(\check{\beta}^{t+1}) = \rho(r_K, r_K(B(\beta^t)))$ for all $t \geq 1$. Combining this inequality with the fact that the K -th coordinate is the greedy one yields

$$\varphi(\beta^t) - \varphi(\check{\beta}^{t+1}) \geq \frac{1}{m} \left(\sum_{k=1}^m \rho(r_k, r_k(B(\beta^t))) \right).$$

Using the Pinsker inequality (Cover and Thomas, 2012), we derive Eq. (46) as desired.

By the definition of β^t , we have $\varphi(\check{\beta}^t) \geq \varphi(\beta^t)$. Plugging this inequality into Eq. (46) together with the Cauchy-Schwarz inequality yields

$$\varphi(\check{\beta}^t) - \varphi(\check{\beta}^{t+1}) \geq \frac{1}{2m} \left(\sum_{k=1}^m \|r_k(B(\beta^t)) - r_k\|_1^2 \right) \geq \frac{1}{2} \left(\frac{E_t}{m} \right)^2.$$

Therefore, we conclude that

$$\varphi(\check{\beta}^j) - \varphi(\check{\beta}^{t+1}) \geq \frac{1}{2m^2} \left(\sum_{i=j}^t E_i^2 \right) \text{ for any } j \in \{1, 2, \dots, t\}.$$

Since $\varphi(\check{\beta}^{t+1}) \geq \varphi(\beta^*)$ for all $t \geq 1$, we have $\varphi(\check{\beta}^j) - \varphi(\check{\beta}^{t+1}) \leq \delta_j$. Then Lemma 20 implies

$$\sum_{i=j}^t E_i^2 \leq \frac{4m^4 n R^2}{(j+1)^2}.$$

Putting these pieces together with the fact that $E_t \geq \varepsilon'$ as soon as the stopping criterion is not fulfilled yields

$$\frac{4m^4 n R^2}{(j+1)^2 (t-j+1)} \geq (\varepsilon')^2.$$

Since this inequality holds true for all $j \in \{1, 2, \dots, t\}$, we assume without loss of generality that t is even and let $j = t/2$. Then, we obtain that

$$t \leq 1 + 4 \left(\frac{\sqrt{nm^2 R}}{\varepsilon'} \right)^{2/3}.$$

This completes the proof. ■

We are ready to present the complexity bound of Algorithm 5 for solving the MOT problem in Eq. (2). Note that $\varepsilon' = \varepsilon/(8\|C\|_\infty)$ is defined using the desired accuracy $\varepsilon > 0$.

Theorem 22. *Algorithm 5 returns an ε -approximate multimarginal transportation plan $\widehat{X} \in \mathbb{R}^{n \times \dots \times n}$ within*

$$O \left(\frac{m^3 n^{m+1/3} \|C\|_\infty^{4/3} (\log(n))^{1/3}}{\varepsilon^{4/3}} \right)$$

arithmetic operations.

Proof Applying the same argument which is used in Theorem 16, we obtain that $\langle C, \widehat{X} \rangle - \langle C, X^* \rangle \leq \varepsilon$ where \widehat{X} is returned by Algorithm 5.

It remains to bound the number of iterations required by Algorithm 4 to reach the criterion $E_t \leq \varepsilon'/2$ (cf. **Step 2** in Algorithm 5). Using Theorem 21, we have

$$t \leq 1 + 4 \left(\frac{\sqrt{nm} R}{\varepsilon'} \right)^{2/3}.$$

By the definition of R (cf. Lemma 3), $\eta = \varepsilon/(2m \log(n))$ and $\varepsilon' = \varepsilon/(8\|C\|_\infty)$, we have

