

Double-Bounded Optimal Transport for Advanced Clustering and Classification

Liangliang Shi, Zhaoqi Shen, Junchi Yan*

Department of Computer Science and Engineering,
MoE Key Lab of Artificial Intelligence, Shanghai Jiao Tong University
{shiliangliang, shenzhaoqi2271621336, yanjunchi}@sjtu.edu.cn

Abstract

Optimal transport (OT) is attracting increasing attention in machine learning. It aims to transport a source distribution to a target one at minimal cost. In its vanilla form, the source and target distributions are predetermined, which contrasts to the real-world case involving undetermined targets. In this paper, we propose Doubly Bounded Optimal Transport (DB-OT), which assumes that the target distribution is restricted within two boundaries instead of a fixed one, thus giving more freedom for the transport to find solutions. Based on the entropic regularization of DB-OT, three scaling-based algorithms are devised for calculating the optimal solution. We also show that our DB-OT is helpful for barycenter-based clustering, which can avoid the excessive concentration of samples in a single cluster. Then we further develop DB-OT techniques for long-tailed classification which is an emerging and open problem. We first propose a connection between OT and classification, that is, in the classification task, training involves optimizing the Inverse OT to learn the representations, while testing involves optimizing the OT for predictions. With this OT perspective, we first apply DB-OT to improve the loss, and the Balanced Softmax is shown as a special case. Then we apply DB-OT for inference in the testing process. Even with vanilla Softmax trained features, our extensive experimental results show that our method can achieve good results with our improved inference scheme in the testing stage.

Introduction

Optimal transport (OT) (Cuturi 2013) has been widely applied in machine learning (Cui et al. 2019a; Wang et al. 2013). For instance, Wasserstein distance (Arjovsky, Chintala, and Bottou 2017; Gulrajani et al. 2017) is applied with the dual form of OT to minimize the gap between the generated and real distributions via min-max adversarial optimization (Li et al. 2022). SwAV (Caron et al. 2020) employs the Sinkhorn algorithm for online clustering in self-supervised contrastive learning (Chen et al. 2020; Khosla et al. 2020; Wang and Liu 2021). OT-LDA (Huynh, Zhao, and Phung 2020) devises a Wasserstein barycenter-based topic model, while in (Xu et al. 2019), Gromov-Wasserstein distance is used for graph matching (Wang, Yan, and Yang

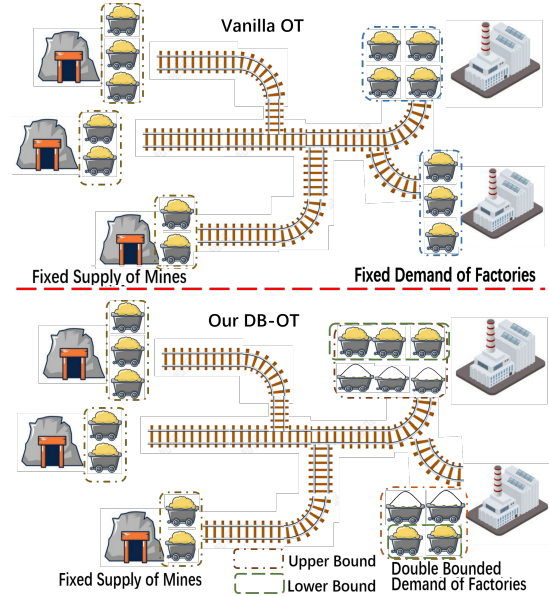


Figure 1: Illustration for the difference between vanilla OT and our DB-OT using the example of mines and factories as source and target, respectively. Vanilla OT assumes the equivalence between the supply and demand. In our DB-OT, we assume that the demand of the factory is not a fixed value, but rather a certain range by upper and lower bounds.

2019, 2021; Hu et al. 2020). These works all assume the source and target distributions are fixed.

However, in many real-world cases, the source or target distribution often varies over time and becomes uncertain. As shown in Fig. 1, in vanilla OT, it assumes that the supply from mines and the demand from factories are fixed and their total masses are equal to each other. While in our realistic setting we assume the demand from factories vary within a double-bounded range. When the demand exceeds the upper boundary, the factory reaches its production capacity limit. Conversely, when the demand falls below the lower boundary, the factory operates below its optimal capacity.¹

¹When both sides are not fixed, it degenerates to the lower bound. We discuss it in details in Sec. .

*Correspondence author.

Unbalanced OT (Chizat et al. 2018) is an OT variant to mitigate this inconsistency between source and target distributions. It relaxes the equality by only penalizing marginal deviation using divergences e.g. l_1 , l_2 , KL divergence, and etc. This method transforms the original hard constraints (equations) into penalty terms, allowing for a relaxation of the constraints and increasing the degree of freedom in transportation. However, the degree of relaxation can only be controlled via the penalty coefficient, which is not directly concerning the desired range of transportation results.

In this paper, we propose and term a new variant of OT called Double Bounded Optimal Transport (DB-OT), which allows the sum of coupling column to be in a range rather than a number, i.e., replacing the original equality constraint with two inequality constraints. Under this formulation, we further provide theoretical results with proof for DB-OT. Additionally, we propose three Sinkhorn-like algorithms based on entropic DB-OT. We also show that our results also inspire downstream applications.

Inspired by the OT-based topic model in OT-LDA (Huynh, Zhao, and Phung 2020), we propose a barycenter-based clustering method using DB-OT. This method allows for control of the sample quantity in each cluster, thereby avoiding the clustering of isolated or minority points within a cluster (Wang and Su 2011). Note that our barycenter-based clustering method is not limited to Euclidean space but can also be extended to other spaces, such as the Wasserstein space (Agueh and Carlier 2011). This extension allows our method to be generalized to handle more complex data, including text data.

Another application is using DB-OT for unbalanced image recognition. Specifically, to our best knowledge, we uncover the relation between classification and OT: training a classifier is equivalent to optimizing Inverse OT (Shi et al. 2023a; Stuart and Wolfram 2020; Chiu, Wang, and Shafto 2022), while inference in testing corresponds to optimizing OT problem. With this view, we first apply DB-OT to improve the loss for long-tailed data training. Different from traditional view using Softmax-based cross-entropy loss for learning, we increase the column-sum constraints for learning the representations and thus the Softmax generalizes to Sinkhorn-like iterations and the well-known balanced Softmax (Ren et al. 2020) can be a special case with only one iteration. Then we apply our DB-OT technique to inference in testing process and our approach is agnostic to the way how the classifier is trained. We successfully adopt our DB-OT inference technique to long-tailed, uniform, and reverse long-tailed datasets. **The contributions of our work are:**

1) We propose DB-OT to handle the case that the targets of transportation fall in a double-bounded range, which is an important yet under-studied setting for practical machine learning, e.g. cluster size-controllable clustering and long-tailed classification. We formulate DB-OT with the Kantorovich form by replacing the column equality constraints with double-bounded inequalities.

