

On Unbalanced Optimal Transport: Gradient Methods, Sparsity and Approximation Error

Quang Minh Nguyen

NMQUANG@MIT.EDU

*Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA, USA*

Hoang H. Nguyen

HNGUYEN455@GATECH.EDU

*Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA, USA*

Yi Zhou

YI.ZHOU@IBM.COM

*IBM Research
Almaden Research Center
San Jose, CA, USA*

Lam M. Nguyen

LAMNGUYEN.MLTD@IBM.COM

*IBM Research
Thomas J. Watson Research Center
Yorktown Heights, NY, USA*

Editor: Silvia Villa

Abstract

We study the Unbalanced Optimal Transport (UOT) between two measures of possibly different masses with at most n components, where the marginal constraints of standard Optimal Transport (OT) are relaxed via Kullback-Leibler divergence with regularization factor τ . Although only Sinkhorn-based UOT solvers have been analyzed in the literature with the iteration complexity of $O(\frac{\tau \log(n)}{\varepsilon} \log(\frac{\log(n)}{\varepsilon}))$ and per-iteration cost of $O(n^2)$ for achieving the desired error ε , their positively dense output transportation plans strongly hinder the practicality. On the other hand, while being vastly used as heuristics for computing UOT in modern deep learning applications and having shown success in sparse OT problem, gradient methods applied to UOT have not been formally studied. In this paper, we propose a novel algorithm based on Gradient Extrapolation Method (GEM-UOT) to find an ε -approximate solution to the UOT problem in $O(\kappa \log(\frac{\tau n}{\varepsilon}))$ iterations with $\tilde{O}(n^2)$ per-iteration cost, where κ is the condition number depending on only the two input measures. Our proof technique is based on a novel dual formulation of the squared ℓ_2 -norm UOT objective, which fills the lack of sparse UOT literature and also leads to a new characterization of approximation error between UOT and OT. To this end, we further present a novel approach of OT retrieval from UOT, which is based on GEM-UOT with fine tuned τ and a post-process projection step. Extensive experiments on synthetic and real datasets validate our theories and demonstrate the favorable performance of our methods in practice. We showcase GEM-UOT on the task of color transfer in terms of both the quality of the transfer image and the sparsity of the transportation plan.

Keywords: Unbalanced Optimal Transport; Gradient Methods; Convex Optimization

1. Introduction

The optimal transport (OT) problem originated from the need to find the optimal cost to transport masses from one distribution to another distribution (Villani, 2008). While initially developed by theorists, OT has found widespread applications in statistics and machine learning (ML) (see e.g. (Ho et al., 2017; Arjovsky et al., 2017; Rabin and Papadakis, 2015)). However, standard OT requires the restricted assumption that the input measures are normalized to unit mass, which facilitates the development of the Unbalanced Optimal Transport (UOT) problem between two measures of possibly different masses (Chizat et al., 2018). The class of UOT problem relaxes OT’s marginal constraints. Specifically, it is a regularized version of Kantorovich formulation placing penalty functions on the marginal distributions based on some given divergences (Liero et al., 2017). While there have been several divergences considered by the literature, such as squared ℓ_2 norm (Blondel et al., 2018), ℓ_1 norm (Caffarelli and McCann, 2010), or general ℓ_p norm (Lee et al., 2019), UOT with Kullback-Leiber (KL) divergence (Chizat et al., 2018) is the most prominent and has been used in statistics and machine learning (Frogner et al., 2015), deep learning (Yang and Uhler, 2019), domain adaptation (Balaji et al., 2020; Fatras et al., 2021), bioinformatics (Schiebinger et al., 2019), and OT robustness (Balaji et al., 2020; Le et al., 2021). Throughout this paper, we refer to UOT penalized by KL divergence as simply UOT, unless otherwise specified. We hereby define some notations and formally present our problem of interest.

Notations: We let $[n]$ stand for the set $\{1, 2, \dots, n\}$ while \mathbb{R}_+^n stands for the set of all vectors in \mathbb{R}^n with nonnegative entries. For a vector $\mathbf{x} \in \mathbb{R}^n$ and $p \in [1, \infty)$, we denote $\|\mathbf{x}\|_p$ as its ℓ_p -norm and $diag(\mathbf{x}) \in \mathbb{R}^{n \times n}$ as the diagonal matrix with $diag(\mathbf{x})_{ii} = x_i$. Let A and B be two matrices of size $n \times n$, we denote their Frobenius inner product as: $\langle A, B \rangle = \sum_{i,j=1}^n A_{ij}B_{ij}$. $\mathbf{1}_n$ stands for a vector of length n with all of its components equal to 1. The KL divergence between two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ is defined as $\mathbf{KL}(x||y) = \sum_{i=1}^n x_i \log\left(\frac{x_i}{y_i}\right) - x_i + y_i$. The entropy of a matrix $X \in \mathbb{R}_+^{n \times n}$ is given by $H(X) = -\sum_{i,j=1}^n X_{ij}(\log(X_{ij}) - 1)$.

Unbalanced Optimal Transport: For a couple of finite measures with possibly different total mass $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{R}_+^n$ and $\mathbf{b} = (b_1, \dots, b_n) \in \mathbb{R}_+^n$, we denote $\alpha = \sum_{i=1}^n a_i$ and $\beta = \sum_{i=1}^n b_i$ as the total masses, and $a_{min} = \min_{1 \leq i \leq n} \{a_i\}$ and $b_{min} = \min_{1 \leq i \leq n} \{b_i\}$ as the minimum masses. The UOT problem can be written as:

$$\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) = \min_{X \in \mathbb{R}_+^{n \times n}} \{f(X) := \langle C, X \rangle + \tau \mathbf{KL}(X \mathbf{1}_n || \mathbf{a}) + \tau \mathbf{KL}(X^\top \mathbf{1}_n || \mathbf{b})\}, \quad (1)$$

where C is a given cost matrix, X is a transportation plan, $\tau > 0$ is a given regularization parameter. When $\mathbf{a}^\top \mathbf{1}_n = \mathbf{b}^\top \mathbf{1}_n$ and $\tau \rightarrow \infty$, (1) reduces to a standard OT problem. Let $X_f \in \operatorname{argmin}_{X \in \mathbb{R}_+^{n \times n}} f(X)$ be optimal transportation plan of the UOT problem (1).

Definition 1 For $\varepsilon > 0$, X is an ε -approximate transportation plan of $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ if:

$$f(X) = \langle C, X \rangle + \tau \mathbf{KL}(X \mathbf{1}_n || \mathbf{a}) + \tau \mathbf{KL}(X^\top \mathbf{1}_n || \mathbf{b}) \leq \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) + \varepsilon. \quad (2)$$

1.1 Open Problems and Motivation

We hereby discuss the open problems within the UOT problem to better facilitate the motivation underpinning our work.

Nascent literature of gradient methods for UOT: While the existing work has well studied the class of OT problems using gradient methods or variations of the Sinkhorn algorithm (Guminov et al., 2021; Altschuler et al., 2017), the complexity theory and algorithms for UOT remain nascent despite its recent emergence. Pham et al. (2020) shows that Sinkhorn algorithm can solve the UOT problem in $O\left(\frac{\tau n^2 \log(n)}{\varepsilon} \log\left(\frac{\log(n)}{\varepsilon}\right)\right)$ up to the desired error ε . Chapel et al. (2021) proposes Majorization-Minimization (MM) algorithm, which is Sinkhorn-based, as a numerical solver with $O(n^{3.27})$ empirical complexity. Thanks to their compatibility, gradient methods have been vastly used as heuristics for computing UOT in deep learning applications (Balaji et al., 2020; Yang and Uhler, 2019) and multi-label learning (Frogner et al., 2015). Nevertheless, no work has formally studied gradient methods applied to UOT problems. Principled study of this topic could lead to preliminary theoretical justification and guide the more refined design of gradient methods for UOT.

Lack of sparse UOT literature: There is currently no study of sparse UOT in the literature, which nullifies its usage in many applications (Pitié et al., 2007; Courty et al., 2017; Muzellec et al., 2016) in which only sparse transportation plan is of interest. On a side note, entropic regularization imposed by the Sinkhorn algorithm keeps the transportation plan dense and strictly positive (Essid and Solomon, 2018; Blondel et al., 2018).

Approximation of OT: The approximability of standard OT via UOT has remained an open problem. Recently, (Chapel et al., 2021) discusses the possibility to approximate OT transportation plan using that of UOT and empirically verifies it. Despite the well known fact that OT is recovered from UOT when $\tau \rightarrow \infty$ and the masses are balanced, no work has analyzed the rate under which UOT converges to OT¹. Characterization of approximation error between UOT and OT could give rise to the new perspective of OT retrieval from UOT, and facilitate further study of the UOT geometry and statistical bounds from the lens of OT (Bernton et al., 2022) in the regime of large τ .

Bottleneck of robust OT computation via UOT: The UOT formulation with small τ has been vastly used as a robust variant of OT² (Fattras et al., 2021; Balaji et al., 2020; Le et al., 2021). In this paper, we further highlight a major bottleneck of this approach having been neglected by the literature: though relaxing OT as UOT admits robustness to outliers, the computed UOT distance far deviates from the original OT distance. Under the naive choice of $\tau = 1$ (Fattras et al., 2021; Balaji et al., 2020; Le et al., 2021), we observe that on CIFAR-10 as an example real dataset the UOT distance differs from OT by order of dozens (see Figure 14). Such large deviation from the original metric can be detrimental to target applications where the distance itself is of interest (Genevay et al., 2018). Another inherent limitation is that the solution returned by UOT is only a “relaxed” transportation plan that does not respect the marginal constraints of standard OT. This naturally raises the question whether despite the relaxation of OT as UOT, there is a way to retrieve the original OT distance and transportation plan.

1. For squared ℓ_2 norm penalized UOT variant, (Blondel et al., 2018) showed that error in terms of transport distance is $O(n^2/\tau)$.

2. While Fattras et al. (2021) uses exactly UOT as robust OT, Le et al. (2021) considers UOT with an additional normalization constraint, i.e. $\|X\|_1 = 1$, which can be solved with usual UOT algorithm with an extra normalization step.

1.2 Contributions

In this paper, we provide a comprehensive study of UOT that addresses all the aforementioned open problems and challenges of UOT literature. We consider the setting where KL divergences are used to penalize the marginal constraints. Our contributions can be summarized as follows:

- We provide a novel dual formulation of the squared ℓ_2 -norm regularized UOT objective, which is the basis of our algorithms. This could facilitate further algorithmic development for solving sparse or standard UOT, since the current literature is limited to the dual formulation of entropic-regularized UOT problem (Chizat et al., 2018).
- Based on the Gradient Extrapolation Method (GEM), we propose GEM-UOT algorithm for finding an ε -approximate solution to the UOT problem in $O(\kappa \log(\frac{\tau n}{\varepsilon}))$ iterations with $\tilde{O}(n^2)$ per-iteration cost, where κ is the condition number that depends on only the two input measures. Through GEM-UOT, we present the first i) principled study of gradient methods for UOT, ii) sparse UOT solver, and iii) algorithm that lifts the linear dependence on τ of Sinkhorn (Pham et al., 2020), a bottleneck in the regime of large τ (Sejourne et al., 2022).
- To the best of our knowledge, we establish the first characterization of the approximation error between UOT and OT (in the context of KL-penalized UOT). In particular, we show that both of UOT’s transportation plan and transport distance converge to OT’s marginal constraints and transport distance with the rate $O(\frac{\text{poly}(n)}{\tau})$. This result opens up directions that use UOT to approximate standard OT and accentuate the importance of our proposed GEM-UOT, which is the first to achieve logarithmic dependence on τ in the literature.
- Inspired by our results on approximation error, we bring up a new perspective of OT retrieval from UOT with fine-tuned τ . In particular, we present GEM-OT, which obtains an ε -approximate solution to the standard OT problem by performing a post-process projection of UOT solution. GEM-OT is of order $\tilde{O}(\kappa n^2)$. While lots of work have naively relaxed OT into UOT to enjoy its unconstrained optimization structure and robustness to outliers, our results on τ make a further significant step by equipping such relaxation with provable guarantees.

Paper Organization: The rest of the paper is organized as follows. We introduce the background of regularized UOT problems and present our novel dual formulation in Section 2. In Section 3, we analyze the complexities our proposed algorithms GEM-UOT. The results on approximability of OT via UOT are established in Section 4. In Section 5, we experiment on both synthetic and real datasets to compare our algorithm with the state-of-the-art Sinkhorn algorithm in terms of both the performance and the induced sparsity, empirically verify our theories on approximation error, and showcase our framework on the task of color transfer. We conclude the paper in Section 6.

2. Background

In this Section, we first present the entropic regularized UOT problem, used by Sinkhorn algorithm. Then we consider the squared ℓ_2 -norm regularized UOT and derive a new dual formulation, which is the basis for our algorithmic design.

2.1 Entropic Regularized UOT

Inspired by the literature of the entropic regularized OT problem, the entropic version of the UOT problem has been considered. The problem is formulated as:

$$\min_{X \in \mathbb{R}_+^{n \times n}} \langle C, X \rangle - \eta H(X) + \tau \mathbf{KL}(X \mathbf{1}_n \| \mathbf{a}) + \tau \mathbf{KL}(X^\top \mathbf{1}_n \| \mathbf{b}), \quad (3)$$

where $\eta > 0$ is a given regularization parameter. By (Chizat et al., 2018), optimizing the Fenchel-Legendre dual of the above entropic regularized UOT is equivalent to:

$$\min_{\mathbf{u}, \mathbf{v} \in \mathbb{R}^n} \eta \sum_{i,j=1}^n \exp\left(\frac{u_i + v_j - C_{ij}}{\eta}\right) + \tau \langle e^{-\mathbf{u}/\tau}, \mathbf{a} \rangle + \tau \langle e^{-\mathbf{v}/\tau}, \mathbf{b} \rangle. \quad (4)$$

Remark 2 We note that existing algorithm (Pham et al., 2020) must set small $\eta = O(\frac{\varepsilon}{\log(n)})$ to drive the regularizing term $-\eta H(X)$ small, while the smoothness condition number of the above dual objective is large at least at the exponential order of η^{-1} . On the other hand, the original UOT objective is non-smooth due to KL divergences. Therefore, direct application of gradient methods to solve either the primal or the dual of entropic regularized UOT would not result in competitive convergence rate, which perhaps is the reason for the limited literature of gradient methods in the context of UOT. Recently, (Essid and Solomon, 2018; Blondel et al., 2018) also observe that entropic regularization imposed by the Sinkhorn algorithm keeps the transportation plan dense and strictly positive.

2.2 Squared ℓ_2 -norm Regularized UOT and Sparsity

Remark 2 motivates our choice of squared ℓ_2 -norm regularization, which would later result in better condition number in the dual formulation with only linear dependence on η^{-1} . Of practical interest, it also leads to sparse transportation plan empirically as later demonstrated in our experimental Section 5.3.

Remark 3 While squared ℓ_2 -norm was well-studied as a sparsity-induced regularization in standard OT (Blondel et al., 2018; Peyré and Cuturi, 2019), for which dedicated solvers have been developed (Lorenz et al., 2019; Essid and Solomon, 2018), to the best of our knowledge, no work has considered squared ℓ_2 -norm regularization in the context of UOT. This paper is the first to propose the squared ℓ_2 -norm UOT problem, consequently derive feasible algorithms, and empirically investigate the sparsity of the output couplings.

The squared ℓ_2 -norm UOT problem is formulated as:

$$\mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) = \min_{X \in \mathbb{R}_+^{n \times n}} \{g_\eta(X) := \langle C, X \rangle + \eta \|X\|_2^2 + \tau \mathbf{KL}(X \mathbf{1}_n \| \mathbf{a}) + \tau \mathbf{KL}(X^\top \mathbf{1}_n \| \mathbf{b})\}, \quad (5)$$

where $\eta > 0$ is a given regularization parameter. Let $X^\eta = \operatorname{argmin}_{X \in \mathbb{R}_+^{n \times n}} g_\eta(X)$, which is unique by the strong convexity of $g_\eta(X)$. We proceed to establish the duality of the squared ℓ_2 -norm UOT problem, which is based on the Fenchel-Rockafellar Theorem and would be the basis for our algorithmic development.