$$\begin{aligned} t &\leq 1 + 4 \left(\frac{\sqrt{nm^2 R}}{\varepsilon'} \right)^{2/3} \\ &\leq 1 + 4 \left[\frac{8\sqrt{nm^2} \|C\|_\infty}{\varepsilon} \left(\frac{\|C\|_\infty}{\eta} + (m-1) \log(n) - 2 \log \left(\min_{1 \leq i \leq m, 1 \leq j \leq n} \tilde{r}_{ij} \right) \right) \right]^{2/3} \\ &\leq 1 + 4 \left[\frac{8\sqrt{nm^2} \|C\|_\infty}{\varepsilon} \left(\frac{2m \log(n) \|C\|_\infty}{\varepsilon} + (m-1) \log(n) - 2 \log \left(\frac{\varepsilon}{32mn \|C\|_\infty} \right) \right) \right]^{2/3} \\ &= O \left(\frac{m^2 n^{1/3} \|C\|_\infty^{4/3} (\log(n))^{1/3}}{\varepsilon^{4/3}} \right). \end{aligned}$$

Since each iteration of Algorithm 4 requires $O(mn^m)$ arithmetic operations, the total arithmetic operations required by **Step 2** of Algorithm 5 is $O(m^3 n^{m+1/3} \|C\|_\infty^{4/3} (\log(n))^{1/3} \varepsilon^{-4/3})$. In addition, computing a set of vectors $\{\tilde{r}_k\}_{k \in [m]}$ requires $O(mn)$ arithmetic operations and Algorithm 2 requires $O(mn^m)$ arithmetic operations. Putting these pieces together yields that the complexity bound of Algorithm 5 is $O(m^3 n^{m+1/3} \|C\|_\infty^{4/3} (\log(n))^{1/3} \varepsilon^{-4/3})$. ■

Remark 23. *Theorem 22 demonstrates that the complexity bound of Algorithm 5 is better than that of Algorithm 3 in terms of $1/\varepsilon$ but not near-linear in n^m . To be more specific, Algorithm 5 is recommended when $n \in (0, 1/\varepsilon^2)$. This occurs if the desired solution accuracy is relatively small, saying 10^{-4} , and the examples include the application problems from economics, physics and generalized Euler flows. In contrast, Algorithm 3 is recommended when $n \in (1/\varepsilon^2, +\infty)$. This occurs if the desired solution accuracy is relatively large, saying 10^{-2} , and the examples include the application problems from image processing.*

Remark 24. *The complexity bound has the same dependence m^3 as that of Algorithm 3. However, the improvement seems possible and can be achieved if we implement **Step 2** of Algorithm 4 in distributed parallel manner and choose the greedy coordinate in **Step 6** using the implementation trick we have mentioned before. Each iteration of Algorithm 4 requires $O(n^m)$ arithmetic operations and thus Algorithm 5 achieves the complexity bound of $O(m^2 n^{m+1/3} \|C\|_\infty^{4/3} (\log(n))^{1/3} \varepsilon^{-4/3})$. Further, it seems possible to improve the dependence of n by extending other algorithmic frameworks to the MOT setting (Blanchet et al., 2018; Lahn et al., 2019; Jambulapati et al., 2019). However, such extension is challenging since we are not clear whether these frameworks heavily depend on the minimum-cost flow structure of the OT problem or not. As such, we leave this topic to the future work.*

6. Experiments

In this section, we evaluate our new algorithms on both synthetic data and real images. In particular, we compute the free-support Wasserstein barycenter based on the OT distance with the quadratic Euclidean distance ground cost function and compare our algorithms with the commercial linear programming (LP) solver GUROBI. All the experiments are conducted in MATLAB R2020a on a workstation with an Intel Core i5-9400F (6 cores and 6 threads) and 32GB memory, equipped with Ubuntu 18.04.