2) We solve DB-OT using entropic regularization and provide theoretical results, such as the static Schrödinger form, solution property, and dual form. To calculate the optimal solution, we propose a Bregman iterative algorithm based on

the static Schrödinger form. Additionally, we develop matrix scaling-based methods using Lagrange methods, with two space-efficient variants derived from the primal and dual forms of the problem, respectively. Note that the motivation behind proposing these three algorithms is to align with the vanilla entropic OT for DB-OT.

3) We propose our barycenter-based method for clustering using DB-OT. It helps control the size of cluster and thereby avoiding the unwanted over-scattered clustering with isolated or very few samples in clusters. Besides, compared with the previous barycenter-based topic model, a reweighting modification is also proposed to avoid the bias for barycenter calculation especially in Euclidean space.

4) Finally, our DB-OT is employed in the long-tailed classification. From the OT perspective, we observe that training can be viewed as optimizing the Inverse Optimal Transport (IOT), while inference in the testing can be regarded as minimizing the OT itself. Building upon this insight, we apply DB-OT to enhance the classification loss during training and also utilize it in the testing process based on a trained model.

Preliminaries and Related Work

Basics of Optimal Transport

As originated from (Kantorovich 1942), the Kantorovich's Optimal Transport is to solve a linear program, which is widely used for many classical problems such as matching (Wang et al. 2013). Specifically, given the cost matrix \mathbf{C} and two histograms (\mathbf{a}, \mathbf{b}) , Kantorovich's OT involves solving the coupling \mathbf{P} (i.e., the joint probability matrix) by

$$\min_{\mathbf{P} \in U(\mathbf{a}, \mathbf{b})} \langle \mathbf{C}, \mathbf{P} \rangle, \quad (1)$$

where $U(\mathbf{a}, \mathbf{b}) = \{\mathbf{P} \in R_{mn}^+ | \mathbf{P}\mathbf{1}_n = \mathbf{a}, \mathbf{P}^\top \mathbf{1}_m = \mathbf{b}\}$. Relaxing with the entropic regularization (Wilson 1969) is one of the simple yet efficient methods for solving OT, which can be formulated as (Liero, Mielke, and Savaré 2018):

$$\min_{\mathbf{P} \in U(\mathbf{a}, \mathbf{b})} \langle \mathbf{C}, \mathbf{P} \rangle - \epsilon H(\mathbf{P}), \quad (2)$$

where $\epsilon > 0$ is the coefficient for entropic regularization $H(\mathbf{P})$, and the regularization $H(\mathbf{P})$ can be specified as $H(\mathbf{P}) = -\langle \mathbf{P}, \log \mathbf{P} - \mathbf{1}_{m \times n} \rangle$. The objective in Eq. 2 is ϵ -strongly convex, and thus it has a unique solution, which can be solved by Sinkhorn algorithms as discussed in (Curi 2013; Benamou et al. 2015).

In this paper, beyond vanilla OT, we present our formulation for DB-OT and also propose the Sinkhorn algorithm.

Optimal Transport w/ Inequality Constraints

The vanilla OT only considers the equality constraints in $U(\mathbf{a}, \mathbf{b})$ yet inequality constraints are also handled in a few studies (Caffarelli and McCann 2010; Benamou et al. 2015). For instance, the optimal partial transport (OPT) problem (Caffarelli and McCann 2010) assumes some mass variation or partial mass displacement should be handled for transportation. Thus OPT focuses on transporting only a fraction of mass $s \in [0, \min(\|\mathbf{a}\|_1, \|\mathbf{b}\|_1)]$ as cheaply as possible. Then the constraints in this case is specified as:

$$C_{OPT} = \{\mathbf{P} \in R_{mn}^+ | \mathbf{P}\mathbf{1}_n \leq \mathbf{a}, \mathbf{P}^\top \mathbf{1}_m \leq \mathbf{b}, \mathbf{1}_m^\top \mathbf{P}\mathbf{1}_n = s\}. \quad (3)$$

One can get the Partial-Wasserstein distance as $PW(\mathbf{a}, \mathbf{b}) = \min_{\mathbf{P} \in C_{OPT}} \langle \mathbf{C}, \mathbf{P} \rangle$. This problem has been studied in (Caffarelli and McCann 2010; Figalli 2010; Chapel, Alaya, and Gasso 2020). In particular, (Chizat et al. 2018; Benamou et al. 2015) propose the numerical solutions of OPT with entropic regularization. Different from the above constraints which only give upper bound in the constraints for mass transportation, in this paper, both upper and lower bounds are assumed for transportation, which allows the transportation results to fluctuate within a certain range rather than an exact number.

Unbalanced Optimal Transport

OT in its vanilla form requires the two histograms \mathbf{a} and \mathbf{b} to have the same total mass. Many works have tried to go beyond this assumption and following the setting in Unbalanced OT (Liero, Mielke, and Savaré 2018; Bai et al. 2023), the Kantorovich formulation in Eq. 1 is "relaxed" by only penalizing marginal deviation using some divergence \mathcal{D}_ψ :

$$\min_{\mathbf{P} \in \mathbb{R}_{p, mn}^+} \langle \mathbf{C}, \mathbf{P} \rangle + \tau_1 \mathcal{D}_\psi(\mathbf{P}\mathbf{1}_n | \mathbf{a}) + \tau_2 \mathcal{D}_\psi(\mathbf{P}^\top \mathbf{1}_m | \mathbf{b}) \quad (4)$$

where τ_1 and τ_2 control how much mass variations are penalized as opposed to transportation of the mass and when $\tau_1 = \tau_2 \rightarrow +\infty$. This formulation equals to the original Kantorovich for $\sum_i \mathbf{a}_i = \sum_j \mathbf{b}_j$. When $\tau_1 \rightarrow +\infty$ and $\tau_2 < +\infty$, Unbalanced OT and DB-OT share a similar characteristic where the source distribution remains fixed while the target becomes unfixed. However, unlike DB-OT, it can only passively adjust the target distribution through τ_2 lacking a direct mechanism to control the targets within a range.

Unbalanced Image Recognition

Unbalanced classification, particularly in the case of long-tailed recognition, is a well-known challenge that has garnered significant attention in vision and machine learning (Zhang et al. 2023; He and Garcia 2009; Lin et al. 2017).

Various approaches (Tan et al. 2020; Cui et al. 2019b; Lin et al. 2017; Kang et al. 2019; Zhang et al. 2021) have been proposed to address this issue. One popular strategy involves rebalancing the class distribution in the training data (Ren et al. 2020; Park et al. 2021). Some methods employ techniques such as re-sampling (Kang et al. 2019) or re-weighting (Cui et al. 2019b) to ensure that the model pays more attention to minority classes during training. Data augmentation techniques, such as synthesizing additional samples for underrepresented classes, have also been utilized (Kim, Jeong, and Shin 2020). In this paper, we solve the unbalanced recognition with the OT perspective. We apply DB-OT both for learning the representations and inference in testing process, which is different to previous Bayesian view for classification.

Double-Bounded Optimal Transport

Vanilla OT usually involves two sets of equality constraints. Our motivation is to modify one type of constraints to the double bounded form which can be applicable to advanced clustering and classification.