Lemma 4 *The dual problem to (5) can be written as:*

$$\begin{aligned} \max_{\mathbf{x}=(\mathbf{u},\mathbf{v}) \in \mathbb{R}^{2n}} \left\{ F_a(\mathbf{x}) := -\frac{1}{4\eta} \sum_{i,j=1}^n \max\{0, u_i + v_j - C_{ij}\}^2 \right. \\ \left. - \tau \left\langle e^{-\mathbf{u}/\tau}, \mathbf{a} \right\rangle - \tau \left\langle e^{-\mathbf{v}/\tau}, \mathbf{b} \right\rangle + \tau \mathbf{a}^\top \mathbf{1}_n + \tau \mathbf{b}^\top \mathbf{1}_n \right\}. \end{aligned} \quad (6)$$

Let $(\mathbf{u}^*, \mathbf{v}^*)$ be an optimal solution to (6), then the primal solution to (5) is given by:

$$X_{ij}^\eta = \frac{1}{2\eta} \max\{0, u_i^* + v_j^* - C_{ij}\}, \quad \forall i, j \in [n]. \quad (7)$$

Moreover, we have $\forall i, j \in [n]$:

$$-\frac{u_i^*}{\tau} + \log(a_i) = \log\left(\sum_{k=1}^n X_{ik}^\eta\right), \quad (8)$$

$$-\frac{v_j^*}{\tau} + \log(b_j) = \log\left(\sum_{k=1}^n X_{kj}^\eta\right). \quad (9)$$

Proof Given a convex set S , we consider the convex indicator function:

$$\mathbb{1}_S(X) = \begin{cases} 0 & , \text{ if } X \in S \\ +\infty, & \text{ otherwise} \end{cases},$$

and write the primal objective (5) as:

$$\mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) = \min_{X \in \mathbb{R}^{n \times n}} \{ \langle C, X \rangle + \eta \|X\|_2^2 + \mathbb{1}_S(X) + \tau \mathbf{KL}(X \mathbf{1}_n \| \mathbf{a}) + \tau \mathbf{KL}(X^\top \mathbf{1}_n \| \mathbf{b}) \} \quad (10)$$

with $S = \{X \in \mathbb{R}^{n \times n} | X_{ij} \geq 0, \forall i, j \in [n]\}$. Next, we consider the functions:

$$G(X) = \frac{1}{\eta} \langle C, X \rangle + \|X\|_2^2 + \frac{1}{\eta} \mathcal{X}_S(X), \quad F_1(\mathbf{y}) = \tau \mathbf{KL}(\mathbf{y} \| \mathbf{a}), \quad F_2(\mathbf{y}) = \tau \mathbf{KL}(\mathbf{y} \| \mathbf{b}),$$

whose convex conjugates are as follows:

$$G^*(p) := \sup_{X \in \mathbb{R}^{n \times n}} \{ \langle p, X \rangle - G(X) \} = \frac{1}{4} \sum_{i,j=1}^n \max\{0, p_{ij} - \frac{1}{\eta} C_{ij}\}^2,$$

$$F_1^*(\mathbf{u}) := \sup_{\mathbf{y} \in \mathbb{R}^n} \{ \langle \mathbf{u}, \mathbf{y} \rangle - F_1(\mathbf{y}) \} = \tau \left\langle e^{\mathbf{u}/\tau}, \mathbf{a} \right\rangle - \tau \mathbf{a}^\top \mathbf{1}_n,$$

$$F_2^*(\mathbf{v}) := \sup_{\mathbf{y} \in \mathbb{R}^n} \{ \langle \mathbf{v}, \mathbf{y} \rangle - F_2(\mathbf{y}) \} = \tau \left\langle e^{\mathbf{v}/\tau}, \mathbf{b} \right\rangle - \tau \mathbf{b}^\top \mathbf{1}_n.$$

Consider the linear operator $A : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{n \times n}$ that maps $A(\mathbf{u}, \mathbf{v}) = X$ with $X_{ij} = u_i + v_j$, then A is continuous and its adjoint $A^* : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{2n}$ is $A^*(X) = (X\mathbf{1}_n, X^T\mathbf{1}_n)$. Now note that the problem (6) can be rewritten as:

$$\max_{\mathbf{u}, \mathbf{v} \in \mathbb{R}^n} -F_1^*(-\mathbf{u}) - F_2^*(-\mathbf{v}) - \eta G^*\left(\frac{A(\mathbf{u}, \mathbf{v})}{\eta}\right).$$

By the Fenchel-Rockafellar Theorem (restated as Theorem 29 in Appendix B), we obtain its dual problem as:

$$\inf_{X \in \mathbb{R}^{n \times n}} F_1(X\mathbf{1}_n) + F_2(X^T\mathbf{1}_n) + \eta G(X),$$

which is the optimization problem (10) and obtain that:

$$\begin{aligned} X^\eta \in \partial G^*\left(\frac{A(\mathbf{u}^*, \mathbf{v}^*)}{\eta}\right) &\implies \forall i, j \in [n] : X_{ij}^\eta = \frac{1}{2\eta} \max\{0, u_i^* + v_j^* - C_{ij}\}, \\ -\mathbf{u}^* \in \partial F_1(X^\eta \mathbf{1}_n) &\implies \forall i : -\frac{u_i^*}{\tau} + \log(a_i) = \log\left(\sum_{k=1}^n X_{ik}^\eta\right), \\ -\mathbf{v}^* \in \partial F_2((X^\eta)^T \mathbf{1}_n) &\implies \forall j : -\frac{v_j^*}{\tau} + \log(b_j) = \log\left(\sum_{k=1}^n X_{kj}^\eta\right), \end{aligned}$$

which conclude the proof of the Lemma. ■

An equivalent dual form of (5) can be deduced from the above Lemma (4) and presented in the next Corollary.

Corollary 5 *The dual problem to (5) can be written as:*

$$\max_{(\mathbf{u}, \mathbf{v}, \mathbf{t}) \in \mathcal{X}} \left\{ -\frac{1}{4\eta} \sum_{i,j=1}^n t_{ij}^2 - \tau \langle e^{-\mathbf{u}/\tau}, \mathbf{a} \rangle - \tau \langle e^{-\mathbf{v}/\tau}, \mathbf{b} \rangle + \tau \mathbf{a}^\top \mathbf{1}_n + \tau \mathbf{b}^\top \mathbf{1}_n \right\}, \quad (11)$$

where $\mathcal{X} = \{(\mathbf{u}, \mathbf{v}, \mathbf{t}) \mid \mathbf{u}, \mathbf{v} \in \mathbb{R}^n, \mathbf{t} \in \mathbb{R}^{n \times n} : t_{ij} \geq 0, t_{ij} \geq u_i + v_j - C_{ij} \ \forall i, j\}$. Let $(\mathbf{u}^*, \mathbf{v}^*, \mathbf{t}^*)$ be an optimal solution to (11), then the primal solution to (5) is given by:

$$X_{ij}^\eta = \frac{1}{2\eta} t_{ij}^*. \quad (12)$$

Moreover, we have $\forall i, j \in [n]$:

$$-\frac{u_i^*}{\tau} + \log(a_i) = \log\left(\sum_{k=1}^n X_{ik}^\eta\right), \quad (13)$$

$$-\frac{v_j^*}{\tau} + \log(b_j) = \log\left(\sum_{k=1}^n X_{kj}^\eta\right), \quad (14)$$

$$t_{ij}^* = \max\{0, u_i^* + v_j^* - C_{ij}\}. \quad (15)$$

Proof From the dual objective (6) in Lemma 4, we replace $\max\{0, u_i + v_j - C_{ij}\}$ with the dummy variable t_{ij} constrained by $t_{ij} \geq 0$ and $t_{ij} \geq u_i + v_j - C_{ij}$. We have just rewritten our problem (5) as the optimization problem (11) over the convex set \mathcal{X} . We thus have (15) at the optimal solution, while (12), (13) and (14) directly follow from (7), (8) and (9). ■

Remark 6 We note that if $(\mathbf{u}^*, \mathbf{v}^*, \mathbf{t}^*)$ is an optimal solution to (11), then $(\mathbf{u}^*, \mathbf{v}^*)$ is an optimal solution to (6).

3. Complexity Analysis of Approximating UOT

In this Section, we provide the algorithmic development based on Gradient Extrapolation Method (GEM) for approximating UOT. We first define the quantities within interest and present the regularity conditions in Section 3.1. Next, we prove several helpful properties of squared ℓ_2 -norm UOT in Section 3.2. We then characterize the problem as composite optimization in Section 3.3, on which GEM-UOT is derived. The complexity analysis for GEM-UOT follows in Section 3.4. In Section 3.5, we consider a relaxed UOT problem that is relevant to certain practical applications, and present a easy-to-implement and practical algorithm, termed GEM-RUOT, for such setting.

3.1 List of Quantities and Assumptions

Given the two masses $\mathbf{a}, \mathbf{b} \in \mathbb{R}_+^n$, we define the following notations and quantities:

$$\begin{aligned} a_{min} &= \min_{1 \leq i \leq n} \{a_i\}, & b_{min} &= \min_{1 \leq j \leq n} \{b_j\}, & a_{max} &= \|\mathbf{a}\|_\infty, & b_{max} &= \|\mathbf{b}\|_\infty \\ \kappa &= \frac{1}{\min\{a_{min}, b_{min}\}}, & R &= \frac{(\alpha + \beta)^2}{4}, & p &= \frac{1}{2} \min\{a_{min}, b_{min}\} e^{-\frac{D}{\tau}}, & q &= \alpha + \beta, \\ D &= \|C\|_\infty + \eta(\alpha + \beta) + \tau \log\left(\frac{\alpha + \beta}{2}\right) - \tau \min\{\log(a_{min}), \log(b_{min})\}, \\ L_1 &= \|C\|_\infty + 2\eta q + 2\tau |\log(p)| + 2\tau |\log(q)| + \tau \max_i |\log(a_i)| + \tau \max_i |\log(b_i)|. \end{aligned}$$

We hereby present the assumptions (A1-A3) required by our algorithm. For interpretation, we also restate the regularity conditions of Sinkhorn (Pham et al., 2020) for solving UOT, where detailed discussion on how their assumptions in the original paper are equivalent to (S1-S4) is deferred to Appendix A.1.

Regularity Conditions of this Paper	Regularity Conditions of Sinkhorn
(A1) $a_{min} > 0, b_{min} > 0$.	(S1) $a_{min} > 0, b_{min} > 0$.
(A2) $ \log(a_{min}) = O(\log(n)),$ $ \log(b_{min}) = O(\log(n)).$	(S2) $ \log(a_{min}) = O(\log(n)),$ $ \log(b_{min}) = O(\log(n)).$
(A3) $\tau = \Omega(\min\{\frac{1}{\alpha + \beta}, \ C\ _\infty\})$.	(S3) $ \log(a_{max}) = O(\log(n)),$ $ \log(b_{max}) = O(\log(n)).$
	(S4) α, β, τ are positive constants.

Remark 7 Compared to the regularity conditions of Sinkhorn, ours lift the strict assumptions (S3) and (S4) that put an upper bound on τ and the input masses. Thus our complexity

analysis supports much more flexibility for the input masses and is still suitable to applications requiring large τ to enforce the marginal constraints. On the other hand, our method requires very mild condition (A3) on τ . For example, (A3) is naturally satisfied by (S4) combined with the boundedness of the grounded cost matrix C which is widely assumed by the literature (Altschuler et al., 2017; Dvurechensky et al., 2018). We strongly note that our method works for any $\alpha, \beta, \|C\|_\infty$ as long as τ , a parameter under our control, satisfies (A3).

3.2 Properties of Squared ℓ_2 -norm UOT

We next present several useful properties and related bounds for the squared ℓ_2 -norm UOT problem.

Lemma 8 *The following identities hold:*

$$g_\eta(X^\eta) + 2\tau\|X^\eta\|_1 + \eta\|X^\eta\|_2^2 = \tau(\alpha + \beta), \quad (16)$$

$$f(X_f) + 2\tau\|X_f\|_1 = \tau(\alpha + \beta). \quad (17)$$

Proof The identity (17) follows from (Pham et al., 2020, Lemma 4). To prove (16), we consider the function $g_\eta(tX^\eta)$, where $t \in \mathbb{R}^+$, we have

$$g_\eta(tX^\eta) = \langle C, tX^\eta \rangle + \eta\|tX^\eta\|_2^2 + \tau\mathbf{KL}(tX^\eta\mathbf{1}_n\|\mathbf{a}) + \tau\mathbf{KL}((tX^\eta)^\top\mathbf{1}_n\|\mathbf{b}).$$

Simple algebraic manipulation gives:

$$\begin{aligned} \mathbf{KL}(tX^\eta\mathbf{1}_n\|\mathbf{a}) &= t\mathbf{KL}(X^\eta\mathbf{1}_n\|\mathbf{a}) + (1-t)\alpha + \|X^\eta\|_1 t \log(t) \\ \mathbf{KL}((tX^\eta)^\top\mathbf{1}_n\|\mathbf{b}) &= t\mathbf{KL}((X^\eta)^\top\mathbf{1}_n\|\mathbf{b}) + (1-t)\beta + \|X^\eta\|_1 t \log(t). \end{aligned}$$

We thus obtain that:

$$g_\eta(tX^\eta) = tg_\eta(X^\eta) + \tau(1-t)(\alpha + \beta) + 2\tau\|X^\eta\|_1 t \log(t) + \eta(t^2 - t)\|X^\eta\|_2^2.$$

Differentiating $g_\eta(tX^\eta)$ with respect to t :

$$\frac{\partial g_\eta(tX^\eta)}{\partial t} = g_\eta(X^\eta) - \tau(\alpha + \beta) + 2\tau\|X^\eta\|_1(1 + \log(t)) + \eta(2t - 1)\|X^\eta\|_2^2.$$

From the above analysis, we can see that $g_\eta(tX^\eta)$ is well-defined for all $t \in \mathbb{R}^+$ and attains its minimum at $t = 1$. Setting $\frac{\partial g_\eta(tX^\eta)}{\partial t}|_{t=1} = 0$, we obtain the identity (16). \blacksquare

The next Lemma provides bounds for both the primal and dual solutions to the squared ℓ_2 -norm UOT problem.

Lemma 9 *We have the following bounds for the optimal solution $(\mathbf{u}^*, \mathbf{v}^*, \mathbf{t}^*)$ of (11):*

$$\|X^\eta\|_1 \leq \frac{\alpha + \beta}{2}, \quad (18)$$

$$\|X_f\|_1 \leq \frac{\alpha + \beta}{2}, \quad (19)$$

$$u_i^* \geq \tau \log \left(\frac{2a_i}{\alpha + \beta} \right) \forall i \in [n], \quad (20)$$

$$v_j^* \geq \tau \log \left(\frac{2b_j}{\alpha + \beta} \right) \forall j \in [n], \quad (21)$$

$$\|\mathbf{u}^*\|_\infty, \|\mathbf{v}^*\|_\infty \leq D, \quad (22)$$

$$\min_{i,j} \left\{ \sum_{k=1}^n X_{ik}^\eta, \sum_{k=1}^n X_{kj}^\eta \right\} \geq \min\{a_{min}, b_{min}\} e^{-\frac{D}{\tau}}. \quad (23)$$

Proof Noting that $f(X)$ and $g_\eta(X)$ are non-negative for $X \in \mathbb{R}_+^{n \times n}$, we directly obtain (18) and (19) from Lemma 8. By $X_{ij}^\eta \geq 0, \forall i, j = 1, \dots, n$, and (18), we have:

$$\frac{\alpha + \beta}{2} \geq \|X^\eta\|_1 = \sum_{i,j=1}^n X_{ij}^\eta \geq \max \left\{ \sum_{k=1}^n X_{ik}^\eta, \sum_{k=1}^n X_{kj}^\eta \right\}. \quad (24)$$

We obtain in view of Lemma 4 that:

$$\log(a_i) - \frac{u_i^*}{\tau} = \log \left(\sum_{j=1}^n X_{ij}^\eta \right) \leq \log \left(\frac{\alpha + \beta}{2} \right), \quad i = 1, \dots, n, \quad (25)$$

$$\log(b_j) - \frac{v_j^*}{\tau} = \log \left(\sum_{i=1}^n X_{ij}^\eta \right) \leq \log \left(\frac{\alpha + \beta}{2} \right), \quad j = 1, \dots, n, \quad (26)$$

which are equivalent to:

$$u_i^* \geq \tau \log \left(\frac{2a_i}{\alpha + \beta} \right) \geq \tau \left(\log(a_{min}) - \log \left(\frac{\alpha + \beta}{2} \right) \right), \quad i = 1, \dots, n, \quad (27)$$

$$v_j^* \geq \tau \log \left(\frac{2b_j}{\alpha + \beta} \right) \geq \tau \left(\log(b_{min}) - \log \left(\frac{\alpha + \beta}{2} \right) \right), \quad j = 1, \dots, n. \quad (28)$$

By (12), we have $\frac{1}{2\eta}(u_i^* + v_j^* - C_{ij}) \leq X_{ij}^\eta, \forall i, j = 1, \dots, n$. Hence,

$$u_i^* \leq 2\eta X_{ij}^\eta + C_{ij} - v_j^* \stackrel{(28)}{\leq} 2\eta X_{ij}^\eta + C_{ij} - \tau \left(\log(b_{min}) - \log \left(\frac{\alpha + \beta}{2} \right) \right), \quad (29)$$

$$v_j^* \leq 2\eta X_{ij}^\eta + C_{ij} - u_i^* \stackrel{(27)}{\leq} 2\eta X_{ij}^\eta + C_{ij} - \tau \left(\log(a_{min}) - \log \left(\frac{\alpha + \beta}{2} \right) \right). \quad (30)$$

Note that $X_{ij}^\eta \leq \max_{i,j} |X_{ij}^\eta| = \|X^\eta\|_\infty \leq \|X^\eta\|_1 \leq \frac{\alpha+\beta}{2}$. We have

$$\begin{aligned} \tau \left(\log(a_{\min}) - \log\left(\frac{\alpha+\beta}{2}\right) \right) &\leq u_i^* \leq \eta(\alpha+\beta) + C_{ij} - \tau \left(\log(b_{\min}) - \log\left(\frac{\alpha+\beta}{2}\right) \right), \\ \tau \left(\log(b_{\min}) - \log\left(\frac{\alpha+\beta}{2}\right) \right) &\leq v_j^* \leq \eta(\alpha+\beta) + C_{ij} - \tau \left(\log(a_{\min}) - \log\left(\frac{\alpha+\beta}{2}\right) \right), \\ \therefore \max\{\|\mathbf{u}^*\|_\infty, \|\mathbf{v}^*\|_\infty\} &\leq D = \|C\|_\infty + \eta(\alpha+\beta) + \tau \log\left(\frac{\alpha+\beta}{2}\right) \\ &\quad - \tau \min\{\log(a_{\min}), \log(b_{\min})\}. \end{aligned}$$

To prove (23), we note that by Lemma 4, $\forall i : \sum_{k=1}^n X_{ik}^\eta = a_i \cdot e^{-\frac{u_i^*}{\tau}} \geq a_i \cdot e^{-\frac{\|\mathbf{u}^*\|_\infty}{\tau}} \geq a_i \cdot e^{-\frac{D}{\tau}}$. Similarly, we have $\forall j : \sum_{k=1}^n X_{kj}^\eta \geq b_j \cdot e^{-\frac{D}{\tau}}$ and thus obtain that:

$$\min_{i,j} \left\{ \sum_{k=1}^n X_{ik}^\eta, \sum_{k=1}^n X_{kj}^\eta \right\} \geq \min\{a_{\min}, b_{\min}\} e^{-\frac{D}{\tau}},$$

which concludes the proof of this Lemma. \blacksquare

The local Lipschitzness of the squared ℓ_2 -norm UOT objective is established in the following Lemma.

Lemma 10 *For $0 < p \leq q$ defined in Section 3.1, let $U_{p,q} = \{X \in \mathbb{R}_+^{n \times n} \mid \forall i, j : \sum_{k=1}^n X_{ik} \in [p, q], \sum_{k=1}^n X_{kj} \in [p, q]\}$. Then $g_\eta(X)$ is L_g -Lipschitz in $U_{p,q}$, i.e. $\forall X, Y \in U_{p,q}$:*

$$|g_\eta(X) - g_\eta(Y)| \leq L_1 \|X - Y\|_1,$$

where $L_1 = \|C\|_\infty + 2\eta q + 2\tau |\log(p)| + 2\tau |\log(q)| + \tau \max_i |\log(a_i)| + \tau \max_i |\log(b_i)|$.