6.1 Experiments on synthetic data

We follow the setup in Altschuler et al. (2017) in order to compare different algorithms on the synthetic images. More specifically, we generate a triple of random grayscale images, each normalized to have unit total mass. The marginals r_1, r_2 and r_3 represent three images, and the cost tensor C is generated by

$$C_{i_1, i_2, i_3} = \frac{1}{2} \left(\sum_{k=1}^3 \lambda_k \|x_{i_k} - A_{i_1, i_2, i_3}(x)\|^2 \right) \text{ for all } (i_1, i_2, i_3) \in [n] \times [n] \times [n],$$

where $A_{i_1, i_2, i_3}(x) = \sum_{k=1}^3 \lambda_k x_{i_k}$ is the Euclidean barycenter and $x = \{x_i\}_{i \in [n]} \subseteq \mathbb{R}^d$ are pixel locations in the images. Moreover, $\lambda = (\lambda_1, \lambda_2, \lambda_3) \in \Delta^3$ is a weight vector and set as $(1/3, 1/3, 1/3)$ consistently in this subsection.

Each of the images has n pixel locations in total and is generated based on randomly positioning a foreground square in otherwise black background. We utilize a uniform distribution on $[0, 1]$ for the intensities of the background pixels and a uniform distribution on $[0, 50]$ for the foreground pixels. We set the proportion of the size of the square is as 10% of the image and implement all the algorithms on the synthetic images with different size n .




Figure 1: Performance of multimarginal Sinkhorn v.s. accelerated multimarginal Sinkhorn on the randomly generated synthetic images. Number of pixel locations in each synthetic image is set as $n = 25$ (top) and $n = 100$ (bottom).

We generalize two metrics proposed by Altschuler et al. (2017) and use them to quantitatively measure the performance of different algorithms. The first metric is the distance between the output of the algorithm, X , and the transportation polytope between the marginals r_1 , r_2 and r_3 . Formally, we have

$$d(X) = \|r_1(X) - r_1\|_1 + \|r_2(X) - r_2\|_1 + \|r_3(X) - r_3\|_1,$$

where $r_1(X)$, $r_2(X)$ and $r_3(X)$ are the marginal vectors of the output X while r_1 , r_2 and r_3 stand for the true marginal vectors. The second metric is the competitive ratio, defined by $\log(d(X_1)/d(X_2))$ where $d(X_1)$ and $d(X_2)$ refer to the distance between the outputs of two algorithms and the transportation polytope.

We perform a pairwise comparative experiment: multimarginal Sinkhorn versus accelerated multimarginal Sinkhorn, by running both algorithms with ten randomly selected pairs of synthetic images with varying size $n \in \{25, 100\}$. In order to have further evaluations with these algorithms, we also compare their performance with different choices of regularization parameter $\eta \in \{1, 0.2, 0.1\}$ while using the value of the MOT problem (without entropic regularization term) as the baseline. The maximum number of iterations is set as 10.

Experimental results. Figure 1 summarizes the results on synthetic images. The images in the first row show the comparative performance of both algorithms in terms of the iteration counts on 10 triples of 5×5 synthetic images. In the leftmost one, the comparison




Figure 2: Computational efficiency of Gurobi v.s. our algorithms as n varies.

uses distance to transportation polytope $d(X)$ where X are returned by the algorithms. In the middle one, the maximum/median/minimum values of the competitive ratios are utilized for the comparison. In the rightmost one, we vary the regularization parameter $\eta \in \{1, 0.2, 0.1\}$ for both algorithms together with the value of the unregularized MOT problem as the baseline. It is clear that accelerated multimarginal Sinkhorn algorithm outperforms multimarginal Sinkhorn algorithm in terms of iteration numbers, illustrating the improvement achieved by using the *estimated sequence* and *monotone search*.

To further compare our algorithms with Gurobi in terms of computational efficiency, we conduct one more experiment with varying number of support points (or pixel locations) $n \in \{25, 100, 144\}$. Figure 2 shows the running time taken by three algorithms across a wide range of n . As n increases, we find that multimarginal Sinkhorn algorithm performs the best, followed by accelerated multimarginal Sinkhorn algorithm, both outperforming Gurobi. This demonstrates that classical LP algorithms might not be suitable for solving the MOT problem, partially confirming our results in Section 3. Moreover, despite fewer iterations, the direct implementation of accelerated multimarginal Sinkhorn algorithm is indeed slower than multimarginal Sinkhorn algorithm. This is mainly due to the heavy computation of gradient and we believe some parallel computing toolbox can be helpful. However, this is beyond the scope of this paper and we leave it to future research.