Formulation of DB-OT

We first assume the source measure $\alpha = \sum_{i=1}^m \mathbf{a}_i \delta_{x_i}$ and two bounds of target measure $\beta^d = \sum_{j=1}^n \mathbf{b}_j^d \delta_{y_j}$, $\beta^u = \sum_{j=1}^n \mathbf{b}_j^u \delta_{y_j}$ where $\mathbf{a} \geq \mathbf{0}$ and $\mathbf{b}^u \geq \mathbf{b}^d \geq \mathbf{0}$. Our purpose is to transport the source samples from α to the target measure which is defined between β^d and β^u . By defining the cost matrix \mathbf{C} between $\{x_i\}$ and $\{y_j\}$ (e.g. $C_{ij} = d(x_i, y_j)$ where $d(\cdot, \cdot)$ is a distance), we can formulate the Double-Bounded Optimal Transport (DB-OT) as

$$\min_{\mathbf{P} \in C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d)} \langle \mathbf{C}, \mathbf{P} \rangle = \sum_{i,j} C_{ij} \mathbf{P}_{ij}, \quad (5)$$

where $C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d)$ is the coupling set defined by constraints:

$$C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d) = \{\mathbf{P} \in \mathbb{R}_{mn}^+ | \mathbf{P}\mathbf{1}_n = \mathbf{a}, \mathbf{b}^d \leq \mathbf{P}^\top \mathbf{1}_m \leq \mathbf{b}^u\}. \quad (6)$$

Similar to vanilla OT, the above model is still a linear program. From Eq. 6, we can see that we replace the original constraint $\mathbf{P}^\top \mathbf{1}_m = \mathbf{b}$ in Eq. 1 with $\mathbf{b}^d \leq \mathbf{P}^\top \mathbf{1}_m \leq \mathbf{b}^u$. That is, we slice $\mathbf{P}^\top \mathbf{1}_m$ between \mathbf{b}^d and \mathbf{b}^u to relax the constraint on $\mathbf{P}^\top \mathbf{1}_m$. It is obvious that when $\mathbf{b} = \mathbf{b}^d = \mathbf{b}^u$, DB-OT degenerates to vanilla OT in Eq. 1. In this paper, we mainly focus on the entropic regularized formulation:

$$\min_{\mathbf{P} \in C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d)} \langle \mathbf{C}, \mathbf{P} \rangle - \epsilon H(\mathbf{P}), \quad (7)$$

where ϵ is the regularization coefficient and $H(\mathbf{P})$ is the entropic regularization. It is obvious that the objective in Eq. 7 and the constraint set $C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d)$ are both convex and thus Eq. 7 has a unique optimal solution. We will discuss its solution and the algorithm of entropic DB-OT in next subsection.

Why not using double-bounded α ? An intuitive question arises: why don't we assume that the α distribution is also constrained within upper and lower bounds? For instance, let's assume $\alpha^u = \sum_{i=1}^m \mathbf{a}_i^u \delta_{x_i}$ and $\alpha^d = \sum_{i=1}^m \mathbf{a}_i^d \delta_{x_i}$, and the coupling satisfies $\mathbf{a}^d \leq \mathbf{P}\mathbf{1}_n \leq \mathbf{a}^u$. In practice, the optimal transportation tends to transport mass vertically and towards the smaller lower bound. Without loss of generality, let's assume $\sum_i \mathbf{a}_i^d < \sum_j \mathbf{b}_j^d$. In this case, the optimal solution must satisfy $\sum_{ij} \mathbf{P}_{ij} = \sum_i \mathbf{a}_i^d$ (proof provided in online Appendix). As a result, setting an upper bound for the source distribution does not hold much significance.

Sinkhorn Algorithm Variants for DB-OT

In this subsection, we introduce several properties of DB-OT and present three corresponding Sinkhorn algorithms, namely the Bregman iterative algorithm, Sinkhorn-Knopp algorithm, and Dual algorithm for DB-OT, respectively, which aims to be **in line with** the algorithms of **vanilla OT**. Although these algorithms are derived from different formulations or properties of DB-OT, they are fundamentally interconnected and converge to same solutions, which validates the effectiveness of the algorithms.

■ Variant-I: Bregman Iterations For DB-OT

Similar to the vanilla entropic OT, our entropic DB-OT in Eq. 7 can also be reformulated to the "static Schrödinger form" (Léonard 2012), which exactly learns a projection under KL divergence.

Proposition 1 (Static Schrödinger Form) *Redefine a general KL divergence in line with (Benamou et al. 2015)* $\widetilde{KL}(\mathbf{P}|\mathbf{K}) = \sum_{ij} \mathbf{P}_{ij} \log \frac{\mathbf{P}_{ij}}{\mathbf{K}_{ij}} - \mathbf{P}_{ij} + \mathbf{K}_{ij}$. Let $\mathbf{K}_{ij} = e^{-C_{ij}/\epsilon}$, the optimization in Eq. 7 is equivalent to the minimization:

$$\mathbf{P}^* = \arg \min_{\mathbf{P} \in C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d)} \widetilde{KL}(\mathbf{P}|\mathbf{K}). \quad (8)$$

Following (Benamou et al. 2015), a new variant of Sinkhorn algorithm can be devised with iterative Bregman projections for DB-OT. By splitting the constraint set as $C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d) = C_1 \cap C_2 \cap C_3$ where $C_1 = \{\mathbf{P} \in \mathbb{R}_{nm}^+ | \mathbf{P}\mathbf{1}_n = \mathbf{a}\}$, $C_2 = \{\mathbf{P} \in \mathbb{R}_{nm}^+ | \mathbf{P}^\top \mathbf{1}_m \geq \mathbf{b}^d\}$, $C_3 = \{\mathbf{P} \in \mathbb{R}_{nm}^+ | \mathbf{P}^\top \mathbf{1}_m \leq \mathbf{b}^u\}$. We can get the KL projection for \mathbf{P} under the constraints as:

$$\begin{aligned} Proj_{C_1}^{KL}(\mathbf{P}) &= \text{diag}\left(\frac{\mathbf{a}}{\mathbf{P}\mathbf{1}_n}\right) \mathbf{P} \\ Proj_{C_2}^{KL}(\mathbf{P}) &= \mathbf{P} \text{diag}\left(\max\left(\frac{\mathbf{b}^d}{\mathbf{P}^\top \mathbf{1}_m}, \mathbf{1}_n\right)\right) \\ Proj_{C_3}^{KL}(\mathbf{P}) &= \mathbf{P} \text{diag}\left(\min\left(\frac{\mathbf{b}^u}{\mathbf{P}^\top \mathbf{1}_m}, \mathbf{1}_n\right)\right), \end{aligned} \quad (9)$$

where the above division operator between two vectors is to be understood entry-wise (See details in online Appendix). Finally, as introduced by (Benamou et al. 2015), assuming $C_l = C_{l+3}$ with positive integer l as the index of Bregman iteration, the minimization in Eq. 8 can be solved by with the iterative projection $\mathbf{P}^{(n)} = Proj_{C_n}^{KL}(\mathbf{P}^{(n-1)})$, starting from $\mathbf{P}^{(0)} = \mathbf{K} = e^{-C/\epsilon}$. Then we can get the solution by $\mathbf{P}^* = \lim_{n \rightarrow \infty} Proj_{C_n}^{KL}(\mathbf{P}^{(n-1)})$, which exactly do the iterations for Eq. 9 and (Bregman 1967) shows its convergence guarantee.