Proof Take any $X, Y \in U_{p,q}$. By Mean Value Theorem, there exists $c \in (0, 1)$ and $Z = (1-c)X + cY \in U_{p,q}$ such that:

$$|g_\eta(X) - g_\eta(Y)| = |\langle \nabla g_\eta(Z), X - Y \rangle| \leq \|\nabla g_\eta(Z)\|_\infty \|X - Y\|_1, \quad (31)$$

where we have used Holder inequality. We also have:

$$\frac{\partial g_\eta}{\partial Z_{ij}} = C_{ij} + 2\eta Z_{ij} + \tau \log\left(\sum_{k=1}^n Z_{ik}\right) + \tau \log\left(\sum_{k=1}^n Z_{kj}\right) - \tau \log(a_i) - \tau \log(b_j).$$

Since $Z \in U_{p,q}$, we obtain $\forall i, j : \sum_{k=1}^n Z_{ik}, \sum_{k=1}^n Z_{kj} \in [p, q]$, which implies that:

$$\left| \log\left(\sum_{k=1}^n Z_{ik}\right) \right|, \left| \log\left(\sum_{k=1}^n Z_{kj}\right) \right| \leq |\log(p)| + |\log(q)|$$

From the above observations, we consequently obtain that:

$$\begin{aligned} \|\nabla g_\eta(Z)\|_\infty &\leq \|C\|_\infty + 2\eta q + 2\tau |\log(p)| + 2\tau |\log(q)| \\ &\quad + \tau \max_i |\log(a_i)| + \tau \max_i |\log(b_i)| = L_g \end{aligned} \quad (32)$$

Plugging (32) into (31) yields the desired inequality. \blacksquare

3.3 Characterization of the Dual Objective

GEM-UOT is designed by solving the dual (11) from Corollary 5, which is equivalent to:

$$\min_{\mathbf{x} \in \mathcal{X}} h_\eta(\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t})) := \frac{1}{4\eta} \sum_{i,j=1}^n t_{ij}^2 + \tau \langle e^{-\mathbf{u}/\tau}, \mathbf{a} \rangle + \tau \langle e^{-\mathbf{v}/\tau}, \mathbf{b} \rangle. \quad (33)$$

For $\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t})$, we consider the functions:

$$f_\eta(\mathbf{x}) = \tau \langle e^{-\mathbf{u}/\tau}, \mathbf{a} \rangle + \tau \langle e^{-\mathbf{v}/\tau}, \mathbf{b} \rangle - \frac{\min\{a_{min}, b_{min}\}}{2\tau} e^{-D/\tau} (\|\mathbf{u}\|_2^2 + \|\mathbf{v}\|_2^2), \quad (34)$$

$$w_\eta(\mathbf{x}) = \frac{\min\{a_{min}, b_{min}\}}{2\tau} e^{-D/\tau} (\|\mathbf{u}\|_2^2 + \|\mathbf{v}\|_2^2) + \frac{1}{4\eta} \sum_{i,j=1}^n t_{ij}^2. \quad (35)$$

Then the problem (33) can be rewritten as:

$$\min_{\mathbf{x} \in \mathcal{X}} h_\eta(\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t})) := f_\eta(\mathbf{x}) + w_\eta(\mathbf{x}), \quad (36)$$

which can be characterized as the composite optimization over the sum of the locally convex smooth $f_\eta(\mathbf{x})$ and the locally strongly convex $w_\eta(\mathbf{x})$ by the following Lemma 11, whose proof is deferred to Appendix C.3.

Lemma 11 *Let $V_D = \{\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t}) \in \mathbb{R}^{n^2+2n} : \forall i, j \in [n], \tau \log(\frac{2a_i}{\alpha+\beta}) \leq u_i \leq D, \tau \log(\frac{2b_j}{\alpha+\beta}) \leq v_j \leq D\}$. Then $f_\eta(\mathbf{x})$ is convex and L -smooth in the domain V_D , and $w_\eta(\mathbf{x})$ is μ -strongly convex with:*

$$L = \frac{\alpha + \beta}{2\tau} + \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau},$$

$$\mu = \min \left\{ \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau}, \frac{1}{2\eta} \right\}.$$

The choice of D in the above characterization is motivated by the next Corollary 12, which ensures that the optimal solution \mathbf{x}^* lies within the locality V_D .

Corollary 12 *If \mathbf{x}^* is an optimal solution to (36), then $\mathbf{x}^* \in V_D$.*

Proof The result follows directly from Lemma 9. ■

3.4 Gradient Extrapolation Method for UOT (GEM-UOT)

Algorithm Description: We develop our algorithm based on Gradient Extrapolation Method (GEM), called GEM-UOT, to solve the UOT problem. GEM was proposed by (Lan and Zhou, 2018) to address problem in the form of (36) and viewed as the dual of the Nesterov's accelerated gradient method. We hereby denote the prox-function associated with w_η by:

$$P(\mathbf{x}_0, \mathbf{x}) := \frac{w_\eta(\mathbf{x}) - [w_\eta(\mathbf{x}_0) + \langle \nabla w_\eta(\mathbf{x}_0), \mathbf{x} - \mathbf{x}_0 \rangle]}{\mu},$$

and the proximal mapping associated with the closed convex set \mathcal{Y} and w is defined as:

$$\mathcal{M}_{\mathcal{Y}}(\mathbf{g}, \mathbf{x}_0, \theta) := \operatorname{argmin}_{\mathbf{x}=(\mathbf{u}, \mathbf{v}, \mathbf{t}) \in \mathcal{Y}} \{ \langle \mathbf{g}, \mathbf{x} \rangle + w_{\eta}(\mathbf{x}) + \theta P(\mathbf{x}_0, \mathbf{x}) \}. \quad (37)$$

One challenge for optimizing (36) is the fact that the component function $f_{\eta}(\mathbf{x})$ is smooth and convex only on the locality V_D as in Lemma 11. By Corollary 12, we can further write (36) as:

$$\min_{\mathbf{x} \in \mathcal{X}} h_{\eta}(\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t})) = \min_{\mathbf{x} \in V_D \cap \mathcal{X}} h_{\eta}(\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t})). \quad (38)$$

While minimizing the UOT dual objective, GEM-UOT hence projects its iterates onto the closed convex set $V_D \cap \mathcal{X}$ through the proximal mapping $\mathcal{M}_{V_D \cap \mathcal{X}}(\cdot)$ to enforce the smoothness and convexity of $f_{\eta}(\mathbf{x})$ on the convex hull of all the iterates, while ensuring they are in the feasible domain. Finally, GEM-UOT outputs the approximate solution, computed from the iterates, to the original UOT problem. The steps are summarized in Algorithm 1.

Algorithm 1 GEM-UOT

```

1: Input:  $C, \mathbf{a}, \mathbf{b}, \varepsilon, \tau, \eta$ 
2: Initialization:  $\underline{\mathbf{x}}^0 = \mathbf{x}^0 = \mathbf{0}, \mathbf{y}^{-1} = \mathbf{y}^0 = \mathbf{0}$ 
3: Compute  $L, \mu$  based Lemma 11
4:  $\zeta = 1 - 1/(1 + \sqrt{1 + 16L/\mu})$ 
5:  $\psi = \frac{1}{1-\zeta} - 1, \rho = \frac{\zeta}{1-\zeta}\mu$ 
6: for  $t = 0, 1, 2, \dots, k$  do
7:    $\tilde{\mathbf{y}}^t = \mathbf{y}^{t-1} + \zeta(\mathbf{y}^{t-1} - \mathbf{y}^{t-2})$ 
8:    $\mathbf{x}^t = \mathcal{M}_{V_D \cap \mathcal{X}}(\tilde{\mathbf{y}}^t, \mathbf{x}^{t-1}, \rho)$ 
9:    $\underline{\mathbf{x}}^t = (1 + \psi)^{-1}(\mathbf{x}^t + \psi \underline{\mathbf{x}}^{t-1})$ 
10:   $\mathbf{y}^t = \nabla f_{\eta}(\underline{\mathbf{x}}^t)$ 
11:   $\theta_t = \theta_{t-1} \zeta^{-1}$ 
12: end for
13: Compute  $\underline{\mathbf{x}}^k = (\sum_{t=1}^k \theta_t)^{-1} \sum_{t=1}^k \theta_t \mathbf{x}^t$ 
14: Let  $\underline{\mathbf{u}}^k, \underline{\mathbf{v}}^k, \underline{\mathbf{t}}^k$  be the vectors such that
     $\underline{\mathbf{x}}^k = (\underline{\mathbf{u}}^k, \underline{\mathbf{v}}^k, \underline{\mathbf{t}}^k)$ 
15: for  $i, j = 1 \rightarrow n$  do
16:    $X_{ij}^k = \frac{1}{2\eta} t_{ij}^k$ 
17: end for
18: Return  $X^k$ 
    
```

Algorithm 2 GEM-RUOT

```

1: Input:  $C, \mathbf{a}, \mathbf{b}, \varepsilon, \tau, \eta$ 
2: Initialization:  $\underline{\mathbf{x}}^0 = \mathbf{x}^0 = \mathbf{0}, \mathbf{y}^{-1} = \mathbf{y}^0 = \mathbf{0}$ 
3: Compute  $L_a$  based on Lemma 17
4: for  $t = 0, 1, 2, \dots, k$  do
5:    $\zeta_t = \frac{t-1}{t}, \psi_t = \frac{t-1}{2}, \rho_t = \frac{6L_a}{t}$ 
6:    $\tilde{\mathbf{y}}^t = \mathbf{y}^{t-1} + \zeta_t(\mathbf{y}^{t-1} - \mathbf{y}^{t-2})$ 
7:    $\mathbf{x}^t = \bar{\mathcal{M}}_{V_a}(\tilde{\mathbf{y}}^t, \mathbf{x}^{t-1}, \rho_t)$ 
8:    $\underline{\mathbf{x}}^t = (1 + \psi_t)^{-1}(\mathbf{x}^t + \psi_t \underline{\mathbf{x}}^{t-1})$ 
9:    $\mathbf{y}^t = \nabla h_a(\underline{\mathbf{x}}^t)$ 
10:   $\theta_t = \theta_{t-1} \zeta_t^{-1}$ 
11: end for
12:  $\underline{\mathbf{x}}^k = (\sum_{t=1}^k \theta_t)^{-1} \sum_{t=1}^k \theta_t \mathbf{x}^t$ 
13: Let  $\underline{\mathbf{u}}^k, \underline{\mathbf{v}}^k$  be the vectors such that
     $\underline{\mathbf{x}}^k = (\underline{\mathbf{u}}^k, \underline{\mathbf{v}}^k)$ 
14: for  $i, j = 1 \rightarrow n$  do
15:    $X_{ij}^k = \frac{1}{2\eta} \max\{0, \underline{u}_i^k + \underline{v}_j^k - C_{ij}\}$ 
16: end for
17: Compute  $F_a(\underline{\mathbf{x}}^k)$ 
18: Return  $F_a(\underline{\mathbf{x}}^k), X^k$ 
    
```

Complexity Analysis: We proceed to establish the complexity for approximating UOT. The following quantities are used in our analysis:

$$\sigma_0^2 = \|\nabla f_{\eta}(\mathbf{x}^0)\|_2^2, \quad \Delta_{0, \sigma_0} = \mu P(\mathbf{x}^0, \mathbf{x}^*) + h_{\eta}(\mathbf{x}^0) - h_{\eta}(\mathbf{x}^*) + \frac{\sigma_0^2}{\mu}.$$

In the next Theorem, we further quantify the number of iterations

$$K_0 = 4 \left(1 + \sqrt{1 + 16L/\mu} \right) \log \left(\frac{4n^2 \Delta_{0, \sigma_0}^{1/2}}{\eta \mu} \max \left\{ \frac{L_1}{\varepsilon}, \frac{e^{D/\tau}}{\min\{a_{\min}, b_{\min}\}}, \frac{1}{\alpha + \beta} \right\} \right)$$

that is required for GEM-UOT to output the ε -approximate solution to $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$.

Theorem 13 *For $\eta = \frac{\varepsilon}{2R}$ and $k \geq K_0$, the output X^k of Algorithm 1 is the ε -approximate solution to the UOT problem, i.e.*

$$f(X^k) - f(X_f) = f(X^k) - \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) \leq \varepsilon. \quad (39)$$

Proof Recalling that $X^\eta = \operatorname{argmin}_{X \in \mathbb{R}_+^{n \times n}} g_\eta(X)$, we get:

$$\begin{aligned} f(X^k) - f(X_f) &= g_\eta(X^k) - \eta \|X^k\|_2^2 - g_\eta(X_f) + \eta \|X_f\|_2^2 \\ &\leq (g_\eta(X^k) - g_\eta(X^\eta)) + \eta (\|X_f\|_2^2 - \|X^k\|_2^2). \end{aligned} \quad (40)$$

From Lemma 9, we have $\|X_f\|_2^2 \leq \|X_f\|_1^2 \leq \frac{(\alpha+\beta)^2}{4} = R$. Since $\eta = \frac{\varepsilon}{2R}$, we further obtain that:

$$\eta (\|X_f\|_2^2 - \|X^k\|_2^2) \leq \frac{\varepsilon}{2}. \quad (41)$$

Plugging (41) into (40), we have:

$$f(X^k) - f(X_f) \leq g_\eta(X^k) - g_\eta(X^\eta) + \frac{\varepsilon}{2}. \quad (42)$$

Thus, it is left to bound $g_\eta(X^k) - g_\eta(X^\eta)$. Next, we proceed establish the bound on $\|X^k - X^\eta\|_1$ in the following Lemma 14.

Lemma 14 *If $k \geq K_0$, the output X^k of Algorithm 1 satisfies:*

$$\|X^k - X^\eta\|_1 \leq \min \left\{ \frac{\varepsilon}{2L_1}, \frac{\min\{a_{\min}, b_{\min}\} e^{-D/\tau}}{2}, \frac{\alpha + \beta}{2} \right\}, \quad (43)$$

The proof of Lemma 14 can be found in Appendix C.1 where we first bound the quantity in terms of the distance $\|\underline{\mathbf{x}}^k - \underline{\mathbf{x}}^*\|_2$ between the dual iterate and dual solution, which provably vanishes to the desired error $\frac{\varepsilon}{2L_1}$ for large enough k thanks to the guarantee from GEM being applied to the dual objective (37). Lemma 14 also helps in deriving the next Lemma 15, whose proof is in Appendix C.2.

Lemma 15 *For $k \geq K_0$, we have $X^k, X^\eta \in U_{p,q}$.*

Lemma 15 is necessary to invoke Lemma 10 on the Lipschitzness of $g_\eta(X)$ within the convex domain $U_{p,q}$. In particular, we obtain from Lemma 10 with $X^k, X^\eta \in U_{p,q}$ that:

$$g_\eta(X^k) - g_\eta(X^\eta) \leq L_1 \|X^k - X^\eta\|_1 \quad (44)$$

From (43) and (44), we get:

$$g_\eta(X^k) - g_\eta(X^\eta) \leq \frac{\varepsilon}{2}. \quad (45)$$

By plugging (45) into (42), we conclude the desired inequality (39). \blacksquare

Finally, the complexity of GEM-UOT is given by the following Corollary 16.

Corollary 16 *Under the Assumptions (A1-A3), Algorithm 1 has the iteration complexity of $O\left((\alpha + \beta)\kappa \log\left(\frac{\tau \cdot n(\alpha + \beta)}{\varepsilon}\right)\right)$ and $\tilde{O}(n^2)$ cost per iteration³.*

3. $\tilde{O}()$ excludes the logarithmic terms in the complexity.

Proof Note that Algorithm 1 sets $\eta = \frac{\varepsilon}{2R} = \frac{2\varepsilon}{(\alpha+\beta)^2}$. Under the Assumptions (A1-A3), we have:

$$\begin{aligned} \frac{D}{\tau} &= \frac{\|C\|_\infty}{\tau} + \frac{\varepsilon}{\tau(\alpha+\beta)} + \log\left(\frac{\alpha+\beta}{2}\right) - \min\{\log(a_{\min}), \log(b_{\min})\} \\ &= O(1) + \log\left(\frac{\alpha+\beta}{2}\right) - \min\{\log(a_{\min}), \log(b_{\min})\} \end{aligned}$$

By Lemma 11 and under the Assumptions (A1-A2), we obtain that:

$$\begin{aligned} L &= O\left(\frac{\alpha+\beta}{2\tau}\right) \\ \mu &= \Omega\left(\min\left\{\frac{(\alpha+\beta)^2}{4\varepsilon}, \frac{1}{\tau} \min\{a_{\min}, b_{\min}\} e^{-\log\left(\frac{\alpha+\beta}{2}\right) + \min\{\log(a_{\min}), \log(b_{\min})\}}\right\}\right) \\ &= \Omega\left(\min\left\{\frac{(\alpha+\beta)^2}{4\varepsilon}, \frac{1}{\tau} \min\{a_{\min}, b_{\min}\}^2\right\}\right) \\ &\geq \Omega\left(\min\left\{\frac{n^2(2\min\{a_{\min}, b_{\min}\})^2}{4\varepsilon}, \frac{2\min\{a_{\min}, b_{\min}\}^2}{\tau(\alpha+\beta)}\right\}\right) \\ &= \Omega\left(\frac{\min\{a_{\min}, b_{\min}\}^2}{\tau(\alpha+\beta)}\right) \text{ (since } \tau = \Omega\left(\frac{1}{\alpha+\beta}\right)\text{)} \\ \sqrt{\frac{L}{\mu}} &= O\left(\frac{(\alpha+\beta)}{\min\{a_{\min}, b_{\min}\}}\right) = O((\alpha+\beta)\kappa). \end{aligned}$$

For convenience, we recall that $P(\mathbf{x}_0, \mathbf{x}) = (w_\eta(\mathbf{x}) - [w_\eta(\mathbf{x}_0) + \langle \nabla w_\eta(\mathbf{x}_0), \mathbf{x} - \mathbf{x}_0 \rangle]) / \mu$. From the initialization $\mathbf{x}^0 = (\mathbf{u}^0, \mathbf{v}^0, \mathbf{t}^0) = \mathbf{0}$, $\eta = \frac{\varepsilon}{2R} = \frac{2\varepsilon}{(\alpha+\beta)^2}$ and Lemma 9, we have:

$$\begin{aligned} \frac{D}{\tau} &= O(1) + \log\left(\frac{\alpha+\beta}{2}\right) - \min\{\log(a_{\min}), \log(b_{\min})\} \leq O(\log(n(\alpha+\beta))) \\ \sigma_0^2 &= \|\nabla f_\eta(\mathbf{x}^0)\|_2^2 = \|\mathbf{a}\|_2^2 + \|\mathbf{b}\|_2^2 \leq (\alpha+\beta)^2 \\ h_\eta(\mathbf{x}^0) &= \tau(\alpha+\beta) \geq h_\eta(\mathbf{x}^*) \geq 0, \quad w_\eta(\mathbf{x}^0) = 0, \quad \nabla w_\eta(\mathbf{x}^0) = \mathbf{0} \\ 0 \leq w_\eta(\mathbf{x}^*) &\leq \frac{\min\{a_{\min}, b_{\min}\}}{2\tau} (n\|\mathbf{u}^*\|_\infty^2 + n\|\mathbf{v}\|_\infty^2) + \frac{1}{4\eta} \|X^\eta\|_2^2 \\ &\leq O(\min\{a_{\min}, b_{\min}\} \tau n \log(n)^2 + (\alpha+\beta)^4 / \varepsilon) \\ \Delta_{0, \sigma_0} &= \mu P(\mathbf{x}^0, \mathbf{x}^*) + h_\eta(\mathbf{x}^0) - h_\eta(\mathbf{x}^*) + \frac{\sigma_0^2}{\mu} \\ &= O\left(\tau \min\{a_{\min}, b_{\min}\} n \log(n)^2 + \tau(\alpha+\beta) + \frac{\tau(\alpha+\beta)^2}{\min\{a_{\min}, b_{\min}\}^2}\right). \end{aligned}$$

We thus obtain the asymptotic complexity under (A1-A3) for the following term inside K_0 :

$$\log\left(\frac{4n^2 \Delta_{0, \sigma_0}^{1/2}}{\eta \mu} \max\left\{\frac{L_1}{\varepsilon}, \frac{e^{D/\tau}}{\min\{a_{\min}, b_{\min}\}}, \frac{1}{\alpha+\beta}\right\}\right) = O\left(\log\left(\frac{\tau \cdot n(\alpha+\beta)}{\varepsilon}\right)\right).$$

Combining with the bounds of L and μ above, we obtain that the iteration complexity is $K_0 = O\left((\alpha + \beta)\kappa \log\left(\frac{\tau \cdot n(\alpha + \beta)}{\varepsilon}\right)\right)$ under the assumptions (A1-A3).