6.2 Experiments on real images

We conduct the experiment with the same setup and MNIST dataset. The MNIST dataset consists of 60,000 images of handwritten digits of size 28 by 28 pixels. We add a very small noise term (10^{-6}) to all the zero elements in the measures and then normalize them such that their sum becomes one. We also vary the regularization parameter $\eta \in \{1, 0.05, 0.02\}$

. Available in <http://yann.lecun.com/exdb/mnist/>




Figure 3: Performance of multimarginal Sinkhorn v.s. accelerated multimarginal Sinkhorn on MNIST images. Number of pixel locations in each MNIST image is set as $n = 576$.

for both algorithms but cannot run GUROBI. Indeed, the LP constructed from the MOT problem using 3 MNIST images is so large that GUROBI is out of memory. Figure 3 presents the comparative performance of our algorithms on the MNIST images, and we find that it is consistent with the performance on the randomly generated synthetic images.

In order to better visualize the quality of approximate barycenters obtained by each algorithm, we run our algorithms with $\eta = 0.05$ to compute the free-support Wasserstein barycenter of two triple of real images with different weight vectors. Indeed, we solve the MOT problem as before and form the barycenter as follows,

$$\mu_\lambda = \sum_{k=1}^3 \sum_{1 \leq i_k \leq n} \gamma_{i_1, i_2, i_3} \delta_{A_{i_1, i_2, i_3}}(x),$$

where $A_{i_1, i_2, i_3}(x) = \sum_{k=1}^3 \lambda_k x_{i_k}$ is the Euclidean barycenter, and $x = \{x_i\}_{i \in [n]} \subseteq \mathbb{R}^d$ are pixel locations in the images and $\gamma \in \mathbb{R}^{n \times n \times n}$ is an optimal multimarginal transportation plan that solves the MOT problem.

Figure 4 presents the approximate barycenters obtained by running our algorithms. These results demonstrate that our algorithms can successfully capture the free-support barycenters of high quality by solving the MOT problem and are at least competitive with the existing algorithms (Benamou et al., 2015, 2019; Peyré and Cuturi, 2019) in practice.

7. Conclusion

In this paper, we have studied the multimarginal optimal transport (MOT) problem, providing new algorithms and complexity bounds for approximating this problem. We demonstrated that the standard linear programming (LP) form of the MOT problem is not a minimum-cost flow problem when $m \geq 3$. This encourages us to study the alternatives to combinatorial algorithms and standard *deterministic* interior-point algorithms. In particular, we considered an entropic regularized version of the MOT problem, developing two *deterministic* algorithms — the multimarginal Sinkhorn and accelerated multimarginal Sinkhorn algorithms — for solving it. Combined with a new rounding scheme, the multimarginal Sinkhorn algorithm can solve the MOT problem and achieves a near-linear time




Figure 4: Approximate barycenters obtained by running the multimarginal Sinkhorn (top) and accelerated multimarginal Sinkhorn (bottom) algorithms.

complexity bound of $O(m^3 n^m \|C\|_\infty^2 \log(n) \varepsilon^{-2})$. For the accelerated multimarginal Sinkhorn algorithm, the complexity bound is $O(m^3 n^{m+1/3} \|C\|_\infty^{4/3} (\log(n))^{1/3} \varepsilon^{-4/3})$ which is not near-linear in the number of variables n^m but has better dependence on $1/\varepsilon$ than that of the multimarginal Sinkhorn algorithm.

We now discuss a few directions that arise naturally from our work. First, the complexity bounds of the proposed algorithms in this paper do not incorporate low-rank approximation framework for the cost tensor C . Intuitively, these low-rank approaches will lead to an improvement of these complexity bounds in terms of the number of support points n . Therefore, with the low-rank approaches, the implementation of these algorithms will be feasible under the large-scale settings of the MOT problem. Second, as mentioned in the paper, one drawback of the entropic regularization is that the sparsity of the solution is lost.