■ Variant-II: Sinkhorn-Knopp algorithm for DB-OT

In addition to analyzing DB-OT with KL divergence and Bergman iterative projections, Lagrangian methods can also solve the problem. Following (Cuturi 2013), we show that the form of solution is specified as follows.

Proposition 2 (Solution Property) *The optimal solution of the optimization in Eq. 7 is unique and has the form*

$$\mathbf{P}^* = \text{diag}(\mathbf{u}) \mathbf{K} \text{diag}(\mathbf{q} \odot \mathbf{v}) \quad (10)$$

where $\mathbf{q} \odot \mathbf{v} = (\mathbf{q}_i \mathbf{v}_i) \in \mathbb{R}_n^+$ and the three scaling variables $(\mathbf{u}, \mathbf{q}, \mathbf{v})$ satisfy $\mathbf{u} \in \mathbb{R}_n^+$, $\mathbf{0}_n \leq \mathbf{v} \leq \mathbf{1}_n$ and $\mathbf{q} \geq \mathbf{1}_n$.

The proof is given in uploaded online appendix. Comparing to vanilla OT with equality constraints, three scaling variables need to be iterated. We then also derive the Sinkhorn_Knopp algorithm of DB-OT given the iteration number l :

$$\begin{aligned} \mathbf{u}^{(l+1)} &= \frac{\mathbf{a}}{\mathbf{K}(\mathbf{v}^{(l)} \odot \mathbf{q}^{(l)})}, \\ \mathbf{q}^{(l+1)} &= \max\left(\frac{\mathbf{b}^d}{(\mathbf{u}^{(l+1)} \mathbf{K}) \odot \mathbf{v}^{(l)}}, \mathbf{1}_n\right), \\ \mathbf{v}^{(l+1)} &= \min\left(\frac{\mathbf{b}^u}{(\mathbf{u}^{(l+1)} \mathbf{K}) \odot \mathbf{q}^{(l+1)}}, \mathbf{1}_n\right). \end{aligned} \quad (11)$$

By initializing the scaling variables $\mathbf{0}_n \leq \mathbf{v}^{(0)} \leq \mathbf{1}_n$ and $\mathbf{q}^{(0)} \geq \mathbf{1}_n$, we can iterate by Eq. 11 to compute $(\mathbf{u}^{(1)}, \mathbf{q}^{(1)}, \mathbf{v}^{(1)}), (\mathbf{u}^{(2)}, \mathbf{q}^{(2)}, \mathbf{v}^{(2)}), \dots$ until convergence to

$(\mathbf{u}^*, \mathbf{q}^*, \mathbf{v}^*)$. Then the solution can be obtained as $\mathbf{P}^* = \text{diag}(\mathbf{u}^*) \mathbf{K} \text{diag}(\mathbf{q}^* \odot \mathbf{v}^*)$.

■ Variant-III: The Dual Sinkhorn algorithm

We also analyze entropic DB-OT with dual formulation, which is devised as follows.

Proposition 3 (Dual Formulation) *From the optimization in Eq. 7, we can get its dual formulation by maximizing*

$$L = \langle \mathbf{f}, \mathbf{a} \rangle + \langle \mathbf{g}, \mathbf{b}^d \rangle + \langle \mathbf{h}, \mathbf{b}^u \rangle - \epsilon B(\mathbf{f}, \mathbf{g}, \mathbf{h}) \quad (12)$$

where $\mathbf{f} \in \mathbb{R}^n$, $\mathbf{g} \leq \mathbf{0}$ and $\mathbf{h} \geq \mathbf{0}$ are the corresponding dual variables and $B(\mathbf{f}, \mathbf{g}, \mathbf{h}) = \langle e^{\mathbf{f}/\epsilon}, \mathbf{K} e^{(\mathbf{g}+\mathbf{h})/\epsilon} \rangle$.

Then we can get the new algorithm with a block coordinate ascent on the dual:

$$\begin{aligned} \Delta_{\mathbf{f}} L &= \mathbf{a} - e^{\mathbf{f}/\epsilon} \odot \mathbf{K} e^{(\mathbf{g}+\mathbf{h})/\epsilon} \\ \Delta_{\mathbf{g}} L &= \mathbf{b}^d - e^{(\mathbf{g}+\mathbf{h})/\epsilon} \odot \mathbf{K}^\top e^{\mathbf{f}/\epsilon}. \\ \Delta_{\mathbf{h}} L &= \mathbf{b}^u - e^{(\mathbf{g}+\mathbf{h})/\epsilon} \odot \mathbf{K}^\top e^{\mathbf{f}/\epsilon} \end{aligned} \quad (13)$$

Starting from arbitrary $\mathbf{g}^{(0)} \leq \mathbf{0}$ and $\mathbf{h}^{(0)} \geq \mathbf{0}$, we get the dual algorithm by iterating the following with iteration number l :

$$\begin{aligned} \mathbf{f}^{(l+1)} &= \epsilon \log \mathbf{a} - \epsilon \log \left(\mathbf{K} e^{(\mathbf{g}^{(l)} + \mathbf{h}^{(l)})/\epsilon} \right) \\ \mathbf{g}^{(l+1)} &= \max \left\{ \epsilon \log \mathbf{b}^d - \epsilon \log(\mathbf{K}^\top e^{\mathbf{f}^{(l+1)}/\epsilon}) - \mathbf{h}^{(l)}, \mathbf{0} \right\} \\ \mathbf{h}^{(l+1)} &= \min \left\{ \epsilon \log \mathbf{b}^u - \epsilon \log(\mathbf{K}^\top e^{\mathbf{f}^{(l+1)}/\epsilon}) - \mathbf{g}^{(l+1)}, \mathbf{0} \right\} \end{aligned} \quad (14)$$

The iteration continues till convergence to $(\mathbf{f}^*, \mathbf{g}^*, \mathbf{h}^*)$ and the optimal solution equals $\mathbf{P}^* = \text{diag}(e^{\mathbf{f}^*/\epsilon}) \mathbf{K} \text{diag}(e^{(\mathbf{g}^* + \mathbf{h}^*)/\epsilon})$.

Exactly, Eq. 14 are mathematically equivalent to the Sinkhorn iterations in Eq. 11 when considering the primal-dual relations. Indeed, we recover that at any iteration: $(\mathbf{f}^{(l)}, \mathbf{g}^{(l)}, \mathbf{h}^{(l)}) = \epsilon (\log(\mathbf{u}^{(l)}), \log(\mathbf{q}^{(l)}), \log(\mathbf{v}^{(l)}))$. In this paper, we utilize the Sinkhorn_Knopp algorithm for clustering tasks and the Bregman iterations for long-tail classification.