Lines 7 and 9 are $O(1)$ point-wise update steps of vectors with $n^2 + 2n$ entries, thereby incurring $O(n^2)$ complexity. The proximal mapping on line 8 is equivalent to solving a sparse separable quadratic program of size $O(n^2)$ with sparse linear constraints and can be solved efficiently in $\tilde{O}(n^2)$ time and with $O(n^2)$ space (see Appendix C.4 for details). The computation of the gradient on line 10 takes $O(n^2)$ since the partial derivative ∂f_η with respect to any of its $n^2 + 2n$ variables takes $O(1)$ operations. Thus, the cost per iteration is $\tilde{O}(n^2)$ in total. ■

The complexities of GEM-UOT and other algorithms in the literature are summarized in Table 1. GEM-UOT is asymptotically better than Sinkhorn in τ and ε , and, as mentioned before, has the edge on practicality over Sinkhorn.

Table 1: Complexity of algorithms for solving UOT problems.

Algorithms	Assumptions	Cost per iteration	Iteration complexity
MM (Chapel et al., 2021)	Unspecified iteration complexity	$O(n^2)$	$O(n^{1.27})$ - Empirical Complexity
Sinkhorn (Pham et al., 2020)	(S1), (S2), (S3), (S4)	$O(n^2)$	$O\left((\alpha + \beta)\frac{\tau \log(n)}{\varepsilon} \log\left(\frac{\log(n)(\alpha + \beta)}{\varepsilon}\right)\right)$
GEM-UOT (this paper)	(A1)≡(S1), (A2)≡(S2), (A3)	$\tilde{O}(n^2)$	$O\left((\alpha + \beta)\kappa \log\left(\frac{\tau \cdot n(\alpha + \beta)}{\varepsilon}\right)\right)$

Novelty and Practicality: GEM-UOT is predicated on the novel dual formulation of the squared ℓ_2 -norm regularized UOT and the intricate design of smooth locality V_D (Lemma 11). We note that naive application of GEM would not lead to competitive complexity. Beside function decomposition, we design the smooth locality V_D to enforce tight bound for the condition number L . The techniques can benefit follow-up work tackling squared ℓ_2 -norm UOT. Next, we highlight the favorable features of GEM-UOT that Sinkhorn is short of:

- *Compatibility with modern applications:* Gradient methods are specifically congruent with and thus deployed as heuristics without theoretical guarantee in many emerging UOT applications (Balaji et al., 2020; Yang and Uhler, 2019; Frogner et al., 2015). Applications (Pitié et al., 2007; Courty et al., 2017; Muzellec et al., 2016) in which only sparse transportation plan is of interest also nullify the usage of Sinkhorn, since entropic regularization imposed by it keeps the transportation plan dense and strictly positive (Essid and Solomon, 2018; Blondel et al., 2018). GEM-UOT addresses both of these problems by being the first gradient-based sparse UOT solver in the literature.
- *Logarithmic dependence on τ :* In the regime of large τ , Sinkhorn’s linear dependence on τ much hinders its practicality, motivating (Sejourne et al., 2022) to alleviate this

issue in the specific case of 1-D UOT. However, no algorithm for *general* UOT problem that could achieve logarithmic dependence on τ had existed before GEM-UOT. From the practical perspective, τ can be large in certain applications (Schiebinger et al., 2019) and, for a specific case of UOT penalized by squared ℓ_2 norm⁴, be of the order of thousands (Blondel et al., 2018).

- *Logarithmic dependence on ε^{-1}* : GEM-UOT is the first algorithm in the literature that achieves logarithmic dependence on ε^{-1} . As the class of scaling algorithms, especially subsuming Sinkhorn, is known to suffer from numerical instability (De Plaen et al., 2023), GEM-UOT thus emerges as a practical solver for the high-accuracy regime.

3.5 Relaxed UOT Problem and Practical Perspectives

In this Section, we discuss the relaxed UOT problem where only the UOT distance is within interest, and propose a practical algorithm, termed GEM-RUOT, for this setting.

Motivation of Relaxed UOT Problem: Certain generative model applications (Nguyen et al., 2021) recently have adopted either OT or UOT as loss metrics for training, in which only the computation of the distance is required, while the transportation plan itself is not within interest. To this end, we consider the *relaxed* UOT (RUOT) problem, where the goal is to approximate the UOT distance oblivious of the transportation plan.

Practicality: Gradient methods have been shown to achieve competitive complexity in the OT literature (Guminov et al., 2021; Dvurechensky et al., 2018). The heuristic proposed by (Yang and Uhler, 2019) and based on gradient descent achieved favorable performance for solving the UOT problem. Through our complexity analysis of GEM-UOT, we want to establish the preliminary understanding of the class of gradient-based optimization when applied to UOT problem. Nevertheless, the convoluted function decomposition (36) of GEM-UOT may make it hard to implement in practice. The strong convexity constant μ (Lemma 11) can be small on certain real datasets, thereby hindering the empirical performance. We thus develop an easy-to-implement algorithm, called GEM-RUOT, for practical purposes that avoids the dependency on μ in complexity and is specialized for the RUOT problem. One interesting direction is to incorporate Adaptive Gradient Descent (Malitsky and Mishchenko, 2020) that automatically adapts to local geometry and is left for future work.

Alternative Dual Formulation: GEM-RUOT does not require intricate function decomposition and optimizes directly over the following dual formulation (6) in Lemma 4. Optimizing (6) is equivalent to:

$$\min_{\mathbf{u}, \mathbf{v} \in \mathbb{R}^n} h_a(\mathbf{x} = (\mathbf{u}, \mathbf{v})) := \frac{1}{4\eta} \sum_{i,j=1}^n \max\{0, u_i + v_j - C_{ij}\}^2 + \tau \langle e^{-\mathbf{u}/\tau}, \mathbf{a} \rangle + \tau \langle e^{-\mathbf{v}/\tau}, \mathbf{b} \rangle, \quad (46)$$

which is locally smooth by the following Lemma 17, whose proof is given in Appendix C.5.

Lemma 17 *Let $V_a = \{\mathbf{x} = (\mathbf{u}, \mathbf{v}) \in \mathbb{R}^{2n} : \forall i, j \in [n], \tau \log(\frac{2a_i}{\alpha+\beta}) \leq u_i \leq D, \tau \log(\frac{2b_j}{\alpha+\beta}) \leq v_j \leq D\}$. Then $h_a(\mathbf{x})$ is L_a -smooth and convex in V_a with: $L_a = \frac{\alpha+\beta}{\tau} + \frac{2\sqrt{n}}{\eta}$.*

4. We recall that our paper considers UOT penalized by KL divergence, which is the most prevalent UOT variant (see more details in Section 1).

GEM-RUOT adopts the ℓ_2 distance for its prox-function $\bar{P}(\mathbf{x}_0, \mathbf{x}) = \frac{1}{2}\|\mathbf{x} - \mathbf{x}_0\|_2^2$, and consequently uses the standard ℓ_2 projection operator for the proximal mapping associated with the closed convex set \mathcal{Y} as $\mathcal{M}_{\mathcal{Y}}(\mathbf{g}, \mathbf{x}_0, \theta) := \operatorname{argmin}_{\mathbf{x}=(\mathbf{u}, \mathbf{v}) \in \mathcal{Y}} \{\langle \mathbf{g}, \mathbf{x} \rangle + \theta \bar{P}(\mathbf{x}_0, \mathbf{x})\}$. Using the convex version of GEM, GEM-RUOT optimizes directly over (46) and projects its iterates onto the closed convex set V_a through the proximal mapping $\bar{\mathcal{M}}_{V_a}(\cdot)$ to enforce the smoothness on the convex hull of all the iterates. Finally, it returns $F_a(\underline{\mathbf{x}}^k)$ as an ε -approximation of the distance $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ and a heuristic transportation plan X^k . The steps are summarized in Algorithm 2 and GEM-RUOT's complexity is given in Theorem 18.

Theorem 18 *If $k \geq \sqrt{\frac{12L_a n D^2}{\varepsilon}}$, then $|F_a(\underline{\mathbf{x}}^k) - \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})| \leq \varepsilon$. Under the assumptions (A1-A3), the complexity of Algorithm 2 is $O\left((\alpha + \beta) \cdot \frac{\tau n^{2.75}}{\varepsilon} \cdot \log(n(\alpha + \beta))\right)$.*

Proof Let $\mathbf{x}^* = (\mathbf{u}^*, \mathbf{v}^*)$ be an optimal solution to (46) and thus (6). Since $\|\mathbf{x}^*\|_\infty \leq D$ by Lemma 9 and Remark 6, we have $\bar{P}(\mathbf{x}^0, \mathbf{x}^*) = \frac{1}{2}\|\mathbf{x}^*\|_2^2 \leq \frac{1}{2}nD^2$. From corollary 3.6 in (Lan and Zhou, 2018), we obtain:

$$\begin{aligned} 0 \leq h_a(\underline{\mathbf{x}}^k) - h_a(\mathbf{x}^*) &= F_a(\mathbf{x}^*) - F_a(\underline{\mathbf{x}}^k) = \mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) - F_a(\underline{\mathbf{x}}^k) \leq \frac{12L_a}{k(k+1)}\bar{P}(\mathbf{x}^0, \mathbf{x}^*) \\ &\therefore 0 \leq \mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) - F_a(\underline{\mathbf{x}}^k) \leq \frac{6L_a n D^2}{k(k+1)} \end{aligned} \quad (47)$$

Using $\|X_f\|_1 \leq \frac{\alpha+\beta}{2}$ (Lemma 9) and $\eta = \frac{\varepsilon}{2R} = \frac{2\varepsilon}{(\alpha+\beta)^2}$ set by the Algorithm 2, we have:

$$\begin{aligned} \mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) - \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) &= g_\eta(X^\eta) - f(X_f) \\ &= (f(X^\eta) - f(X_f)) + \eta\|X^\eta\|_2^2 \geq 0 \\ \mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) - \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) &= g_\eta(X^\eta) - g_\eta(X_f) + \eta\|X_f\|_2^2 \\ &\leq \eta\|X_f\|_2^2 \leq \eta \frac{(\alpha + \beta)^2}{4} = \frac{\varepsilon}{2} \\ \therefore 0 \leq \mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) - \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) &\leq \frac{\varepsilon}{2} \end{aligned} \quad (48)$$

Combining (47), (48), we obtain:

$$|F_a(\underline{\mathbf{x}}^k) - \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})| \leq \frac{\varepsilon}{2} + \frac{6L_a n D^2}{k(k+1)} \quad (49)$$

The RHS of (49) is less than ε if $k \geq \sqrt{\frac{12L_a n D^2}{\varepsilon}} = K_a$. Under the assumptions (A1-A3) and and $\eta = \frac{\varepsilon}{2R} = \frac{2\varepsilon}{(\alpha+\beta)^2}$, we have:

$$\begin{aligned} D &= O(\tau \log(n(\alpha + \beta))), \quad L_a = O\left(\frac{(\alpha + \beta)^2 \sqrt{n}}{\varepsilon}\right) \\ \therefore K_a &= O\left((\alpha + \beta) \cdot \frac{\tau n^{0.75}}{\varepsilon} \cdot \log(n(\alpha + \beta))\right) \end{aligned}$$

Since it costs $O(n^2)$ per iteration to compute the gradient $h_a(\mathbf{x})$, the total complexity of Algorithm 2 is $O\left((\alpha + \beta) \cdot \frac{\tau n^{2.75}}{\varepsilon} \cdot \log(n(\alpha + \beta))\right)$. \blacksquare

4. Approximability of Standard OT

For balanced masses and $\tau \rightarrow \infty$, UOT problem reduces to standard OT. In this Section, we establish the very first characterization of the vanishing approximation error between UOT and OT in terms of τ . To facilitate our discussion, we formally define the standard OT problem in Section 4.1. We prove several some approximability properties of UOT in Section 4.2, which are helpful for later deriving the main results. Then in Section 4.3, we show that UOT's transportation plan converges to the marginal constraints of standard OT in the ℓ_1 sense with the rate $O(\frac{n}{\tau})$. Beside transportation plan, we upperbound the difference in transport plan between OT and UOT distances by $O(\frac{n^2}{\tau})$ in Section 4.4.

4.1 Standard OT

When the masses lie in the probability simplex, i.e. $\mathbf{a}, \mathbf{b} \in \Delta^n := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_1 = 1\}$, the standard OT problem (Kantorovich, 1942) can be formulated as:

$$\mathbf{OT}(\mathbf{a}, \mathbf{b}) = \min_{X \in \Pi(\mathbf{a}, \mathbf{b})} \langle C, X \rangle \quad (50)$$

where $\Pi(\mathbf{a}, \mathbf{b}) = \{X \in \mathbb{R}^{n \times n} : X \mathbf{1}_n = \mathbf{a}, X^\top \mathbf{1}_n = \mathbf{b}\}$.

Definition 19 For $\varepsilon > 0$, X is an ε -approximation transportation plan of $\mathbf{OT}(\mathbf{a}, \mathbf{b})$ if:

$$\langle C, X \rangle \leq \langle C, X^{OT} \rangle + \varepsilon, \quad (51)$$

where $X^{OT} = \operatorname{argmin}_{X \in \Pi(\mathbf{a}, \mathbf{b})} \langle C, X \rangle$ is the optimal transportation plan of the OT problem.

4.2 Approximability Properties of UOT

For any $\eta > 0$, consider the optimal solution $\mathbf{x}^* = (\mathbf{u}^*, \mathbf{v}^*)$ to (6). Note that since $\mathbf{a}, \mathbf{b} \in \Delta^n$, we have $\alpha = \beta = 1$. By Lemma 9, we obtain:

$$\|X^\eta\|_1 \leq 1. \quad (52)$$

The following Lemmas are useful for the proofs of the main Theorem 23 and 24 later.

Lemma 20 The followings hold $\forall i, j \in [n]$:

$$u_i^* \leq \|C\|_\infty + 2\eta \quad (53)$$

$$v_j^* \leq \|C\|_\infty + 2\eta \quad (54)$$

Proof Firstly, we show that there must exist some $l \in [n]$ such that $v_l \geq 0$. Assume for the sake of contradiction that $\forall j \in [n] : v_j < 0$. Then from (14), we obtain $\forall j \in [n]$:

$$\begin{aligned} \log\left(\sum_{k=1}^n X_{kj}^\eta\right) &> \log(b_j) \\ (X^{\eta\top} \mathbf{1}_n)_j &= \sum_{k=1}^n X_{kj}^\eta > b_j \\ \therefore \|X^\eta\|_1 &= \sum_{j=1}^n (X^{\eta\top} \mathbf{1}_n)_j > \sum_{j=1}^n b_j = 1, \end{aligned}$$

where the last line contradicts (52). From (52), (12) and taking any $v_l \geq 0$, we have $\forall i \in [n]$:

$$1 \geq X_{il}^\eta = \frac{1}{2\eta} \max\{0, u_i^* + v_l^* - C_{il}\} \geq \frac{1}{2\eta}(u_i^* + 0 - \|C\|_\infty) \implies u_i^* \leq \|C\|_\infty + 2\eta$$

Similarly, we can prove that $\forall j \in [n]$: $v_j^* \leq \|C\|_\infty + 2\eta$. ■

Lemma 21 *Define $\gamma = \|C\|_\infty + 2\eta$. Then the followings hold $\forall i, j \in [n]$:*

$$a_i e^{-\gamma/\tau} \leq \sum_{j=1}^n X_{ij}^\eta = (X^\eta \mathbf{1}_n)_i \leq 1 - e^{-\gamma/\tau}(1 - a_i) \quad (55)$$

$$b_j e^{-\gamma/\tau} \leq \sum_{i=1}^n X_{ij}^\eta = (X^{\eta\top} \mathbf{1}_n)_j \leq 1 - e^{-\gamma/\tau}(1 - b_j) \quad (56)$$

Proof From (13), we have $\forall i \in [n]$:

$$(X^\eta \mathbf{1}_n)_i = a_i e^{-u_i^*/\tau} \stackrel{(53)}{\geq} a_i e^{-\gamma/\tau} \quad (57)$$

For the upper bound, we have $\forall i \in [n]$:

$$(X^\eta \mathbf{1}_n)_i = \|X^\eta\|_1 - \sum_{k \neq i} (X^\eta \mathbf{1}_n)_k \stackrel{(52)+(57)}{\leq} 1 - \sum_{k \neq i} a_k e^{-\gamma/\tau} = 1 - e^{-\gamma/\tau}(1 - a_i),$$

which has proved (55). Similarly, (56) is obtained. ■

Lemma 22 *Define $\gamma = \|C\|_\infty + 2\eta$. Then the followings hold:*

$$\|X^\eta \mathbf{1}_n - \mathbf{a}\|_1 \leq \frac{n\gamma}{\tau} \quad (58)$$

$$\|X^{\eta\top} \mathbf{1}_n - \mathbf{b}\|_1 \leq \frac{n\gamma}{\tau} \quad (59)$$

Proof From (55), we have:

$$\begin{aligned} -a_i(1 - e^{-\gamma/\tau}) &\leq (X^\eta \mathbf{1}_n)_i - a_i \leq (1 - a_i)(1 - e^{-\gamma/\tau}) \\ \therefore |(X^\eta \mathbf{1}_n)_i - a_i| &\leq \max\{a_i, 1 - a_i\}(1 - e^{-\gamma/\tau}) \leq \gamma/\tau, \end{aligned}$$

where the last line follows from $\max\{a_i, 1 - a_i\} \leq 1$ and $e^x \geq x + 1, \forall x \in \mathbb{R}$. We obtain that:

$$\|X^\eta \mathbf{1}_n - \mathbf{a}\|_1 = \sum_{i=1}^n |(X^\eta \mathbf{1}_n)_i - a_i| \leq \frac{n\gamma}{\tau},$$

which has proved (58). Similarly, (59) is obtained. ■

4.3 Approximation of OT Transportation Plan

We prove in the next Theorem the rate at which UOT's transportation plan converges to the marginal constraints of standard OT.