Even though an ε -approximate transportation plan can be obtained efficiently, it is not clear how different the resulting sparsity pattern of the obtained solution is with respect to the solution of the actual MOT problem. An important direction is to incorporate sparsity penalty functions to the entropic regularized MOT problem such that an ε -approximate sparse transportation plan is achieved. Third, the MOT problem suffers from curse of dimensionality, demonstrating the importance of efficient dimension reduction frameworks in both theory and practice. Finally, it is of interest to extend the current algorithms in the paper to the multimarginal optimal transport among general measures, which are not necessarily probability measures, such as multimarginal unbalanced optimal transport (Pham et al., 2020) or multimarginal partial optimal transport (Le et al., 2022).

8. Acknowledgements

This work was supported in part by the Mathematical Data Science program of the Office of Naval Research under grant number N00014-18-1-2764 to MJ, and by the NSF IFML 2019844 award and research gifts by UT Austin ML grant to NH.

References

- M. Agueh and G. Carlier. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- Z. Allen-Zhu, Z. Qu, P. Richtárik, and Y. Yuan. Even faster accelerated coordinate descent using non-uniform sampling. In *ICML*, pages 1110–1119, 2016.
- J. Altschuler, J. Weed, and P. Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *NeurIPS*, pages 1964–1974, 2017.
- J. M. Altschuler and E. Boix-Adsera. Hardness results for multimarginal optimal transport problems. *Discrete Optimization*, 42:100669, 2021a.
- J. M. Altschuler and E. Boix-Adsera. Wasserstein barycenters can be computed in polynomial time in fixed dimension. *Journal of Machine Learning Research*, 22:1–19, 2021b.
- J. M. Altschuler and E. Boix-Adserà. Wasserstein barycenters are NP-hard to compute. *SIAM Journal on Mathematics of Data Science*, 4(1):179–203, 2022.
- E. Anderes, S. Borgwardt, and J. Miller. Discrete Wasserstein barycenters: Optimal transport for discrete data. *Mathematical Methods of Operations Research*, 84(2):389–409, 2016.
- J-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- J-D. Benamou, G. Carlier, and L. Nenna. Generalized incompressible flows, multi-marginal transport and Sinkhorn algorithm. *Numerische Mathematik*, 142(1):33–54, 2019.
- C. Berge. *The Theory of Graphs*. Courier Corporation, 2001.

- J. Blanchet, A. Jambulapati, C. Kent, and A. Sidford. Towards optimal running times for optimal transport. *ArXiv Preprint: 1810.07717*, 2018.
- Y. Brenier. The least action principle and the related concept of generalized flows for incompressible perfect fluids. *Journal of the American Mathematical Society*, 2(2):225–255, 1989.
- Y. Brenier. Minimal geodesics on groups of volume-preserving maps and generalized solutions of the Euler equations. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 52(4):411–452, 1999.
- Y. Brenier. Generalized solutions and hydrostatic approximation of the Euler equations. *Physica D: Nonlinear Phenomena*, 237(14-17):1982–1988, 2008.
- G. Buttazzo, L. D. Pascale, and P. Gori-Giorgi. Optimal-transport formulation of electronic density-functional theory. *Physical Review A*, 85(6), 2012.
- J. Cao, L. Mo, Y. Zhang, K. Jia, C. Shen, and M. Tan. Multi-marginal Wasserstein GAN. In *NeurIPS*, pages 1776–1786, 2019.
- G. Carlier and I. Ekeland. Matching for teams. *Economic Theory*, 42(2):397–418, 2010a.
- G. Carlier and I. Ekeland. Hedonic price equilibria, stable matching and optimal transport: Equivalence, topology and uniqueness. *Economic Theory*, 42(2):317–354, 2010b.
- G. Carlier, A. Oberman, and E. Oudet. Numerical methods for matching for teams and Wasserstein barycenters. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1621–1642, 2015.
- Y. Choi, M. Choi, M. Kim, J-W. Ha, S. Kim, and J. Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, pages 8789–8797, 2018.
- S. Clatici, E. Chien, and J. Solomon. Stochastic Wasserstein barycenters. In *ICML*, pages 999–1008. PMLR, 2018.
- M. B. Cohen, Y. T. Lee, and Z. Song. Solving linear programs in the current matrix multiplication time. In *STOC*, pages 938–942, 2019.
- C. Cotar, G. Friesecke, and C. Klüppelberg. Density functional theory and optimal transportation with Coulomb cost. *Communications on Pure and Applied Mathematics*, 66(4): 548–599, 2013.
- T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.
- M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NeurIPS*, pages 2292–2300, 2013.
- M. Cuturi and A. Doucet. Fast computation of Wasserstein barycenters. In *ICML*, pages 685–693, 2014.