DB-OT for Clustering and Classification

Barycenter-based Clustering

Inspired by the OT-based topic model (Huynh, Zhao, and Phung 2020), we show the application of DB-OT to an advanced clustering task whereby the number of samples in each cluster is softly required to fall into a certain range. Specifically given samples $\{\alpha_s\}$, it learns the coupling \mathbf{P} to match the samples to centroids, and the clustered centroids $\{\beta_t\}$ are viewed as barycenters which needs to optimize. The optimization is specified as:

$$\min_{\mathbf{P} \in C(\mathbf{a}, \mathbf{b}^d, \mathbf{b}^u), \{\beta_t\}} \sum_{st} \mathbf{P}_{st} D(\alpha_s, \beta_t) - \epsilon H(\mathbf{P}), \quad (15)$$

where $\mathbf{a} = \mathbf{1}$ (here we break the probability measure assumption as done in (Saad-Eldin et al. 2021)) and $\mathbf{b}^d, \mathbf{b}^u$ are used to denote the given bounds of clustered centroids i.e. the minimum/maximum "quantity" of samples in each cluster. $D(\alpha_s, \beta_t)$ is the distance between α_s and β_t (e.g. Euclidean or Wasserstein distance), and $H(\mathbf{P})$ is the entropic regularization which aims to relax the solution of coupling. This problem can be solved by alternating optimization. For each iteration, we first fix the centroid distributions $\{\beta_t\}$ to

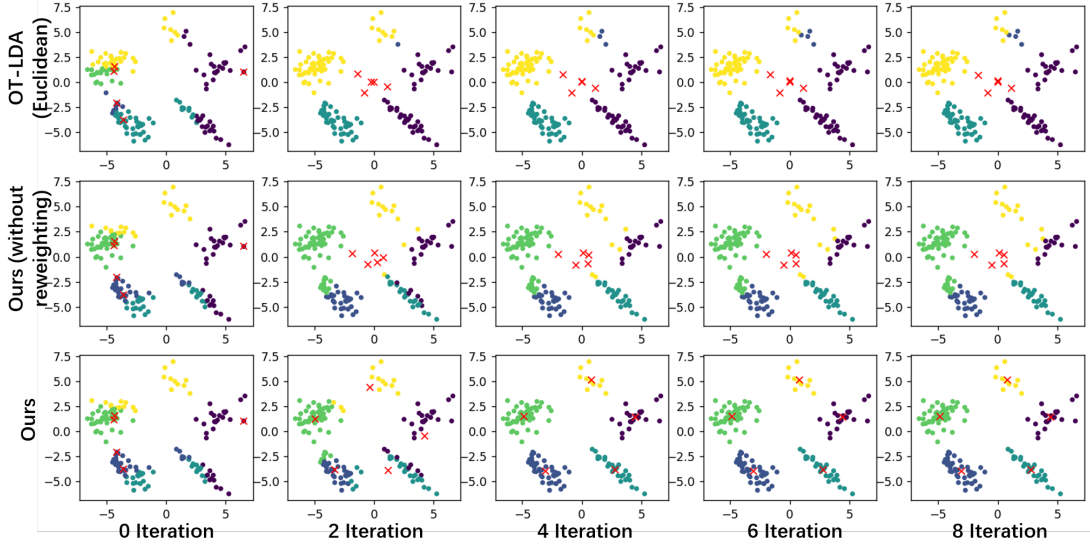


Figure 2: The results of Barycenter-based clustering, which is performed on data points sampled from 5 Gaussian distributions. The colors represent the cluster assignments of the samples, and the red crosses denote the centroids/barycenters. Note that in both OT-LDA and our method without reweighting the barycenter weights, the calculated centroids exhibit a noticeable bias.

compute the distance $D(\alpha_s, \beta_t)$ between sample i and centroid j and then transportation matrix \mathbf{P} can be learned from the source samples to clustered centroids.

Note in our method, we have $\mathbf{b}^d \leq \mathbf{P}^\top \mathbf{1}_n \leq \mathbf{b}^u$, which means the coupling between samples and centroids are controlled within a desired range, avoiding clusters with isolated or very few samples, as well as dominating clusters.

Specifically, for solving the optimization to clustering in Eq. 15, we first initialize $\{\beta_t\}$ with initialization of Kmeans++ or set them as randomly from $\{\alpha_s\}$. Then we can learn by iterating the following two steps:

1) Fixing $\{\beta_t\}$ to compute \mathbf{P} . If the barycenter $\{\beta_t\}$ is known, then one can calculate the matrix \mathbf{D} where $\mathbf{D}_{st} = D(\alpha_s, \beta_t)$. Viewing the matrix \mathbf{D} as the cost function, the optimization of \mathbf{P} in Eq. 15 is equal to DB-OT optimization in Eq. 7 and the algorithm proposed in Sec. can be used to get the result of \mathbf{P} . Here we adopt the Sinkhorn_Knopp algorithm as given in Eq. 11.

2) Fixing \mathbf{P} to compute $\{\beta_t\}$. Assuming the transportation \mathbf{P} is known, our goal is now to update $\{\beta_t\}$ as the solution of the following optimization problem

$$\beta_t = \arg \min_{\beta} \sum_s \frac{\mathbf{P}_{st} \cdot R_{st}}{\sum_{s'} \mathbf{P}_{s't} \cdot R_{s't}} D(\alpha_s, \beta) \quad (16)$$

where $R_{st} = 1$ if $t = \arg \max_{t'} \mathbf{P}_{st'}$ otherwise $R_{st} = 0$ given the sample α_s . Here R_{st} are used for re-weighting the weights of barycenters, which aims to debias the influence for samples from different clusters. Fig. 2 is the clustering results and we can find a large bias for centroids without re-weighting the barycenter weights.

Under Different Metric Space for $D(\cdot, \cdot)$. Clustering can be performed in different metric spaces by setting the distance function $D(\cdot, \cdot)$ to adapt to various metric settings or datasets. For example, when $D(\cdot, \cdot)$ represents the Euclidean distance

and α_s corresponds to feature points in Euclidean space, we aim to cluster α_s in the Euclidean space. The centroids β_t can be directly computed as follows:

$$\beta_t = \sum_s \frac{\mathbf{P}_{st} \cdot R_{st}}{\sum_{s'} \mathbf{P}_{s't} \cdot R_{s't}} \alpha_s. \quad (17)$$

Exactly if we set $R_{st} = 1$ for all s and t (i.e. without reweighting as shown in the second row of Fig. 2), the barycenters can be easily influenced by samples from other clusters. This can lead to a bias in the barycenter calculation, particularly when the value of ϵ is not sufficiently small.

In addition to Euclidean space, clustering can also be performed in Wasserstein space. In this case, we assume that α_s is no longer a set of feature points but a collection of probability measures, i.e., $\alpha_s = \sum_i \mathbf{a}_i^s \delta_{x_i}$. In this setting, the computation of the barycenter in Eq. 16 does not have an analytical solution but requires iterative methods, as proposed in (Benamou et al. 2015) for barycenter calculation.