Theorem 23 *For $\mathbf{a}, \mathbf{b} \in \Delta^n$, the problem $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ admits some optimal transportation plan $X_f = \operatorname{argmin}_{X \in \mathbb{R}_+^{n \times n}} f(X)$ such that:*

$$\|X_f \mathbf{1}_n - \mathbf{a}\|_1 + \|X_f^\top \mathbf{1}_n - \mathbf{b}\|_1 \leq \frac{2n\|C\|_\infty}{\tau}. \quad (60)$$

Proof From Lemma 9, we know that X^η is bounded by $\|X^\eta\|_1 \leq \frac{\alpha+\beta}{2}$. By the Bolzano–Weierstrass Theorem, there exists a convergent subsequence $\{X^{\eta_t}\}_{t=1}^\infty$ where $\lim_{t \rightarrow \infty} \eta_t = 0$. Define the limit of this subsequence as $X_f = \lim_{\eta_t \rightarrow 0} X^{\eta_t}$. We first prove that X_f is a transportation plan of the $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ problem by showing that X_f is a minimizer of f . By the boundedness of $\|X^{\eta_t}\|_1 \leq \frac{\alpha+\beta}{2}$ (Lemma 9), we have $\lim_{\eta_t \rightarrow 0} \eta_t \|X^{\eta_t}\|_2^2 = 0$ and obtain:

$$\lim_{\eta_t \rightarrow 0} g_{\eta_t}(X^{\eta_t}) = \lim_{\eta_t \rightarrow 0} [f(X^{\eta_t}) + \eta_t \|X^{\eta_t}\|_2^2] = \lim_{\eta_t \rightarrow 0} f(X^{\eta_t}) = f(\lim_{\eta_t \rightarrow 0} X^{\eta_t}) = f(X_f) \quad (61)$$

where the last equality is by the continuity of f . Now take any $\tilde{X} \in \mathbb{R}_+^{n \times n}$. we have:

$$\begin{aligned} g_{\eta_t}(X^{\eta_t}) &\leq g_{\eta_t}(\tilde{X}) = f(\tilde{X}) + \eta_t \|\tilde{X}\|_2^2 \\ \therefore \lim_{\eta_t \rightarrow 0} g_{\eta_t}(X^{\eta_t}) &\leq \lim_{\eta_t \rightarrow 0} g_{\eta_t}(\tilde{X}) = f(\tilde{X}). \end{aligned} \quad (62)$$

Combining (61) and (62), we have $f(X_f) \leq f(\tilde{X})$, thereby proving the minimality of X_f . Thus, X_f is a transportation plan of the $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ problem. Now by Lemma 22, we have:

$$\|X^{\eta_t} \mathbf{1}_n - \mathbf{a}\|_1 + \|X^{\eta_t \top} \mathbf{1}_n - \mathbf{b}\|_1 \leq \frac{2n(\|C\|_\infty + 2\eta_t)}{\tau}$$

Taking $\eta_t \rightarrow 0$ and noting that $X_f = \lim_{\eta_t \rightarrow 0} X^{\eta_t}$, we obtain the desired inequality (60). ■

The quantity $\|X \mathbf{1}_n - \mathbf{a}\|_1 + \|X^\top \mathbf{1}_n - \mathbf{b}\|_1$ measures the closeness of X to $\Pi(\mathbf{a}, \mathbf{b})$ in ℓ_1 distance and is also used as the stopping criteria in algorithms solving standard OT (Altschuler et al., 2018; Dvurechensky et al., 2018). Therefore, Theorem 23 characterizes the rate $O(\frac{n}{\tau})$ by which the transportation plan of $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ converges to the marginal constraints $\Pi(\mathbf{a}, \mathbf{b})$ of $\mathbf{OT}(\mathbf{a}, \mathbf{b})$. Utilizing this observation, we present the algorithm GEM-OT (Algorithm 3), which solves the $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ problem via GEM-UOT with fine tuned τ and finally projects the UOT solution onto $\Pi(\mathbf{a}, \mathbf{b})$ (via Algorithm 4), to find an ε -approximate solution to the standard OT problem $\mathbf{OT}(\mathbf{a}, \mathbf{b})$. The following Theorem characterizes the complexity of GEM-OT.

Algorithm 3 GEM-OT

- 1: **Input:** $C, \mathbf{a}, \mathbf{b}, \varepsilon$
 - 2: $\varepsilon' = \varepsilon/16$
 - 3: $\eta = \frac{2\varepsilon'}{(\alpha+\beta)^2} = \varepsilon'/2$
 - 4: $\gamma = \|C\|_\infty + \eta$
 - 5: $\tau = \frac{16\|C\|_\infty n \gamma}{\varepsilon}$
 - 6: $\bar{X} = \text{GEM-UOT}(C, \mathbf{a}, \mathbf{b}, \varepsilon', \tau, \eta)$
 - 7: $Y = \text{PROJ}(\bar{X}, \mathbf{a}, \mathbf{b})$
 - 8: **Return** Y
-

Algorithm 4 PROJ (Altschuler et al., 2017, Algorithm 2)

- 1: **Input:** $X \in \mathbb{R}^{n \times n}, \mathbf{a} \in \mathbb{R}^n, \mathbf{b} \in \mathbb{R}^n$
 - 2: $P \leftarrow \text{diag}(\mathbf{x})$ with $x_i = \frac{a_i}{(X\mathbf{1}_n)_i} \wedge 1$
 - 3: $X' \leftarrow PX$
 - 4: $Q \leftarrow \text{diag}(\mathbf{y})$ with $y_j = \frac{b_j}{(X'^\top \mathbf{1}_n)_j} \wedge 1$
 - 5: $X'' \leftarrow X'Q$
 - 6: $\text{err}_r \leftarrow \mathbf{a} - X''\mathbf{1}_n, \text{err}_c \leftarrow \mathbf{b} - X''^\top \mathbf{1}_n$
 - 7: **Return** $Y \leftarrow X'' + \text{err}_r \text{err}_c^\top / \|\text{err}_r\|_1$
-

Theorem 24 *Algorithm 3 outputs $Y \in \Pi(\mathbf{a}, \mathbf{b})$ such that:*

$$\langle C, Y \rangle \leq \langle C, X^{OT} \rangle + \varepsilon.$$

In other words, Y is an ε -approximate transportation plan of $\text{OT}(\mathbf{a}, \mathbf{b})$. Under assumptions (A1-A2), the complexity of Algorithm 3 is $\tilde{O}(\kappa n^2)$.

Proof We obtain via simple algebra and Holder inequality that:

$$\begin{aligned} \langle C, Y \rangle - \langle C, X^{OT} \rangle &= \langle C, \bar{X} - X^\eta \rangle + \langle C, Y - \bar{X} \rangle + \langle C, X^\eta - X^{OT} \rangle \\ &\leq \|C\|_\infty \|\bar{X} - X^\eta\|_1 + \|C\|_\infty \|Y - \bar{X}\|_1 + \langle C, X^\eta - X^{OT} \rangle \end{aligned} \quad (63)$$

We proceed to upper-bound each of the three terms in the RHS of (63). We note that \bar{X} is the output of GEM-UOT fed with input error $\varepsilon/16$, and GEM-UOT (in the context of Algorithm 3) thus sets $\eta = \frac{2\varepsilon/16}{(\alpha+\beta)^2} = \frac{\varepsilon}{32}$. The first term can be bounded by Lemma 14 as:

$$\|C\|_\infty \|\bar{X} - X^\eta\|_1 \leq \|C\|_\infty \frac{\varepsilon/16}{2L_1} \leq \frac{\varepsilon}{8} \quad (64)$$

where $L_1 \geq \|C\|_\infty$ holds directly from the definition of L_1 . For the second term, we obtain from Lemma 30 that:

$$\begin{aligned} \|C\|_\infty \|Y - \bar{X}\|_1 &\leq 2\|C\|_\infty [\|\bar{X}\mathbf{1}_n - \mathbf{a}\|_1 + \|\bar{X}^\top \mathbf{1}_n - \mathbf{b}\|_1] \\ &\leq 2\|C\|_\infty [\|X^\eta \mathbf{1}_n - \mathbf{a}\|_1 + \|X^{\eta \top} \mathbf{1}_n - \mathbf{b}\|_1] \\ &\quad + 2\|C\|_\infty [\|(X^\eta - \bar{X})\mathbf{1}_n\|_1 + \|(X^{\eta \top} - \bar{X}^\top)\mathbf{1}_n\|_1] \\ &\leq \frac{4\|C\|_\infty n \gamma}{\tau} + 4\|C\|_\infty \|\bar{X} - X^\eta\|_1 \end{aligned} \quad (65)$$

$$\leq \frac{\varepsilon}{4} + \frac{4\varepsilon}{8} = \frac{3\varepsilon}{4} \quad (66)$$

where (65) follows Lemma 22 with $\gamma = \|C\|_\infty + 2\eta$ and triangular inequality, and (66) is by the definition of τ and (64). Since $X^{OT} \in \Pi(\mathbf{a}, \mathbf{b})$, we have: $\mathbf{KL}(X^{OT} \mathbf{1}_n \|\mathbf{a}) = \mathbf{KL}(X^{OT \top} \mathbf{1}_n \|\mathbf{b}) = 0$ and $\eta \|X^{OT}\|_2^2 \leq \eta \|X^{OT}\|_1^2 = \eta$. We thus obtain that:

$$g_\eta(X^{OT}) = \langle C, X^{OT} \rangle + \eta \|X^{OT}\|_2^2 \leq \langle C, X^{OT} \rangle + \eta \quad (67)$$

Since the **KL** divergence and ℓ_2 -norm are non-negative, we have:

$$g_\eta(X^\eta) \geq \langle C, X^\eta \rangle. \quad (68)$$

We recall that $X^\eta = \operatorname{argmin}_{X \in \mathbb{R}_+^{n \times n}} g_\eta(X)$, which implies:

$$g_\eta(X^{OT}) \geq g_\eta(X^\eta). \quad (69)$$

Combining (67), (68) and (69), we upper-bound the third term in the RHS of (63):

$$\langle C, X^\eta - X^{OT} \rangle \leq \eta = \frac{\varepsilon}{32}. \quad (70)$$

Plugging (64), (66) and (70) into (63), we have:

$$\langle C, Y \rangle - \langle C, X^{OT} \rangle \leq \frac{\varepsilon}{8} + \frac{3\varepsilon}{4} + \frac{\varepsilon}{32} < \varepsilon.$$

The complexity of GEM-OT is the total of $O(n^2)$ for the PROJ() algorithm and the complexity of GEM-UOT for the specified τ in the algorithm. By Corollary 16, which establishes the complexity of GEM-UOT under the assumptions (A1-A3), we conclude that the complexity of GEM-OT under (A1-A2) (where (A3) is naturally satisfied by our choice of τ) is $\tilde{O}(\kappa n^2)$. \blacksquare

Remark 25 *The best-known complexities in the literature of OT are respectively $\tilde{O}(n^{2.5})$ (Lee and Sidford, 2014) for finding exact solution, and $\tilde{O}(\frac{n^2}{\varepsilon})$ (Luo et al., 2023; Xie et al., 2023) for finding the ε -approximate solution. We highlight that the complexity in Theorem 24 is the first to achieve logarithmic dependence on ε^{-1} to obtain the ε -approximate solution, while still enjoying better dependence on n than the exact solver (Lee and Sidford, 2014). As the prevalent entropy-regularized methods to approximate OT (Luo et al., 2023; Altschuler et al., 2017; Dvurechensky et al., 2018) have been known to suffer from numerical instability and blurring issues (Zhou and Parno, 2023), GEM-OT can emerge as a practical solver for OT specialized in the high-accuracy regime (Dong et al., 2020).*

Theorem 23 would elucidate the tuning of τ , which has been a heuristic process in all prior work, such that the transportation plan of $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ well respects the original structure of that of $\mathbf{OT}(\mathbf{a}, \mathbf{b})$. Furthermore, we provide a novel recipe for OT retrieval from UOT, which combines Theorem 23 with a low-cost post-process projection step and can be integrated with any UOT solver beside GEM-UOT as considered in this paper.

4.4 Approximation Error between UOT and OT Distances

As discussed in Section 1.1, while a lot of work have used UOT with small $\tau = 1$ as a relaxed variant of OT (Fattras et al., 2021; Balaji et al., 2020; Le et al., 2021), we could empirically illustrate noticeably large approximation error between the UOT and OT distances on real dataset. Such deviation can be detrimental to target applications including generative modelling (Fattras et al., 2021) which recently has adopted UOT in place of OT as a loss metrics for training. This further necessitates the understanding of approximation error between UOT and OT in the sense of purely transport distance. We next establish an upper bound on the approximation error between UOT and OT distances in Theorem 26.

Theorem 26 For $\mathbf{a}, \mathbf{b} \in \Delta^n$, we define $M = \log(2)\|C\|_\infty^2 (n + 3\kappa)^2 + 2n\|C\|_\infty^2$ and have the following bound:

$$0 \leq \mathbf{OT}(\mathbf{a}, \mathbf{b}) - \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) \leq \frac{M}{\tau}. \quad (71)$$

Proof The lower bound is straightforward. Since X^{OT} (the transportation plan of $\mathbf{OT}(\mathbf{a}, \mathbf{b})$) is a feasible solution to the optimization problem $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ and $\mathbf{KL}(X^{OT}\mathbf{1}_n|\mathbf{a}) = \mathbf{KL}(X^{OT\top}\mathbf{1}_n|\mathbf{b}) = 0$ (as $X^{OT} \in \Pi(\mathbf{a}, \mathbf{b})$), we obtain that:

$$\mathbf{OT}(\mathbf{a}, \mathbf{b}) = \langle C, X^{OT} \rangle = f(X^{OT}) \geq f(X_f) = \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$$

We proceed to prove the upper bound of the error approximation. We first define a variant of $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ problem where the $\mathbf{KL}()$ divergence is replaced by squared ℓ_2 norm:

$$\mathbf{UOT}_{\ell_2}(\mathbf{a}, \mathbf{b}) = \min_{X \in \mathbb{R}_+^{n \times n}} \left\{ f_{\ell_2}(X) := \langle C, X \rangle + \frac{\tau}{2 \log(2)} \|X\mathbf{1}_n - \mathbf{a}\|_2^2 + \frac{\tau}{2 \log(2)} \|X^\top \mathbf{1}_n - \mathbf{b}\|_2^2 \right\}, \quad (72)$$

By (Blondel et al., 2018, Theorem 2), for $M_1 = \log(2)\|C\|_\infty^2 (n + 3\kappa)^2$, we have:

$$0 \leq \mathbf{OT}(\mathbf{a}, \mathbf{b}) - \mathbf{UOT}_{\ell_2}(\mathbf{a}, \mathbf{b}) \leq \frac{M_1}{\tau}. \quad (73)$$

Remark 27 We highlight that deriving the approximation error for $\mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ is harder than $\mathbf{UOT}_{\ell_2}(\mathbf{a}, \mathbf{b})$. First, KL divergence is non-smooth, which is stated in (Blondel et al., 2018) as the main reason for their choice of squared ℓ_2 norm instead of KL to relax the marginal constraints. Second, the proof in (Blondel et al., 2018) is based on interpreting the dual of $\mathbf{UOT}_{\ell_2}(\mathbf{a}, \mathbf{b})$ as the sum of the dual of $\mathbf{OT}(\mathbf{a}, \mathbf{b})$ and the a squared ℓ_2 -norm regularization on the dual variables, so the approximation error is derived by simply bounding the dual variables.