- M. Cuturi and G. Peyré. Semidual regularized optimal transport. *SIAM Review*, 60(4): 941–965, 2018.
- S. I. Daitch and D. A. Spielman. Faster approximate lossy generalized flow via interior point algorithms. In *Foundations of Computer Science*, pages 451–460. ACM, 2008.
- I. S. Dhillon, P. K. Ravikumar, and A. Tewari. Nearest neighbor based greedy coordinate descent. In *NeurIPS*, pages 2160–2168, 2011.
- J. Diakonikolas and L. Orecchia. Alternating randomized block coordinate descent. In *ICML*, pages 1224–1232. PMLR, 2018.
- Y. Dolinsky and M. H. Soner. Robust hedging and martingale optimal transport in continuous time. *Probability Theory and Related Fields*, 160:391–427, 2014.
- R. M. Dudley. The speed of mean Glivenko-Cantelli convergence. *The Annals of Mathematical Statistics*, 40(1):40–50, 1969.
- P. Dvurechensky, D. Dvinskikh, A. Gasnikov, C. A. Uribe, and A. Nedić. Decentralize and randomize: Faster algorithm for Wasserstein barycenters. In *NeurIPS*, pages 10760–10770, 2018a.
- P. Dvurechensky, A. Gasnikov, and A. Kroshnin. Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn’s algorithm. In *ICML*, pages 1367–1376, 2018b.
- J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.
- I. Ekeland. An optimal matching problem. *ESAIM: Control, Optimisation and Calculus of Variations*, 11(1):57–71, 2005.
- T. R. Ervolina and S. T. McCormick. Canceling most helpful total cuts for minimum cost network flow. *Networks*, 23(1):41–52, 1993a.
- T. R. Ervolina and S. T. McCormick. Two strongly polynomial cut cancelling algorithms for minimum cost network flow. *Discrete Applied Mathematics*, 46(2):133–165, 1993b.
- O. Fercoq and P. Richtárik. Accelerated, parallel, and proximal coordinate descent. *SIAM Journal on Optimization*, 25(4):1997–2023, 2015.
- R. Flamary and N. Courty. POT: Python optimal transport library, 2017. URL <https://pythonot.github.io/>.
- N. Fournier and A. Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3-4):707–738, 2015.
- A. Galichon, P. Henry-Labordere, and N. Touz. A stochastic control approach to non-arbitrage bounds given marginals, with an application to Lookback options. *The Annals of Applied Probability*, 24:312–336, 2014.