Advantages for our Barycenter-based Clustering. Our barycenter-based clustering is essentially a unified clustering framework. Compare to the previous clustering works (e.g. Kmeans (Ahmed, Seraj, and Islam 2020) or OT-LDA (Huynh, Zhao, and Phung 2020)), the advantages can be summarized as follows: **1) Our** clustering methods apply DB-OT to calculate the matching between samples and centroids, providing better controllability. As shown in Fig. 3, by setting \mathbf{b}^d and \mathbf{b}^u , we can constrain the number of samples in each cluster to a certain extent, thereby avoiding clusters with isolated or very few samples, as well as clusters that dominate the majority of the data; **2) Our** clustering methods can be applied not only to clustering feature points but also to clustering probability measures. For example, as shown in Fig. 3, we can perform clustering on the MNIST dataset in the Wasserstein space. In this case, we no

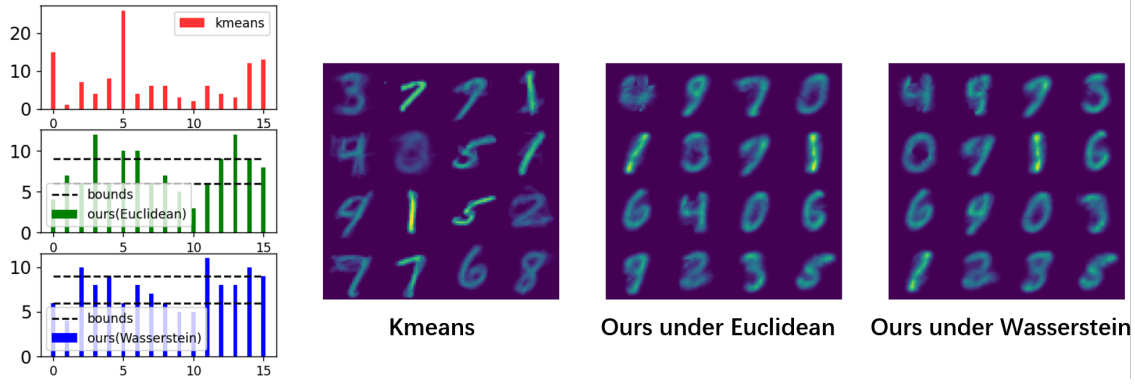


Figure 3: Clustering distribution and the pixel-wise mean centroids (forming into numbers) on MNIST. Our results are well controlled within the bounds, and kmeans cannot satisfy this property resulting in more scattered clusters of varying size.

longer consider images as point vectors but treat the pixel values as histograms, and the pixel locations are used to compute the cost matrix. Moreover, this clustering method can also be used for text clustering (e.g., topic modeling), making it more adaptable to different types of datasets; **3) Compared** to previous work in OT-LDA (Huynh, Zhao, and Phung 2020), we introduce a reweighting of the barycenter weights, which helps to mitigate the biases in centroid computation, as shown in Fig. 2.

DB-OT for Long-tailed Classification

Motivation of using DB-OT in classification. Here we apply DB-OT for Long-tailed unbalanced classification. Following the works (Shi et al. 2023a,b), the Inverse Optimal Transport (Li et al. 2019; Stuart and Wolfram 2020; Chiu, Wang, and Shafto 2022) can be set as a bi-level optimization for classification as

$$\min_{\theta} KL(\tilde{\mathbf{P}}|\mathbf{P}^{\theta}) \quad \text{s.t.} \quad \mathbf{P}^{\theta} = \arg \min_{\mathbf{P} \in U} \langle \mathbf{C}^{\theta}, \mathbf{P} \rangle - \epsilon H(\mathbf{P}). \quad (18)$$

Here U is the set of couplings and when the set $U = U(\mathbf{a}) = \{\mathbf{P} \in \mathbb{R}_{m \times n}^+ | \mathbf{P}\mathbf{1}_m = \mathbf{a}\}$, this optimization is equivalent to:

$$\min_{\theta} \mathcal{L} = - \sum_{ij} \tilde{\mathbf{P}}_{ij} \log \left(\frac{\exp(-\mathbf{C}_{ij}^{\theta}/\epsilon)}{\sum_{k=1}^m \exp(-\mathbf{C}_{ik}^{\theta}/\epsilon)} \right) + \text{constant}, \quad (19)$$

where $\tilde{\mathbf{P}}_{ij}$ is the ground truth matrix defined by the supervised one-hot labels and $\mathbf{a} = \mathbf{1}/m$. The proof process is similarly discussed in (Shi et al. 2023a). Specifically, for the ground truth, $\tilde{\mathbf{P}}_{ij} = 1$ if j is the label index of sample i otherwise $\tilde{\mathbf{P}}_{ij} = 0$. The loss in Eq. 19 is exactly equal to the Softmax-CrossEntropy loss if we set $\mathbf{C}_{ij}^{\theta} = c - l_{ij}$ where c is a large enough constant and l_{ij} is the logits of sample i on class j . The equivalence between Inverse OT under $U(\mathbf{a})$ and Softmax-CrossEntropy loss motivate us that we can adopt different constraints (e.g. $C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d)$) instead of $U(\mathbf{a})$. Then in the following, we show our main idea that classification learning can be viewed as optimizing Inverse OT and Classification inference is learning OT.

Training via Inverse DB-OT. As discussed above, we apply DB-OT by setting $U = C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d)$ in Eq. 18 for long-tailed

classification, which assumes that the labels of training data have a known long-tailed distribution \mathbf{r} . We set $\mathbf{a} = \mathbf{1}/m$ and:

$$\mathbf{b}^u = (1 + \delta)\mathbf{r} \quad \text{and} \quad \mathbf{b}^d = (1 - \delta)\mathbf{r}, \quad (20)$$

where $\delta \in (0, 1)$ is the bound rate for \mathbf{r} . We adopt the Bregman method for calculating \mathbf{P}^{θ} with K iteration. We find when $K = 1$ and $\delta = 0$, the Balanced Softmax (Ren et al. 2020) is a special case of our loss, which validates the effectiveness of our theory.

Testing-time Inference with DB-OT. Here we focus on using DB-OT for inference in testing process. Note our inference method is orthogonal to the training method proposed in Sec. , which means our inference can be used for different classifiers. For the testing inference, we treat all the training as a process of feature learning and then with the feature learned, we can define the cost matrix given each batch data and match the features and labels with OT. For the long-tailed classification, the inference is as follows:

$$\min_{\mathbf{P} \in C(\mathbf{a}, \mathbf{b}^u, \mathbf{b}^d)} \langle \mathbf{C}, \mathbf{P} \rangle - \epsilon H(\mathbf{P}). \quad (21)$$

Here $\mathbf{C}_{ij} = c - l_{ij}$ with large enough value c to guarantee the positiveness of the the cost. Meanwhile, we set \mathbf{b}^u and \mathbf{b}^d as Eq. 20, which enables predictions to fluctuate both upward and downward, taking into account the randomness of batch data sampling. During the testing process, \mathbf{r} can be configured as a known long-tailed distribution, uniform distribution, reverse long-tailed distribution, or any other appropriate distribution, which is a main advantage of our method.