Recall that:

$$\mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) = \min_{X \in \mathbb{R}_+^{n \times n}} \{g_\eta(X) := \langle C, X \rangle + \eta \|X\|_2^2 + \tau \mathbf{KL}(X\mathbf{1}_n|\mathbf{a}) + \tau \mathbf{KL}(X^\top \mathbf{1}_n|\mathbf{b})\},$$

and $X^\eta = \operatorname{argmin}_{X \in \mathbb{R}_+^{n \times n}} g_\eta(X)$ is the solution to $\mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b})$. We further define the problem $\overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b})$ as follows:

$$\overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) = \min_{X \in \mathbb{R}_+^{n \times n}, \|X\|_1=1} \{g_\eta(X) := \langle C, X \rangle + \eta \|X\|_2^2 + \tau \mathbf{KL}(X\mathbf{1}_n|\mathbf{a}) + \tau \mathbf{KL}(X^\top \mathbf{1}_n|\mathbf{b})\},$$

which optimizes over the same objective cost as $\mathbf{UOT}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b})$ yet with the additional constraint that X lies in the probability simplex, i.e. $\|X\|_1 = 1$. We let the solution to $\overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b})$ be $Z^\eta = \operatorname{argmin}_{X \in \mathbb{R}_+^{n \times n}, \|X\|_1=1} g_\eta(X)$. Since $\|Z^\eta\|_1 = 1$, we have $Z^\eta \mathbf{1}_n, Z^{\eta\top} \mathbf{1}_n \in \Delta^n$. Then Pinsker inequality gives:

$$\begin{aligned} \mathbf{KL}(Z^\eta \mathbf{1}_n|\mathbf{a}) &\geq \frac{1}{2 \log(2)} \|Z^\eta \mathbf{1}_n - \mathbf{a}\|_1^2 \geq \frac{1}{2 \log(2)} \|Z^\eta \mathbf{1}_n - \mathbf{a}\|_2^2 \\ \mathbf{KL}(Z^{\eta\top} \mathbf{1}_n|\mathbf{b}) &\geq \frac{1}{2 \log(2)} \|Z^{\eta\top} \mathbf{1}_n - \mathbf{b}\|_1^2 \geq \frac{1}{2 \log(2)} \|Z^{\eta\top} \mathbf{1}_n - \mathbf{b}\|_2^2 \end{aligned}$$

Consequently, we obtain that:

$$\overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) = g_\eta(Z^\eta) \geq f_{\ell_2}(Z^\eta) \geq \mathbf{UOT}_{\ell_2}(\mathbf{a}, \mathbf{b}) \quad (74)$$

For any $\eta > 0$, we consider $Y^\eta = \text{PROJ}(X^\eta, \mathbf{a}, \mathbf{b})$ as the projection of X^η onto $\Pi(\mathbf{a}, \mathbf{b})$ via Algorithm 4. Then $\|Y^\eta\|_1 = \mathbf{a}^\top \mathbf{1}_n = 1$. By Lemma 30 and 22, we have:

$$\|Y^\eta - X^\eta\|_1 \leq 2[\|X^\eta \mathbf{1}_n - \mathbf{a}\|_1 + \|X^{\eta\top} \mathbf{1}_n - \mathbf{b}\|_1] \leq \frac{2n(\|C\|_\infty + 2\eta)}{\tau}. \quad (75)$$

Since Y^η is a feasible solution to the optimization problem $\overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b})$, we obtain that:

$$g_\eta(Y^\eta) \geq \overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}), \quad (76)$$

Note that $\mathbf{KL}(Y^\eta \mathbf{1}_n \| \mathbf{a}) = \mathbf{KL}(Y^{\eta\top} \mathbf{1}_n \| \mathbf{b}) = 0$ as $Y^\eta \in \Pi(\mathbf{a}, \mathbf{b})$. We thus can write $g_\eta(Y^\eta) = \langle C, Y^\eta \rangle + \eta \|Y^\eta\|_2^2$ and obtain via Holder inequality that:

$$\begin{aligned} g_\eta(Y^\eta) - \overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) &= g_\eta(Y^\eta) - g_\eta(X^\eta) \\ &= \langle C, Y^\eta - X^\eta \rangle + \eta(\|Y^\eta\|_2^2 - \|X^\eta\|_2^2) - \tau \mathbf{KL}(X^\eta \mathbf{1}_n \| \mathbf{a}) - \tau \mathbf{KL}(X^{\eta\top} \mathbf{1}_n \| \mathbf{b}) \\ &\leq \|C\|_\infty \|Y^\eta - X^\eta\|_1 + \eta \|Y^\eta\|_2^2 \stackrel{(75)}{\leq} \frac{2n\|C\|_\infty(\|C\|_\infty + 2\eta)}{\tau} + \eta \end{aligned}$$

where for the last inequality we note that $\|Y^\eta\|_2 \leq \|Y^\eta\|_1 = 1$. The above is equivalent to:

$$\overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) \geq g_\eta(Y^\eta) - \frac{2n\|C\|_\infty(\|C\|_\infty + 2\eta)}{\tau} - \eta. \quad (77)$$

Now combining (74), (76) and (77), we have $\forall \eta > 0$:

$$\overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) \geq \mathbf{UOT}_{\ell_2}(\mathbf{a}, \mathbf{b}) - \frac{2n\|C\|_\infty(\|C\|_\infty + 2\eta)}{\tau} - \eta$$

Taking the limit $\eta \rightarrow 0$ on both sides and noting that $\lim_{\eta \rightarrow 0} \overline{\mathbf{UOT}}_{\mathbf{KL}}^\eta(\mathbf{a}, \mathbf{b}) = \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$, we obtain:

$$\begin{aligned} \overline{\mathbf{UOT}}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) &\geq \mathbf{UOT}_{\ell_2}(\mathbf{a}, \mathbf{b}) - \frac{2n\|C\|_\infty^2}{\tau} \\ \mathbf{OT}(\mathbf{a}, \mathbf{b}) - \overline{\mathbf{UOT}}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) &\leq \mathbf{OT}(\mathbf{a}, \mathbf{b}) - \mathbf{UOT}_{\ell_2}(\mathbf{a}, \mathbf{b}) + \frac{2n\|C\|_\infty^2}{\tau} \end{aligned}$$

Combining with (73), we get the desired bound $\mathbf{OT}(\mathbf{a}, \mathbf{b}) - \overline{\mathbf{UOT}}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b}) \leq \frac{M}{\tau}$ where $M = M_1 + 2n\|C\|_\infty^2$. \blacksquare

5. Experiments

In this Section, we empirically verify GEM-UOT's performance and sparsity, and our theories on approximation error. To compute the ground truth value, we use the convex programming package `cvxpy` to find the exact UOT plan (Agrawal et al., 2018). Unless specified otherwise, we always set $\eta = \frac{\varepsilon}{2R} = \frac{2\varepsilon}{(\alpha+\beta)^2}$ (according to Theorem 13).

5.1 Synthetic Data

We set $n = 200, \tau = 55, \alpha = 4, \beta = 5$. Then a_i 's are drawn from uniform distribution and rescaled to ensure $\alpha = 4$, while b_i 's are drawn from normal distribution with $\sigma = 0.1$. Entries of C are drawn uniformly from $[10^{-1}, 1]$. For both GEM-UOT and Sinkhorn, we vary $\varepsilon = 1 \rightarrow 10^{-4}$ to evaluate their time complexities in Figure 1 and test the τ dependency in Figure 2. Both experiments verify GEM-UOT's better theoretical dependence on τ and ε .




Figure 1: Comparison in number of iterations (left) and per-iteration cost (right) between GEM-UOT and Sinkhorn on synthetic data for $\varepsilon = 1 \rightarrow 10^{-4}$.




Figure 2: Scalability of τ (for $\varepsilon = 10^{-2}$) on synthetic data and CIFAR-10. Sinkhorn scales linearly while GEM-UOT scales logarithmically in τ .

Next, to empirically evaluate the efficiency of GEM-OT for approximating standard OT problem, we consider the setting of balanced masses and large τ with $n = 200, \alpha = \beta = 1$ and $\tau = 500$. We compare GEM-OT with Sinkhorn (Cuturi, 2013) in wall-clock time for varying $\varepsilon = 1 \rightarrow 10^{-4}$ in Figure 3.




Figure 3: Wall-clock time comparison between GEM-OT and Sinkhorn on synthetic dataset.

5.2 Real Datasets: CIFAR-10 and Fashion-MNIST

To practically validate our algorithms, we compare the GEM-UOT with Sinkhorn on CIFAR-10 dataset (Krizhevsky and Hinton, 2009). A pair of flattened images corresponds to the marginals \mathbf{a}, \mathbf{b} , whereby the cost matrix C is the matrix of ℓ_1 distances between pixel locations. This is also the setting considered in (Pham et al., 2020; Dvurechensky et al., 2018). We plot the results in Figure 4, which demonstrates GEM-UOT's superior performance. Additional experiments on Fashion-MNIST dataset that further cover GEM-RUOT are presented in Appendix D. We highlight that GEM-UOT only has logarithmic dependence on τ , improving over Sinkhorn's linear dependence on τ which had been its major bottleneck in the regime of large τ (Sejourne et al., 2022). Experiment on the scalability of τ in Figure 2,

which is averaged over 10 randomly chosen image pairs, empirically reasserts our favorable dependence on τ on the CIFAR-10 dataset. In addition, we evaluate the behaviour of GEM-UOT for different levels of ε and thus η in Figure 5.




Figure 4: Primal gap $f - f^*$ of GEM-UOT and Sinkhorn on CIFAR-10.




Figure 5: Algorithmic dependence on η and the corresponding accuracy achieved by the algorithm on CIFAR-10 for 10^6 iterations.

5.3 Sparsity of GEM-UOT

On two real datasets Fashion-MNIST and CIFAR 10, we empirically illustrate the sparseness of the transportation plans produced by GEM-UOT in Figure 6 and Figure 7, while Sinkhorn produces transportation plan with 0% sparsity in those experiments. This result is consistent with (Blondel et al., 2018), where entropic regularization enforces strictly positive and dense transportation plan and squared- ℓ_2 induces sparse transportation plan. To illustrate the effectiveness of ℓ_2 as the sparse regularization, in Figure 8 we compare the heat maps of the optimal transportation plans of squared- ℓ_2 UOT and original UOT on CIFAR 10. We set the sparsity threshold of 10^{-2} and report the sparsity results in the following figures.

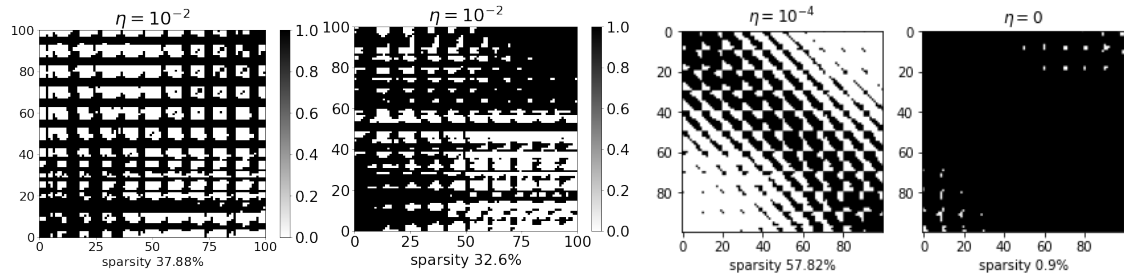


Figure 6: 37.88% sparsity on Fashion-MNIST.

Figure 7: 32.6% sparsity on CIFAR-10.

Figure 8: Sparsity of squared- ℓ_2 UOT plan (left) versus original UOT plan (right) on CIFAR 10.

5.4 Color Transfer

Color transfer is a prominent application where sparse transportation plan is of special interest (Blondel et al., 2018). This experiment shows that squared- ℓ_2 UOT (i.e. sparse UOT) solved by GEM-UOT results in a sparse transportation plan and is thus congruent with this application, while entropic UOT solved by Sinkhorn results in strictly positive

and dense transportation plan, hindering its adoption. In particular, we conduct the UOT between the histograms of two images. Given the source image X and target image Y both of size $l \times h$, we present their pixels in RGB space. Quantizing the images down to n colors, we obtain the color centroids for the source and target images respectively as $S_{source} = \{x_1, x_2, \dots, x_n\}$ and $S_{target} = \{y_1, y_2, \dots, y_n\}$, where $x_i, y_i \in \mathbb{R}^3$. Then the source image's color histogram $\mathbf{a} \in \Delta^n$ is computed as the empirical distribution of pixels over the n centroids, i.e. $a_k = \frac{1}{lh} \cdot (\sum_{i=1}^l \sum_{j=1}^h \mathbb{1}\{X_{ij} \in x_k\})$. The target image's color histogram $\mathbf{b} \in \Delta^n$ is computed similarly. We compute the transportation plan between the two images via sparse UOT solved by GEM-UOT and entropic UOT solved by Sinkhorn, and report their qualities of color transfer in Figure 9, and the sparsity in Figure 10 and Figure 11. To demonstrate how entropic regularized UOT tend to generate dense mappings, we set the sparsity threshold of 10^{-2} for Figure 11 while the sparsity threshold for Figure 10 is 0. (Note that due to the entropic regularization, Sinkhorn has 0% sparsity under such threshold.)




Figure 9: Quality of color transfer by GEM-UOT (solving squared- ℓ_2 UOT) and Sinkhorn (solving entropic UOT).

Figure 10: GEM-UOT gets 99.4% sparsity.

Figure 11: Sinkhorn gets 51% sparsity.

5.5 Approximation Error between UOT and OT

5.5.1 APPROXIMATION OF OT TRANSPORTATION PLAN AND OT RETRIEVAL GAP




Figure 12: Approximability of OT via UOT (Theorem 23) on synthetic data. The gap approaches 0 with rate $O(\frac{n}{\tau})$, as τ grows.




Figure 13: OT retrieval of GEM-OT (Theorem 24) on CIFAR-10. This empirically verifies vanishing OT retrieval gap for $\tau \rightarrow \infty$.




Figure 14: The error $\mathbf{OT}(\mathbf{a}, \mathbf{b}) - \mathbf{UOT}_{\text{KL}}(\mathbf{a}, \mathbf{b})$ (Theorem 26) in terms of the empirical gap (red) and the theoretical gap (black dotted) on the CIFAR-10.

We investigate the approximability of OT via UOT in terms of transportation plan (Theorem 23) on synthetic data where $\alpha = \beta = 1$ and entries of C are drawn uniformly from $[0.5, 1]$. The result is depicted in Figure 12. Next, we verify the error incurred by GEM-OT (Algorithm 3) for retrieving OT solution from UOT. The choice of τ in GEM-OT equivalently means that the retrieval gap ε established in Theorem 24 can be upperbounded

by $\frac{16\|C\|_\infty n^\gamma}{\tau}$ where $\gamma = \|C\|_\infty + \varepsilon/32$. To assert such vanishing effect of the error as τ grows large, we test on CIFAR-10 dataset and report the OT retrieval gap in Figure 13. We note that the OT retrieval gap also upperbounds the approximation error in transportation plan, so Figure 13 also serves as empirical verification of Theorem 23 on CIFAR-10 data.

5.5.2 APPROXIMATION ERROR BETWEEN UOT AND OT DISTANCES

On CIFAR-10 data, we measure the gap $\mathbf{OT}(\mathbf{a}, \mathbf{b}) - \mathbf{UOT}_{\mathbf{KL}}(\mathbf{a}, \mathbf{b})$ of transport distance between OT and UOT (Theorem 26) in Figure 14. The result illustrates the sharpness of our theoretical bound and the trade-off between τ and approximation error, which can further elucidate more refined tuning of τ .

6. Conclusion

We have developed a new algorithm for solving the discrete UOT problem using gradient-based methods. Our method achieves the complexity of $\tilde{O}(\kappa n^2)$, where κ is the condition number depending on only the two input measures. This is significantly better than the only known complexity $O(\frac{\tau n^2 \log(n)}{\varepsilon} \log(\frac{\log(n)}{\varepsilon}))$ of Sinkhorn algorithm (Pham et al., 2020) to the UOT problem in terms of ε and τ dependency. In addition, we are the first to theoretically analyze the approximability of the UOT problem to find an OT solution. Our numerical results show the efficiency of the new method and the tightness of our theoretical bounds. We believe that our analysis framework can be extended to study the convergence results of stochastic methods, widely used in ML applications for handling large-scale datasets (Nguyen et al., 2018) or under nonlinear stochastic control settings (Nguyen and Maguluri, 2023). Our results on approximation error can open up directions that utilize UOT to approximate standard OT and even partial OT (Nguyen et al., 2024) using first-order methods.

APPENDIX

Appendix A. Sinkhorn Algorithm

A.1 Regularity Conditions

In this Section, we will show that the assumptions used by the Sinkhorn (Pham et al., 2020) algorithm are equivalent to our stated regularity conditions (S1)-(S4). (S4) is clearly stated in their list of regularity conditions. To claim their complexity, in (Pham et al., 2020, Corollary1), they further require the condition ' $R = O(\frac{1}{\eta}\|C\|_\infty)$ '. Following their discussion in (Pham et al., 2020, Section 3.1), a necessary condition for it to hold is $\|\log(\mathbf{a})\|_\infty = O(\log(n))$ (and resp. $\|\log(\mathbf{b})\|_\infty = O(\log(n))$). To see that $\|\log(\mathbf{a})\|_\infty = O(\log(n))$ (and resp. $\|\log(\mathbf{b})\|_\infty = O(\log(n))$) is equivalent to (S1)-(S3), we note that if $a_{\min} = 0$ then $\|\log(\mathbf{a})\|_\infty = \infty$ (contradiction!), and $\|\log(\mathbf{a})\|_\infty = \max\{|\log(a_{\min})|, |\log(a_{\max})|\}$.

A.2 The dependency of Sinkhorn complexity on $\log(\alpha + \beta)$

We note that in its pure form, Sinkhorn's complexity has $\log(\alpha + \beta)$, which is excluded in (Pham et al., 2020) under their regularity condition (S4) that α, β are constants. To see

this, note that in the proof of (Pham et al., 2020, Corollary1), their quantity U has the complexity $O((\alpha + \beta) \log(n))$ and appears in Sinkhorn’s final complexity as $\log(U)$.

Appendix B. Supplementary Lemmas and Theorems

Lemma 28 For $x, y \in \mathbb{R}$, we have: $|\max\{0, x\} - \max\{0, y\}| \leq |x - y|$

Proof We have:

$$\begin{aligned} \max\{0, x\} &= \max\{0, x - y + y\} \leq \max\{0, y\} + \max\{0, x - y\} \leq \max\{0, y\} + |x - y| \\ \max\{0, x\} - \max\{0, y\} &\leq |x - y| \end{aligned}$$

By symmetry, $\max\{0, y\} - \max\{0, x\} \leq |x - y|$. Thus, $|\max\{0, x\} - \max\{0, y\}| \leq |x - y|$. ■

Theorem 29 (Rockafellar (1967)) Let (E, E^*) and (F, F^*) be two couples of topologically paired spaces. Let $A : E \rightarrow F$ be a continuous linear operator and $A^* : F^* \rightarrow E^*$ its adjoint. Let f and g be lower semicontinuous and proper convex functions defined on E and F respectively. If there exists $x \in \text{dom} f$ such that g is continuous at Ax , then

$$\sup_{x \in E} -f(-x) - g(Ax) = \inf_{y^* \in F^*} f^*(A^* y^*) + g^*(y^*)$$

Moreover, if there exists a maximizer $x \in E$ then there exists $y^* \in F^*$ satisfying $Ax \in \partial g^*(y^*)$ and $A^* y^* \in \partial f(-x)$.