- Z. Galil and É. Tardos. An $o(n^2(m+n\log n)\log n)$ min-cost flow algorithm. *Journal of the ACM (JACM)*, 35(2):374–386, 1988.
- W. Gangbo and A. Swiech. Optimal maps for the multidimensional Monge-Kantorovich problem. *Communications on Pure and Applied Mathematics*, 51(1):23–45, 1998.
- M. R. Garey and D. S. Johnson. *Computers and Intractability*, volume 29. WH Freeman New York, 2002.
- D. Ge, H. Wang, Z. Xiong, and Y. Ye. Interior-point methods strike back: Solving the Wasserstein barycenter problem. In *NeurIPS*, pages 6894–6905, 2019.
- A. Genevay, L. Chizat, F. Bach, M. Cuturi, and G. Peyré. Sample complexity of Sinkhorn divergences. In *AISTATS*, 2019.
- A. Ghouila-Houri. Caractérisation des matrices totalement unimodulaires. *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences (Paris)*, 254:1192–1194, 1962.
- A. V. Goldberg and S. Rao. Beyond the flow decomposition barrier. *Journal of the ACM (JACM)*, 45(5):783–797, 1998.
- A. V. Goldberg and R. E. Tarjan. Finding minimum-cost circulations by successive approximation. *Mathematics of Operations Research*, 15(3):430–466, 1990.
- S. Guminov, P. Dvurechensky, N. Tupitsa, and A. Gasnikov. Accelerated alternating minimization, accelerated Sinkhorn’s algorithm and accelerated iterative Bregman projections. *ArXiv Preprint: 1906.03622*, 2019.
- R. Hassin. The minimum cost flow problem: a unifying approach to dual algorithms and a new tree-search algorithm. *Mathematical Programming*, 25(2):228–239, 1983.
- R. Hassin. Algorithms for the minimum cost circulation problem based on maximizing the mean improvement. *Operations Research Letters*, 12(4):227–233, 1992.
- Z. He, W. Zuo, M. Kan, S. Shan, and X. Chen. Attgan: Facial attribute editing by only changing what you want. *IEEE Transactions on Image Processing*, 28(11):5464–5478, 2019.
- L. Hui, X. Li, J. Chen, H. He, and J. Yang. Unsupervised multi-domain image translation with domain-specific encoders/decoders. In *ICPR*, pages 2044–2049. IEEE, 2018.
- A. Jambulapati, A. Sidford, and K. Tian. A direct $\tilde{O}(1/\epsilon)$ iteration parallel algorithm for optimal transport. In *NeurIPS*, pages 11355–11366, 2019.
- L. V. Kantorovich. On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)*, volume 37, pages 199–201, 1942.
- R. M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer, 1972.

- M. Klein. A primal method for minimal cost flows with applications to the assignment and transportation problems. *Management Science*, 14(3):205–220, 1967.
- A. Kroshnin, N. Tupitsa, D. Dvinskikh, P. Dvurechensky, A. Gasnikov, and C. Uribe. On the complexity of approximating Wasserstein barycenters. In *ICML*, pages 3530–3540, 2019.
- T. Lacombe, M. Cuturi, and S. Oudot. Large scale computation of means and clusters for persistence diagrams using optimal transport. In *NeurIPS*, 2018.
- N. Lahn, D. Mulchandani, and S. Raghvendra. A graph theoretic additive approximation of optimal transport. In *NeurIPS*, pages 13836–13846, 2019.
- K. Le, H. Nguyen, K. Nguyen, T. Pham, and N. Ho. On multimarginal partial optimal transport: Equivalent forms and computational complexity. In *AISTATS*, 2022.
- Y. T. Lee and A. Sidford. Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\sqrt{\text{rank}})$ iterations and faster algorithms for maximum flow. In *Foundations of Computer Science*, pages 424–433. IEEE, 2014.
- J. Lei. Convergence and concentration of empirical measures under Wasserstein distance in unbounded functional spaces. *Bernoulli*, 26(1):767–798, 2020.
- Q. Lin, Z. Lu, and L. Xiao. An accelerated randomized proximal coordinate gradient method and its application to regularized empirical risk minimization. *SIAM Journal on Optimization*, 25(4):2244–2273, 2015.
- T. Lin, N. Ho, and M. Jordan. On efficient optimal transport: An analysis of greedy and accelerated mirror descent algorithms. In *ICML*, pages 3982–3991, 2019a.
- T. Lin, N. Ho, and M. I. Jordan. On the efficiency of the Sinkhorn and Greenhorn algorithms and their acceleration for optimal transport. *ArXiv Preprint: 1906.01437*, 2019b.
- T. Lin, N. Ho, X. Chen, M. Cuturi, and M. I. Jordan. Fixed-support Wasserstein barycenters: Computational hardness and fast algorithm. In *NeurIPS*, pages 5368–5380, 2020.
- H. Lu, R. Freund, and V. Mirrokni. Accelerating greedy coordinate descent methods. In *ICML*, pages 3257–3266, 2018.
- G. Mena and J. Niles-Weed. Statistical bounds for entropic optimal transport: Sample complexity and the central limit theorem. In *NeurIPS*, pages 4541–4551, 2019.
- C. B. Mendl and L. Lin. Kantorovich dual solution for strictly correlated electrons in atoms and molecules. *Physical Review B*, 87:125106, 2013.
- O. Meshi, A. Globerson, and T. S. Jaakkola. Convergence rate analysis of MAP coordinate minimization algorithms. In *NeurIPS*, pages 3014–3022, 2012.
- L. Mi and J. Bento. Multi-marginal optimal transport defines a generalized metric. *ArXiv Preprint: 2001.11114*, 2020.

- Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- Y. Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- Y. Nesterov. *Lectures on Convex Optimization*, volume 137. Springer, 2018.
- J. Nutini, M. Schmidt, I. Laradji, M. Friedlander, and H. Koepke. Coordinate descent converges faster with the Gauss-Southwell rule than random selection. In *ICML*, pages 1632–1641, 2015.
- J. B. Orlin. A faster strongly polynomial minimum cost flow algorithm. *Operations Research*, 41(2):338–350, 1993.
- J. B. Orlin. A polynomial time primal network simplex algorithm for minimum cost flows. *Mathematical Programming*, 78(2):109–129, 1997.
- B. Pass. Multi-marginal optimal transport: Theory and applications. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1771–1790, 2015.
- G. Peyré and M. Cuturi. *Computational Optimal Transport: With Applications to Data Science*. Foundations and Trends(r) in Machine Learning, 2019.
- K. Pham, K. Le, N. Ho, T. Pham, and H. Bui. On unbalanced optimal transport: An analysis of Sinkhorn algorithm. In *ICML*, pages 7673–7682. PMLR, 2020.
- A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24. Springer Science & Business Media, 2003.
- M. Seidl, P. Gori-Giorgi, and A. Savi. Strictly correlated electrons in density-functional theory: A general formulation with applications to spherical densities. *Physical Review A*, 75:75:042511, 2007.
- S. Srivastava, C. Li, and D. Dunson. Scalable Bayes via barycenter in Wasserstein space. *Journal of Machine Learning Research*, 19(8):1–35, 2018.
- M. Staib, S. Clatici, J. M. Solomon, and S. Jegelka. Parallel streaming Wasserstein barycenters. In *NeurIPS*, pages 2647–2658, 2017.
- É. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5(3):247–255, 1985.
- R. E. Tarjan. Dynamic trees as search trees via euler tours, applied to the network simplex algorithm. *Mathematical Programming*, 78(2):169–177, 1997.
- P. Tseng. On accelerated proximal gradient methods for convex-concave optimization. *submitted to SIAM Journal on Optimization*, 2(3), 2008.
- N. Tupitsa, P. Dvurechensky, A. Gasnikov, and C. A. Uribe. Multimarginal optimal transport by accelerated alternating minimization. In *CDC*, pages 6132–6137. IEEE, 2020.

- C. Villani. *Topics in Optimal Transportation*. American Mathematical Society, Providence, RI, 2003.
- J. Weed and F. Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019.
- S. J. Wright. *Primal-Dual Interior-Point Methods*, volume 54. SIAM, 1997.
- Y. Xie, X. Wang, R. Wang, and H. Zha. A fast proximal point method for computing exact Wasserstein distance. In *UAI*, pages 433–453. PMLR, 2020.