Experiments

We conduct experiments on both clustering and classification. To showcase the advantage of DB-OT on fine-grained controlling the behavior of clustering under bounded sizes of clusters, and classification in long-tail cases.

Experiments on Size-controlled Clustering

We first study our clustering setting which each cluster's size is bounded by a certain range, using our DB-OT function with Barycenter. The experimental results on Gaussian mixture synthetic datasets with 2D 150 points are shown in

| Method | CIFAR10-LT | | | CIFAR100-LT | | | | ImageNet-LT | | | |
|---------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | Many | Few | All | Many | Medium | Few | All | Many | Medium | Few | All |
| Vanilla Softmax | 77.4 | 68.9 | 74.9 | 75.8 | 48.2 | 11.0 | 42.0 | 57.3 | 26.2 | 3.1 | 35.0 |
| LDAM (Cao et al. 2019) | 80.5 | 65.2 | 75.9 | 75.7 | 50.6 | 11.5 | 42.9 | 57.3 | 27.6 | 4.4 | 35.9 |
| Balanced Softmax | 82.2 | 71.6 | 79.0 | 70.3 | 50.4 | 26.5 | 47.0 | 52.5 | 38.6 | 17.8 | 41.1 |
| Focal Loss (Lin et al. 2017) | 79.6 | 58.4 | 73.3 | 76.1 | 46.9 | 11.1 | 41.7 | 57.3 | 27.6 | 4.4 | 35.9 |
| LogitAdjust (Menon et al. 2020) | 80.0 | 35.3 | 66.6 | 75.7 | 39.2 | 4.1 | 36.5 | 54.2 | 14.0 | 0.4 | 27.6 |
| CB-CE (Cui et al. 2019b) | 76.6 | 70.7 | 74.8 | 53.2 | 48.8 | 13.3 | 36.3 | 35.3 | 32.1 | 21.2 | 31.9 |
| CB-FC (Cui et al. 2019b) | 76.6 | 70.7 | 74.8 | 53.2 | 48.8 | 13.3 | 36.3 | 35.3 | 32.1 | 21.2 | 31.9 |
| DB-OT Loss (ours) | 82.4 | 80.8 | 81.9 | 70.4 | 53.0 | 26.6 | 47.9 | 53.5 | 39.0 | 17.4 | 41.6 |

Table 1: Top-1 accuracy (%) for long-tailed image classification with 200 imbalanced factor on three popular LT datasets.

| Testing Inference | Vanilla Softmax Loss | | | Balanced Softmax | | | Focal Loss | | | Logit Adjustment Loss | | |
|------------------------|----------------------|-------------|-------------|------------------|-------------|-------------|-------------|-------------|-------------|-----------------------|-------------|-------------|
| | LT | Uniform | RLT | LT | Uniform | RLT | LT | Uniform | RLT | LT | Uniform | RLT |
| Vanilla Softmax | 71.6 | 42.0 | 18.7 | 66.5 | 47.0 | 20.1 | 68.5 | 41.7 | 18.4 | 69.3 | 36.4 | 4.8 |
| classifier normalize | 70.0 | 45.2 | 16.8 | 54.9 | 42.4 | 28.6 | 67.0 | 41.3 | 13.3 | 70.5 | 40.4 | 7.9 |
| Class-aware bias | 68.0 | 41.3 | 14.5 | 41.5 | 45.0 | 18.6 | 54.6 | 40.0 | 14.1 | 70.8 | 36.4 | 7.7 |
| DB-OT inference (ours) | 71.8 | 48.5 | 36.6 | 71.4 | 47.8 | 34.6 | 68.7 | 44.1 | 34.4 | 71.3 | 44.8 | 32.3 |

Table 2: Top-1 accuracy (%) of CIFAR-100 for the comparison of different testing inference methods given four trained models.

Fig. 2. We can find that the OT-LDA and our method without reweighting exhibit a noticeable bias and re-weighting the barycenter weights can overcome this issue with Eq. 16. We also evaluate the top-1 accuracy within the cluster and the results of OT-LDA, ours (without and with reweighting) are 82.67%, 88.00%, and 100.00% respectively, which shows the superiority of our method. We also do the clustering experiments on MNIST with three different methods (i.e. K-means, DB-OT under Euclidean space and DB-OT under Wasserstein space). We select parts of MNIST data (12 images in each class) and set 16 clusters for the experiments. The results are shown in Fig. 3, and we can find that our method can be more controllable for the number of samples in each cluster. Note the bounds are to control the column probability sum of the coupling and thus it does not strictly meet the bound range for sample quantity.

Experiments on Long-tailed Classification

We evaluate our two methods – both for using DB-OT to improve the loss and adopting the DB-OT for classification inference in the testing process. These two methods are parallel to each other as the former aims to learn a better representation and the latter is to get a more accurate prediction based on a known testing label ratio. We do the experiments on CIFAR10-LT, CIFAR100-LT (Krizhevsky, Hinton et al. 2009), ImageNet-LT (Liu et al. 2019) for image classification. For a fair comparison, all methods share the same network backbone and hyperparameters. More details about the experimental settings is given in the online Appendix.

Improvements of DB-OT based loss. To evaluate the performance of our DB-OT based loss, following (Ren et al. 2020), we start with a vanilla Softmax pretrained model and train with our DB-OT based loss to improve the representations. We use the corresponding balanced testing dataset for evaluation, where its labels are uniformly distributed. We report top-1 accuracy as the evaluation metric. Specifically,

for CIFAR10-LT, we report accuracy on two sets of classes in detail: Many-shot (more than 100 images) and Few-shot (less than 100 images). For CIFAR00-LT and ImageNet-LT, we report accuracy on three sets: Many-shot (more than 100 images), Medium-shot (20 ~ 100 images), and Few-shot (less than 20 images). The experiments for unbalanced image classification are all conducted with an imbalanced factor of less than 200, which is defined as the ratio of the number of training instances in the largest class to the smallest (Ren et al. 2020). The results for long-tailed classification are presented in Table 1. From a comprehensive perspective, our approach achieves the best average accuracy across the entire dataset. When examining the results for different data splits, our method particularly outperforms the others on subsets with fewer images (i.e., median or few-shot).

Performance of testing inference. To evaluate the DB-OT based inference, we use the corresponding long-tailed (LT), uniform, and Reverse long-tailed (RLT) testing dataset for evaluation. We compare our method with vanilla Softmax, classifier normalization (Kang et al. 2019) and class-aware bias (Menon et al. 2020), which are discussed in (Wu et al. 2021). Table 2 shows the comparison results for all the models and different testing data in CIFAR100. Our inference method outperforms and can achieve a great improvement when the testing data is reverse long-tailed distributed and the model trained by vanilla softmax performs the best using our DB-OT testing inference though it may fail with vanilla softmax prediction in testing. More experimental results (e.g. ablation study) are given in online appendix

Conclusion

We have presented the so-called double-bounded optimal transport, with theoretical analysis and further derive three variants of algorithms to solve the problem. We test our technique for the challenging yet realistic tasks of cluster-size bounded clustering, as well as long-tailed image recognition.