Lemma 30 (Lemma 7, (Altschuler et al., 2017)) For the inputs $X \in \mathbb{R}^{n \times n}$, $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{b} \in \mathbb{R}^n$, Algorithm 4 takes $O(n^2)$ time to output a matrix $Y \in \Pi(\mathbf{a}, \mathbf{b})$ satisfying:

$$\|Y - X\|_1 \leq 2[\|X \mathbf{1}_n - \mathbf{a}\|_1 + \|X^\top \mathbf{1}_n - \mathbf{b}\|_1] \quad (78)$$

Appendix C. Missing Proofs

C.1 Proof of Lemma 14

Let $\mathbf{x}^* = (\mathbf{u}^*, \mathbf{v}^*, \mathbf{t}^*)$ be an optimal solution to (11). Because of the projection step on line 8 of Algorithm 1, we have $\underline{\mathbf{x}}^t, \mathbf{x}^t \in V_D$ for $t = 0, \dots, k$. By Lemma 11, $f_\eta(\mathbf{x})$ is L -smooth and $w_\eta(\mathbf{x})$ is μ -strongly convex on $\text{conv}\{\mathbf{x}^*, \underline{\mathbf{x}}^0, \mathbf{x}^0, \underline{\mathbf{x}}^1, \mathbf{x}^1, \dots\}$. The GEM-UOT is based on the GEM being applied to solve the dual objective (38), which, from (Lan and Zhou, 2018, Theorem 2.1) and the proof therein, thus guarantees that the dual iterate $\underline{\mathbf{x}}^k$ converges to the dual solution \mathbf{x}^* of (38) by:

$$\|\underline{\mathbf{x}}^k - \mathbf{x}^*\|_2^2 \leq \frac{4\Delta_{0, \sigma_0} k \zeta^k (1 - \zeta)}{\mu(1 - \zeta^k)}. \quad (79)$$

Furthermore, we have:

$$\begin{aligned}
 \frac{k\zeta^k(1-\zeta)}{1-\zeta^k} &= \left(\sum_{t=1}^k \frac{\zeta^t}{\zeta^t} \right) \frac{\zeta^k(1-\zeta)}{1-\zeta^k} \leq \left(\sum_{t=1}^k \frac{\zeta^t}{\zeta^{1.5t}} \right) \frac{\zeta^k(1-\zeta)}{1-\zeta^k} \\
 &\leq \frac{1-\zeta^{k/2}}{\zeta^{k/2}(1-\zeta^{1/2})} \frac{\zeta^k(1-\zeta)}{1-\zeta^k} \\
 &\leq 2\zeta^{k/2}
 \end{aligned} \tag{80}$$

Now, plugging (80) into (79) gives:

$$\|\underline{\mathbf{x}}^k - \mathbf{x}^*\|_2^2 \leq \frac{8\Delta_{0,\sigma_0}\zeta^{k/2}}{\mu} \tag{81}$$

Line 16 of Algorithm 1 further converts the dual iterate $\underline{\mathbf{x}}^k$ into the primal iterate X^k as the final output. Next, we show that the convergence of the dual iterate in the sense of (81) leads to the convergence of the primal iterate X^k to the primal solution X^η of the primal objective (5). From the definition of X^k in Algorithm 1 and in view of Corollary 5, we have:

$$\begin{aligned}
 \|X^k - X^\eta\|_1 &= \frac{1}{2\eta} \sum_{i,j=1}^n |t_{ij}^k - t_{ij}^*| \\
 &\leq \frac{1}{2\eta} \|\underline{\mathbf{x}}^k - \mathbf{x}^*\|_1 \leq \frac{\sqrt{n^2 + 2n}}{2\eta} \|\underline{\mathbf{x}}^k - \mathbf{x}^*\|_2 \\
 &\stackrel{(81)}{\leq} \frac{\sqrt{2n}}{2\eta\mu} 2\sqrt{2}\Delta_{0,\sigma_0}^{1/2} \zeta^{k/4} = \frac{2n\Delta_{0,\sigma_0}^{1/2}}{\eta\mu} \left(1 - \frac{1}{1 + \sqrt{1 + 16L/\mu}}\right)^{k/4} \\
 &\leq \frac{2n\Delta_{0,\sigma_0}^{1/2}}{\eta\mu} \cdot \exp\left\{-\frac{k}{4(1 + \sqrt{1 + 16L/\mu})}\right\}.
 \end{aligned}$$

The condition that $\frac{2n\Delta_{0,\sigma_0}^{1/2}}{\eta\mu} \cdot \exp\left\{-\frac{k}{4(1 + \sqrt{1 + 16L/\mu})}\right\} \leq \min\left\{\frac{\varepsilon}{2L_1}, \frac{\min\{a_{min}, b_{min}\}e^{-D/\tau}}{2}, \frac{\alpha + \beta}{2}\right\}$ is equivalent to:

$$k \geq 4(1 + \sqrt{1 + 16L/\mu}) \log \left(\frac{4n\Delta_{0,\sigma_0}^{1/2}}{\eta\mu} \max \left\{ \frac{L_1}{\varepsilon}, \frac{e^{D/\tau}}{\min\{a_{min}, b_{min}\}}, (\alpha + \beta)^{-1} \right\} \right) = K_0.$$

Therefore,

$$\|X^k - X^\eta\|_1 \leq \min \left\{ \frac{\varepsilon}{2L_1}, \frac{\min\{a_{min}, b_{min}\}e^{-D/\tau}}{2}, \frac{\alpha + \beta}{2} \right\}. \tag{82}$$

C.2 Proof of Lemma 15

For $p = \frac{1}{2} \min\{a_{min}, b_{min}\}e^{-\frac{D}{\tau}}$ and $q = \alpha + \beta$ given in Section 3.1, we know that $X^\eta \in U_{p,q}$ by Lemma 9. We now proceed to prove that $X^k \in U_{p,q}$. Consider any $i \in [1, n]$:

$$\begin{aligned} \sum_{j=1}^n X_{ij}^k &\geq \sum_{j=1}^n X_{ij}^\eta - \left| \sum_{j=1}^n (X_{ij}^k - X_{ij}^\eta) \right| \geq \sum_{j=1}^n X_{ij}^\eta - \|X^k - X^\eta\|_1 \stackrel{(13)}{=} a_i e^{-u_i/\tau} - \|X^k - X^\eta\|_1 \\ &\stackrel{(22)}{\geq} a_{min} e^{-D/\tau} - \|X^k - X^\eta\|_1 \stackrel{(82)}{\geq} \frac{a_{min} e^{-D/\tau}}{2} \geq p, \\ \sum_{j=1}^n X_{ij}^k &\leq \sum_{j=1}^n X_{ij}^\eta + \left| \sum_{j=1}^n (X_{ij}^k - X_{ij}^\eta) \right| \leq \|X^\eta\|_1 + \|X^k - X^\eta\|_1 \stackrel{(18),(82)}{\leq} \frac{\alpha + \beta}{2} + \frac{\alpha + \beta}{2} = \alpha + \beta. \end{aligned}$$

Similarly, for any $j \in [1, n]$, we have: $p \leq \sum_{i=1}^n X_{ij}^k \leq q$, and thus conclude that $X^k \in U_{p,q}$.

C.3 Proof of Lemma 11

The gradient of $f_\eta(\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t}))$ can be computed as:

$$\begin{aligned} \frac{\partial f_\eta}{\partial u_i} &= -a_i e^{-u_i/\tau} - \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} u_i, \\ \frac{\partial f_\eta}{\partial v_j} &= -b_j e^{-v_j/\tau} - \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} v_j, \\ \frac{\partial f_\eta}{\partial t_{ij}} &= 0. \end{aligned}$$

The Hessian of $f_\eta(\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t}))$ can be computed as:

$$\begin{aligned} \frac{\partial^2 f_\eta}{\partial u_i \partial v_j} &= \frac{\partial^2 f_\eta}{\partial u_i \partial t_{kj}} = \frac{\partial^2 f_\eta}{\partial v_i \partial t_{kj}} = 0, \\ \frac{\partial^2 f_\eta}{\partial u_i^2} &= \frac{a_i}{\tau} e^{-u_i/\tau} - \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau}, \\ \frac{\partial^2 f_\eta}{\partial v_j^2} &= \frac{b_j}{\tau} e^{-v_j/\tau} - \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau}. \end{aligned}$$

If $\mathbf{x} \in V_D$, then $u_i, v_j \leq D$ for all $i, j \in [n]$ which implies $\frac{\partial^2 f_\eta}{\partial u_i^2}, \frac{\partial^2 f_\eta}{\partial v_j^2} \geq 0$ for all $i, j \in [n]$ and thereby $\nabla^2 f_\eta(\mathbf{x}) \succcurlyeq 0$ on V_D . We can then conclude that $f_\eta(\mathbf{x})$ is convex on V_D .

Now let us consider any $\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t}), \mathbf{x}' = (\mathbf{u}', \mathbf{v}', \mathbf{t}') \in V_D$. By Mean Value Theorem, $\exists c_i \in \text{range}(u_i, u'_i)$, such that $e^{-u_i/\tau} - e^{-u'_i/\tau} = -\frac{1}{\tau} e^{-c_i/\tau} (u_i - u'_i)$ and $\exists d_j \in \text{range}(v_j, v'_j)$, such that $e^{-v_j/\tau} - e^{-v'_j/\tau} = -\frac{1}{\tau} e^{-d_j/\tau} (v_j - v'_j)$. We then obtain:

$$\begin{aligned} \|\nabla f_\eta(\mathbf{x}) - \nabla f_\eta(\mathbf{x}')\|_2^2 &= \sum_{i=1}^n \left[\left(\frac{a_i}{\tau} e^{-c_i/\tau} - \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} \right) (u_i - u'_i) \right]^2 \\ &\quad + \sum_{j=1}^n \left[\left(\frac{b_j}{\tau} e^{-d_j/\tau} - \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} \right) (v_j - v'_j) \right]^2 \end{aligned}$$

$$\begin{aligned}
 &\leq \sum_{i=1}^n \left(\frac{a_i}{\tau} e^{-c_i/\tau} + \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} \right)^2 (u_i - u'_i)^2 \\
 &\quad + \sum_{j=1}^n \left(\frac{b_j}{\tau} e^{-d_j/\tau} + \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} \right)^2 (v_j - v'_j)^2 \\
 &\leq \sum_{i=1}^n \left(\frac{a_i}{\tau} \frac{\alpha + \beta}{2a_i} + \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} \right)^2 (u_i - u'_i)^2 \\
 &\quad + \sum_{j=1}^n \left(\frac{b_j}{\tau} \frac{\alpha + \beta}{2b_j} + \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} \right)^2 (v_j - v'_j)^2 \\
 &\quad \text{(since } c_i \in \text{range}(u_i, u'_i), c_i \geq \tau \log(\frac{2a_i}{\alpha + \beta}). \text{ Similarly, } d_j \geq \tau \log(\frac{2b_j}{\alpha + \beta})) \\
 &\leq \left(\frac{\alpha + \beta}{2\tau} + \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} \right)^2 \|\mathbf{x} - \mathbf{x}'\|_2^2. \\
 \|\nabla f_\eta(\mathbf{x}) - \nabla f_\eta(\mathbf{x}')\|_2 &\leq \left(\frac{\alpha + \beta}{2\tau} + \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} \right) \|\mathbf{x} - \mathbf{x}'\|_2,
 \end{aligned}$$

which implies that $f_\eta(\mathbf{x})$ is L -smooth with $L = \frac{\alpha + \beta}{2\tau} + \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau}$.

Now let us consider $w_\eta(\mathbf{x})$. The gradient $w_\eta(\mathbf{x})$ can be computed as:

$$\frac{\partial w}{\partial u_i} = \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} u_i, \quad \frac{\partial w}{\partial v_j} = \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} v_j, \quad \frac{\partial w}{\partial t_{ij}} = \frac{1}{2\eta} t_{ij}.$$

For any $\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t})$, $\mathbf{x}' = (\mathbf{u}', \mathbf{v}', \mathbf{t}') \in V_D$, we have:

$$\begin{aligned}
 \langle \nabla w_\eta(\mathbf{x}) - \nabla w_\eta(\mathbf{x}'), \mathbf{x} - \mathbf{x}' \rangle &= \sum_{i=1}^n \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} (u_i - u'_i)^2 \\
 &\quad + \sum_{j=1}^n \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau} (v_j - v'_j)^2 + \sum_{i,j=1}^n \frac{1}{2\eta} (t_{ij} - t'_{ij})^2 \\
 &\geq \min \left\{ \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau}, \frac{1}{2\eta} \right\} \|\mathbf{x} - \mathbf{x}'\|_2^2.
 \end{aligned}$$

Therefore, $w_\eta(\mathbf{x})$ is μ -strongly convex on V_D with $\mu = \min \left\{ \frac{\min\{a_{min}, b_{min}\}}{\tau} e^{-D/\tau}, \frac{1}{2\eta} \right\}$.

C.4 Projection Step of GEM-UOT

We describe the following projection step (i.e. line 8 in Algorithm 1) in more details:

$$\mathbf{x}^t = \mathcal{M}_{V_D \cap \mathcal{X}}(\tilde{\mathbf{y}}^t, \mathbf{x}^{t-1}, \rho) = \underset{\mathbf{x}=(\mathbf{u}, \mathbf{v}, \mathbf{t}) \in V_D \cap \mathcal{X}}{\operatorname{argmin}} \{G(\mathbf{x}) := \langle \tilde{\mathbf{y}}^t, \mathbf{x} \rangle + w_\eta(\mathbf{x}) + \rho P(\mathbf{x}^{t-1}, \mathbf{x})\} \quad (83)$$

where we further expand the objective $G(\mathbf{x})$ as:

$$G(\mathbf{x}) = \langle \tilde{\mathbf{y}}^t, \mathbf{x} \rangle + w_\eta(\mathbf{x}) + \frac{\rho}{\mu} [w_\eta(\mathbf{x}) - w_\eta(\mathbf{x}^{t-1}) - \langle \nabla w_\eta(\mathbf{x}^{t-1}), \mathbf{x} - \mathbf{x}^{t-1} \rangle]$$

$$= \left(1 + \frac{\rho}{\mu}\right) w_\eta(\mathbf{x}) + \langle \tilde{\mathbf{y}}^t - \frac{\rho}{\mu} \nabla w_\eta(\mathbf{x}^{t-1}), \mathbf{x} \rangle + \frac{\rho}{\mu} \langle \nabla w_\eta(\mathbf{x}^{t-1}), \mathbf{x}^{t-1} \rangle. \quad (84)$$

For abbreviation, we let $\tilde{\mathbf{y}}^t - \frac{\rho}{\mu} \nabla w_\eta(\mathbf{x}^{t-1}) = (\mathbf{u}^c, \mathbf{v}^c, \mathbf{t}^c) \in \mathbb{R}^{n^2+2n}$ where $\mathbf{u}^c \in \mathbb{R}^n$, $\mathbf{v}^c \in \mathbb{R}^n$ and $\mathbf{t}^c \in \mathbb{R}^{n^2}$ respectively are coefficients of the variables $\mathbf{u}, \mathbf{v}, \mathbf{t}$ in the linear term of (84). We also define the constants $q_1 = \left(1 + \frac{\rho}{\mu}\right) \cdot \frac{\min\{a_{\min}, b_{\min}\}}{2\tau} e^{-D/\tau}$ and $q_2 = \left(1 + \frac{\rho}{\mu}\right) / (4\eta)$. The objective $G(\mathbf{x})$ can be written as the sum of quadratic functions as follows:

$$G(\mathbf{x}) = (q_1 \|\mathbf{u}\|_2^2 + \langle \mathbf{u}^c, \mathbf{u} \rangle) + (q_1 \|\mathbf{v}\|_2^2 + \langle \mathbf{v}^c, \mathbf{v} \rangle) + (q_2 \|\mathbf{t}\|_2^2 + \langle \mathbf{t}^c, \mathbf{t} \rangle) + \frac{\rho}{\mu} \langle \nabla w_\eta(\mathbf{x}^{t-1}), \mathbf{x}^{t-1} \rangle. \quad (85)$$

And the projection step is thus equivalent to solving:

$$\min_{\mathbf{x}=(\mathbf{u}, \mathbf{v}, \mathbf{t}) \in V_D \cap \mathcal{X}} G(\mathbf{x}) \quad (86)$$

where $V_D \cap \mathcal{X} = \{(\mathbf{u}, \mathbf{v}, \mathbf{t}) \mid \mathbf{u}, \mathbf{v} \in \mathbb{R}^n, \mathbf{t} \in \mathbb{R}^{n \times n} : \forall i, j \in [1, n],$

$$\tau \log\left(\frac{2a_i}{\alpha + \beta}\right) \leq u_i \leq D$$

$$\tau \log\left(\frac{2b_j}{\alpha + \beta}\right) \leq v_j \leq D$$

$$0 \leq t_{ij},$$

$$u_i + v_j - C_{ij} \leq t_{ij}\}.$$

Next, we aim to derive a simplified equivalence form of (86) that is obtained by expressing the optimal \mathbf{t} with respect to fixed \mathbf{u}, \mathbf{v} , and consequently optimizing over \mathbf{u}, \mathbf{v} . In particular, we consider the mapping $\mathbf{t}(\mathbf{u}, \mathbf{v}) : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{n^2}$ defined by:

$$t_{ij}(\mathbf{u}, \mathbf{v}) = \max\left\{0, -\frac{t_{ij}^c}{q_2}, u_i + v_j - C_{ij}\right\}, \quad \forall i, j \in [n], \quad (87)$$

the restricted domain \tilde{V}_D as follows:

$$\tilde{V}_D = \{(\mathbf{u}, \mathbf{v}) \mid \mathbf{u}, \mathbf{v} \in \mathbb{R}^n : \forall i \in [n],$$

$$\tau \log\left(\frac{2a_i}{\alpha + \beta}\right) \leq u_i \leq D$$

$$\tau \log\left(\frac{2b_i}{\alpha + \beta}\right) \leq v_i \leq D\}, \quad (88)$$

and the objective:

$$\tilde{G}(\mathbf{u}, \mathbf{v}) = G(\mathbf{u}, \mathbf{v}, \mathbf{t}(\mathbf{u}, \mathbf{v})). \quad (89)$$

The next Lemma establishes the equivalence form for the projection problem (86).

Lemma 31 *The following holds:*

$$\min_{\mathbf{x}=(\mathbf{u}, \mathbf{v}, \mathbf{t}) \in V_D \cap \mathcal{X}} G(\mathbf{x}) = \min_{(\mathbf{u}, \mathbf{v}) \in \tilde{V}_D} \tilde{G}(\mathbf{u}, \mathbf{v}) \quad (90)$$

Proof For any $(\mathbf{u}, \mathbf{v}) \in \tilde{V}_D$, we have $\mathbf{t}(\mathbf{u}, \mathbf{v})$ that is defined as in (87) satisfies:

$$\begin{aligned} 0 &\leq t_{ij}(\mathbf{u}, \mathbf{v}), \\ u_i + v_j - C_{ij} &\leq t_{ij}(\mathbf{u}, \mathbf{v}), \end{aligned}$$

Thus, we have $(\mathbf{u}, \mathbf{v}, \mathbf{t}(\mathbf{u}, \mathbf{v})) \in V_D$, which implies:

$$\begin{aligned} \min_{\mathbf{x}=(\mathbf{u}, \mathbf{v}, \mathbf{t}) \in V_D \cap \mathcal{X}} G(\mathbf{x}) &\leq G(\mathbf{u}, \mathbf{v}, \mathbf{t}(\mathbf{u}, \mathbf{v})) = \tilde{G}(\mathbf{u}, \mathbf{v}), \quad \forall (\mathbf{u}, \mathbf{v}) \in \tilde{V}_D \\ \therefore \min_{\mathbf{x}=(\mathbf{u}, \mathbf{v}, \mathbf{t}) \in V_D \cap \mathcal{X}} G(\mathbf{x}) &\leq \min_{(\mathbf{u}, \mathbf{v}) \in \tilde{V}_D} \tilde{G}(\mathbf{u}, \mathbf{v}). \end{aligned} \quad (91)$$

Now, take any $\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t}) \in V_D \cap \mathcal{X}$. We fix \mathbf{u} and \mathbf{v} , and optimize over all \mathbf{t}' such that $(\mathbf{u}, \mathbf{v}, \mathbf{t}') \in V_D \cap \mathcal{X}$, i.e. those \mathbf{t}' with $0 \leq t'_{ij}$ and $u_i + v_j - C_{ij} \leq t'_{ij}$. Then, we have:

$$G(\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t})) \geq \min_{\mathbf{t}' : (\mathbf{u}, \mathbf{v}, \mathbf{t}') \in V_D \cap \mathcal{X}} G(\mathbf{u}, \mathbf{v}, \mathbf{t}'). \quad (92)$$

Note that the optimization problem over \mathbf{t}' in the RHS of (92) has separable quadratic objective and separable linear constraints. Thus, simple algebra gives $\mathbf{t}(\mathbf{u}, \mathbf{v})$ being defined in (87) as the optimal solution to such problem, i.e.

$$\mathbf{t}(\mathbf{u}, \mathbf{v}) = \operatorname{argmin}_{\mathbf{t}' : (\mathbf{u}, \mathbf{v}, \mathbf{t}') \in V_D \cap \mathcal{X}} G(\mathbf{u}, \mathbf{v}, \mathbf{t}'). \quad (93)$$

From (92) and (93), we have:

$$\begin{aligned} G(\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t})) &\geq G(\mathbf{u}, \mathbf{v}, \mathbf{t}(\mathbf{u}, \mathbf{v})) = \tilde{G}(\mathbf{u}, \mathbf{v}) \\ &\geq \min_{(\mathbf{u}, \mathbf{v}) \in \tilde{V}_D} \tilde{G}(\mathbf{u}, \mathbf{v}) \\ \therefore \min_{\mathbf{x}=(\mathbf{u}, \mathbf{v}, \mathbf{t}) \in V_D \cap \mathcal{X}} G(\mathbf{x}) &\geq \min_{(\mathbf{u}, \mathbf{v}) \in \tilde{V}_D} \tilde{G}(\mathbf{u}, \mathbf{v}), \end{aligned} \quad (94)$$

where the second inequality is because any $\mathbf{x} = (\mathbf{u}, \mathbf{v}, \mathbf{t}) \in V_D \cap \mathcal{X}$ implies $(\mathbf{u}, \mathbf{v}) \in \tilde{V}_D$ by definition. Finally, combining (91) and (94) gives the desired statement. \blacksquare

Since the equivalence form $\min_{(\mathbf{u}, \mathbf{v}) \in \tilde{V}_D} \tilde{G}(\mathbf{u}, \mathbf{v})$ of the projection step is strongly convex smooth optimization over box constraints, we can adopt projected quasi-Newton method with superlinear convergence rate (Schmidt et al., 2011, Theorem 11.2), thereby resulting in $\tilde{O}(1)$ iteration complexity. Furthermore, the per iteration cost is dominated by the gradient and Hessian computation of \tilde{G} , which can be achieved efficiently in $O(n^2)$ as follows. In particular, we have for $\forall i, j \in [n]$:

$$\begin{aligned} \frac{\partial \tilde{G}}{\partial u_i} &= 2q_1 u_i + u_i^c + 2 \sum_{j=1}^n \max \left\{ 0, \frac{t_{ij}^c}{q_2}, u_i + v_j - C_{ij} + \frac{t_{ij}^c}{q_2} \right\}, \\ \frac{\partial \tilde{G}}{\partial v_j} &= 2q_1 v_j + v_j^c + 2 \sum_{i=1}^n \max \left\{ 0, \frac{t_{ij}^c}{q_2}, u_i + v_j - C_{ij} + \frac{t_{ij}^c}{q_2} \right\}, \end{aligned}$$

$$\begin{aligned}\frac{\partial^2 \tilde{G}}{\partial u_i \partial u_j} &= \begin{cases} 2q_1 + 2 \sum_{k=1}^n \mathbb{1}\{u_i + v_k - C_{ik} \geq \max\{0, -\frac{t_{ik}^c}{q_2}\}\}, & \text{if } i = j \\ 0, & \text{otherwise} \end{cases}, \\ \frac{\partial^2 \tilde{G}}{\partial v_i \partial v_j} &= \begin{cases} 2q_1 + 2 \sum_{k=1}^n \mathbb{1}\{u_k + v_j - C_{kj} \geq \max\{0, -\frac{t_{kj}^c}{q_2}\}\}, & \text{if } i = j \\ 0, & \text{otherwise} \end{cases}, \\ \frac{\partial^2 \tilde{G}}{\partial u_i \partial v_j} &= 2 \mathbb{1}\{u_i + v_j - C_{ij} \geq \max\{0, -\frac{t_{ij}^c}{q_2}\}\}.\end{aligned}$$

Thus, we conclude that the projection step can be solved in $\tilde{O}(n^2)$ time.

C.5 Proof of Lemma 17

The gradient of $h_a(\mathbf{x} = (\mathbf{u}, \mathbf{v}))$ can be computed as:

$$\begin{aligned}\frac{\partial h_a}{\partial u_i} &= -a_i e^{-u_i/\tau} + \sum_{j=1}^n \frac{\max\{0, u_i + v_j - C_{ij}\}}{2\eta}, \\ \frac{\partial h_a}{\partial v_j} &= -b_j e^{-v_j/\tau} + \sum_{i=1}^n \frac{\max\{0, u_i + v_j - C_{ij}\}}{2\eta}.\end{aligned}$$

Let us consider any $\mathbf{x} = (\mathbf{u}, \mathbf{v}), \mathbf{x}' = (\mathbf{u}', \mathbf{v}') \in V_a$. By Mean Value Theorem, $\exists c_i \in \text{range}(u_i, u'_i)$, such that $e^{-u_i/\tau} - e^{-u'_i/\tau} = -\frac{1}{\tau} e^{-c_i/\tau} (u_i - u'_i)$ and $\exists d_j \in \text{range}(v_j, v'_j)$, such that $e^{-v_j/\tau} - e^{-v'_j/\tau} = -\frac{1}{\tau} e^{-d_j/\tau} (v_j - v'_j)$. We then obtain:

$$\begin{aligned}\|\nabla h_a(\mathbf{x}) - \nabla h_a(\mathbf{x}')\|_2^2 &\leq 2 \sum_{i=1}^n \left[\frac{a_i}{\tau} e^{-c_i/\tau} (u_i - u'_i) \right]^2 + 2 \sum_{j=1}^n \left[\frac{b_j}{\tau} e^{-d_j/\tau} (v_j - v'_j) \right]^2 \\ &+ \sum_{i=1}^n \frac{2}{4\eta^2} \left[\sum_{j=1}^n (u_i - u'_i + v_j - v'_j) \right]^2 + \sum_{j=1}^n \frac{2}{4\eta^2} \left[\sum_{i=1}^n (u_i - u'_i + v_j - v'_j) \right]^2 \\ &\text{(By Lemma 28 and simple algebra } (x+y)^2 \leq 2(x^2 + y^2)\text{)} \\ &\leq \sum_{i=1}^n 2 \left(\frac{a_i}{\tau} e^{-c_i/\tau} \right)^2 (u_i - u'_i)^2 + \sum_{j=1}^n 2 \left(\frac{b_j}{\tau} e^{-d_j/\tau} \right)^2 (v_j - v'_j)^2 \\ &+ \sum_{i=1}^n \frac{1}{\eta^2} \left\{ \sum_{j=1}^n [(u_i - u'_i)^2 + (v_j - v'_j)^2] \right\} + \sum_{j=1}^n \frac{1}{\eta^2} \left\{ \sum_{i=1}^n [(u_i - u'_i)^2 + (v_j - v'_j)^2] \right\} \\ &\leq \sum_{i=1}^n 2 \left(\frac{a_i}{\tau} \frac{\alpha + \beta}{2a_i} \right)^2 (u_i - u'_i)^2 + \sum_{j=1}^n 2 \left(\frac{b_j}{\tau} \frac{\alpha + \beta}{2b_j} \right)^2 (v_j - v'_j)^2 + \frac{2n}{\eta^2} \|\mathbf{x} - \mathbf{x}'\|_2^2 \\ &\text{(since } c_i \in \text{range}(u_i, u'_i), c_i \geq \tau \log(\frac{2a_i}{\alpha + \beta}). \text{ Similarly, } d_j \geq \tau \log(\frac{2b_j}{\alpha + \beta})) \\ &= \left[\frac{1}{2} \cdot \left(\frac{\alpha + \beta}{\tau} \right)^2 + \frac{2n}{\eta^2} \right] \|\mathbf{x} - \mathbf{x}'\|_2^2. \\ \therefore \|\nabla h_a(\mathbf{x}) - \nabla h_a(\mathbf{x}')\|_2 &\leq \left(\frac{\alpha + \beta}{\tau} + \frac{2\sqrt{n}}{\eta} \right) \cdot \|\mathbf{x} - \mathbf{x}'\|_2.\end{aligned}$$

This implies that $h_a(\mathbf{x})$ is L -smooth with $L_a = \frac{\alpha+\beta}{\tau} + \frac{2\sqrt{n}}{\eta}$.

Appendix D. Experimental Result for GEM-RUOT on Fashion-MNIST dataset

We test GEM-UOT (Algorithm 1), GEM-RUOT (Algorithm 2) and Sinkhorn on the Fashion-MNIST dataset using the same setting as the experiment on CIFAR dataset in Section 5.2 and report the result in Figure 15.




Figure 15: Comparison of f primal gap between the Sinkhorn algorithm and GEM algorithms for 1500 iterations.

References

- Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.
- J. Altschuler, F. Bach, A. Rudi, and J. Weed. Approximating the quadratic transportation metric in near-linear time. *ArXiv Preprint: 1810.10046*, 2018.
- Jason Altschuler, Jonathan Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. In *Advances in Neural Information Processing Systems*, pages 1964–1974, 2017.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.
- Yogesh Balaji, Rama Chellappa, and Soheil Feizi. Robust optimal transport with applications in generative modeling and domain adaptation. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural*

- Information Processing Systems*, volume 33, pages 12934–12944. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/9719a00ed0c5709d80dfef33795dcef3-Paper.pdf.
- Espen Bernton, Promit Ghosal, and Marcel Nutz. Entropic optimal transport: Geometry and large deviations. *Duke Mathematical Journal*, 171(16):3363 – 3400, 2022. doi: 10.1215/00127094-2022-0035. URL <https://doi.org/10.1215/00127094-2022-0035>.
- M. Blondel, V. Seguy, and A. Rolet. Smooth and sparse optimal transport. In *AISTATS*, pages 880–889, 2018.
- Luis A. Caffarelli and Robert J. McCann. Free boundaries in optimal transport and monge-ampère obstacle problems. *Annals of Mathematics*, 171(2):673–730, 2010. ISSN 0003486X. URL <http://www.jstor.org/stable/20752228>.
- Laetitia Chapel, Rémi Flamary, Haoran Wu, Cédric Févotte, and Gilles Gasso. Unbalanced optimal transport through non-negative penalized linear regression. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 23270–23282. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/c3c617a9b80b3ae1ebd868b0017cc349-Paper.pdf.
- Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of Computation*, 87(314):2563–2609, 2018.
- N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9): 1853–1865, 2017.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/af21d0c97db2e27e13572cbf59eb343d-Paper.pdf.
- Henri De Plaen, Pierre-François De Plaen, Johan A. K. Suykens, Marc Proesmans, Tinne Tuytelaars, and Luc Van Gool. Unbalanced optimal transport: A unified framework for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3198–3207, June 2023.
- Yihe Dong, Yu Gao, Richard Peng, Ilya Razenshteyn, and Saurabh Sawlani. A study of performance of optimal transport. *ArXiv Preprint: 2005.01182*, 2020.
- P. Dvurechensky, A. Gasnikov, and A. Kroshnin. Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn’s algorithm. In *International conference on machine learning*, pages 1367–1376, 2018.

- Montacer Essid and Justin Solomon. Quadratically regularized optimal transport on graphs. *SIAM Journal on Scientific Computing*, 40(4):A1961–A1986, 2018. doi: 10.1137/17M1132665. URL <https://doi.org/10.1137/17M1132665>.
- Kilian Fatras, Thibault Sejourne, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 3186–3197. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/fatras21a.html>.
- Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya-Polo, and Tomaso Poggio. Learning with a wasserstein loss. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’15, page 2053–2061, Cambridge, MA, USA, 2015. MIT Press.
- Aude Genevay, Gabriel Peyre, and Marco Cuturi. Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617, 2018.
- Sergey Guminov, Pavel E. Dvurechensky, Nazarii Tupitsa, and Alexander V. Gasnikov. On a combination of alternating minimization and nesterov’s momentum. In *International Conference on Machine Learning*, 2021.
- Nhat Ho, XuanLong Nguyen, Mikhail Yurochkin, Hung Hai Bui, Viet Huynh, and Dinh Phung. Multilevel clustering via Wasserstein means. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1501–1509. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/ho17a.html>.
- L. V. Kantorovich. On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)*, volume 37, pages 199–201, 1942.
- A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- Guanghui Lan and Yi Zhou. Random gradient extrapolation for distributed and stochastic optimization. *SIAM Journal on Optimization*, 28(4):2753–2782, 2018.
- Khang Le, Huy Nguyen, Quang Minh Nguyen, Tung Pham, Hung Bui, and Nhat Ho. On robust optimal transport: Computational complexity and barycenter computation. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=xRLT28nn1FV>.
- John Lee, Nicholas P Bertrand, and Christopher J Rozell. Parallel unbalanced optimal transport regularization for large scale imaging problems. *ArXiv Preprint: 1909.00149*, 2019.

- Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\text{vrank})$ iterations and faster algorithms for maximum flow. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 424–433, 2014. doi: 10.1109/FOCS.2014.52.
- Matthias Liero, Alexander Mielke, and Giuseppe Savaré. Optimal entropy-transport problems and a new hellinger–kantorovich distance between positive measures. *Inventiones mathematicae*, 211(3):969–1117, Dec 2017. ISSN 1432-1297. doi: 10.1007/s00222-017-0759-8. URL <http://dx.doi.org/10.1007/s00222-017-0759-8>.
- Dirk A. Lorenz, Paul Manns, and Christian Meyer. Quadratically regularized optimal transport. *Applied Mathematics & Optimization*, 83:1919 – 1949, 2019. URL <https://api.semanticscholar.org/CorpusID:85517311>.
- Yiling Luo, Yiling Xie, and Xiaoming Huo. Improved rate of first order algorithms for entropic optimal transport. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 2723–2750. PMLR, 25–27 Apr 2023. URL <https://proceedings.mlr.press/v206/luo23a.html>.
- Yura Malitsky and Konstantin Mishchenko. Adaptive gradient descent without descent. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 6702–6712, 2020.
- Boris Muzellec, Richard Nock, Giorgio Patrini, and Frank Nielsen. Tsallis regularized optimal transport and ecological inference. In *AAAI Conference on Artificial Intelligence*, 2016.
- Anh Duc Nguyen, Tuan Dung Nguyen, Quang Minh Nguyen, Hoang H. Nguyen, and Kim-Chuan Toh. On partial optimal transport: Revising the infeasibility of sinkhorn and efficient gradient methods. In *AAAI Conference on Artificial Intelligence*, 2024. URL <https://arxiv.org/pdf/2312.13970.pdf>.
- Hoang Huy Nguyen and Siva Theja Maguluri. Stochastic approximation for nonlinear discrete stochastic control: Finite-sample bounds. *ArXiv Preprint: 2304.11854*, 2023.
- Khai Nguyen, Nhat Ho, Tung Pham, and Hung Bui. Distributional sliced-wasserstein and applications to generative modeling. In *International Conference on Learning Representations*, 2021.
- Lam Nguyen, Phuong Ha Nguyen, Marten van Dijk, Peter Richtarik, Katya Scheinberg, and Martin Takac. SGD and hogwild! Convergence without the bounded gradients assumption. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3750–3758. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/nguyen18c.html>.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

- Khiem Pham, Khang Le, Nhat Ho, Tung Pham, and Hung Bui. On unbalanced optimal transport: An analysis of Sinkhorn algorithm. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7673–7682. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/pham20a.html>.
- François Pitié, Anil C. Kokaram, and Rozenn Dahyot. Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding*, 107(1):123–137, 2007. ISSN 1077-3142. doi: <https://doi.org/10.1016/j.cviu.2006.11.011>. URL <https://www.sciencedirect.com/science/article/pii/S1077314206002189>. Special issue on color image processing.
- Julien Rabin and Nicolas Papadakis. Convex color image segmentation with optimal transport distances. In *SSVM*, 2015.
- R. Tyrrell Rockafellar. Duality and stability in extremum problems involving convex functions. *Pacific Journal of Mathematics*, 21(1):167 – 187, 1967. doi: [pjm/1102992608](https://doi.org/10.1016/j.cviu.2006.11.011). URL <https://doi.org/>.
- G. Schiebinger et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176:928–943, 2019.
- Mark Schmidt, Dongmin Kim, and Suvrit Sra. Projected Newton-type Methods in Machine Learning. In *Optimization for Machine Learning*. The MIT Press, 09 2011. ISBN 9780262298773. doi: 10.7551/mitpress/8996.003.0013.
- Thibault Sejourne, Francois-Xavier Vialard, and Gabriel Peyré. Faster unbalanced optimal transport: Translation invariant sinkhorn and 1-d frank-wolfe. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 4995–5021. PMLR, 28–30 Mar 2022. URL <https://proceedings.mlr.press/v151/sejourne22a.html>.
- Cédric Villani. *Optimal transport: Old and New*. Springer, 2008.
- Yiling Xie, Yiling Luo, and Xiaoming Huo. An accelerated stochastic algorithm for solving the optimal transport problem. *ArXiv Preprint: 2203.00813*, 2023.
- K. D. Yang and C. Uhler. Scalable unbalanced optimal transport using generative adversarial networks. In *ICLR*, 2019.
- Bohan Zhou and Matthew Parno. Efficient and exact multimarginal optimal transport with pairwise costs. *ArXiv Preprint: 2208.03025*, 2023.