Acknowledgements

Junchi Yan is the corresponding author, and this work is supported by the National Natural Science Foundation of China (grant numbers 92370201 and 62222607).

References

- Agueh, M.; and Carlier, G. 2011. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2): 904–924.
- Ahmed, M.; Seraj, R.; and Islam, S. M. S. 2020. The k-means algorithm: A comprehensive survey and performance evaluation. *Electronics*, 9(8): 1295.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein GAN. *ICML*.
- Bai, Y.; Schmitzer, B.; Thorpe, M.; and Kolouri, S. 2023. Sliced optimal partial transport. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13681–13690.
- Benamou, J.-D.; Carlier, G.; Cuturi, M.; Nenna, L.; and Peyré, G. 2015. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2): A1111–A1138.
- Bregman, L. M. 1967. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3): 200–217.
- Caffarelli, L. A.; and McCann, R. J. 2010. Free boundaries in optimal transport and Monge-Ampère obstacle problems. *Annals of mathematics*, 673–730.
- Cao, K.; Wei, C.; Gaidon, A.; Arechiga, N.; and Ma, T. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in neural information processing systems*, 32.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33: 9912–9924.
- Chapel, L.; Alaya, M. Z.; and Gasso, G. 2020. Partial optimal transport with applications on positive-unlabeled learning. *Advances in Neural Information Processing Systems*, 33: 2903–2913.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Chiu, W.-T.; Wang, P.; and Shafto, P. 2022. Discrete Probabilistic Inverse Optimal Transport. In *International Conference on Machine Learning*, 3925–3946. PMLR.
- Chizat, L.; Peyré, G.; Schmitzer, B.; and Vialard, F.-X. 2018. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of computation*, 87(314): 2563–2609.
- Cui, L.; Qi, X.; Wen, C.; Lei, N.; Li, X.; Zhang, M.; and Gu, X. 2019a. Spherical optimal transportation. *Computer-Aided Design*, 115: 181–193.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019b. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9268–9277.
- Cuturi, M. 2013. Sinkhorn distances: Lightspeed computation of optimal transportation distances. *arXiv preprint arXiv:1306.0895*.
- Figalli, A. 2010. The optimal partial transport problem. *Archive for rational mechanics and analysis*, 195(2): 533–560.
- Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. 2017. Improved Training of Wasserstein GANs. In *NIPS*.
- He, H.; and Garcia, E. A. 2009. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9): 1263–1284.
- Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; and Leskovec, J. 2020. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33: 22118–22133.
- Huynh, V.; Zhao, H.; and Phung, D. 2020. Otda: A geometry-aware optimal transport approach for topic modeling. *Advances in Neural Information Processing Systems*, 33: 18573–18582.
- Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2019. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*.
- Kantorovich, L. 1942. On the transfer of masses (in Russian). In *Doklady Akademii Nauk*, volume 37, 227–229.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.
- Kim, J.; Jeong, J.; and Shin, J. 2020. M2m: Imbalanced classification via major-to-minor translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13896–13905.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning Multiple Layers of Features from Tiny Images. Technical report, Citeseer.
- Léonard, C. 2012. From the Schrödinger problem to the Monge–Kantorovich problem. *Journal of Functional Analysis*, 262(4): 1879–1920.
- Li, R.; Ye, X.; Zhou, H.; and Zha, H. 2019. Learning to match via inverse optimal transport. *Journal of machine learning research*, 20.
- Li, Y.; Mo, Y.; Shi, L.; and Yan, J. 2022. Improving generative adversarial networks via adversarial learning in latent space. *Advances in Neural Information Processing Systems*, 35: 8868–8881.
- Liero, M.; Mielke, A.; and Savaré, G. 2018. Optimal entropy-transport problems and a new Hellinger–Kantorovich distance between positive measures. *Inventiones mathematicae*, 211(3): 969–1117.

- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Liu, Z.; Miao, Z.; Zhan, X.; Wang, J.; Gong, B.; and Yu, S. X. 2019. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2537–2546.
- Menon, A. K.; Jayasumana, S.; Rawat, A. S.; Jain, H.; Veit, A.; and Kumar, S. 2020. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314*.
- Park, S.; Lim, J.; Jeon, Y.; and Choi, J. Y. 2021. Influence-balanced loss for imbalanced visual classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 735–744.
- Ren, J.; Yu, C.; Ma, X.; Zhao, H.; Yi, S.; et al. 2020. Balanced meta-softmax for long-tailed visual recognition. *Advances in neural information processing systems*, 33: 4175–4186.
- Saad-Eldin, A.; Pedigo, B. D.; Priebe, C. E.; and Vogelstein, J. T. 2021. Graph matching via optimal transport. *arXiv preprint arXiv:2111.05366*.
- Shi, L.; Zhang, G.; Zhen, H.; Fan, J.; and Yan, J. 2023a. Understanding and Generalizing Contrastive Learning from the Inverse Optimal Transport Perspective. in *ICML*.
- Shi, L.; Zhen, H.; Zhang, G.; and Yan, J. 2023b. Relative Entropic Optimal Transport: a (Prior-aware) Matching Perspective to (Unbalanced) Classification. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Stuart, A. M.; and Wolfram, M.-T. 2020. Inverse optimal transport. *SIAM Journal on Applied Mathematics*, 80(1): 599–619.
- Tan, J.; Wang, C.; Li, B.; Li, Q.; Ouyang, W.; Yin, C.; and Yan, J. 2020. Equalization loss for long-tailed object recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11662–11671.
- Wang, F.; and Liu, H. 2021. Understanding the behaviour of contrastive loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2495–2504.
- Wang, J.; and Su, X. 2011. An improved K-Means clustering algorithm. In *2011 IEEE 3rd international conference on communication software and networks*, 44–46. IEEE.
- Wang, R.; Yan, J.; and Yang, X. 2019. Learning Combinatorial Embedding Networks for Deep Graph Matching. 3056–3065.
- Wang, R.; Yan, J.; and Yang, X. 2021. Neural graph matching network: Learning lawler’s quadratic assignment problem with extension to hypergraph and multiple-graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wang, W.; Slepčev, D.; Basu, S.; Ozolek, J. A.; and Rohde, G. K. 2013. A linear optimal transportation framework for quantifying and visualizing variations in sets of images. *International journal of computer vision*, 101(2): 254–269.
- Wilson, A. G. 1969. The use of entropy maximising models, in the theory of trip distribution, mode split and route split. *Journal of transport economics and policy*, 108–126.
- Wu, T.; Liu, Z.; Huang, Q.; Wang, Y.; and Lin, D. 2021. Adversarial robustness under long-tailed distribution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8659–8668.
- Xu, H.; Luo, D.; Zha, H.; and Duke, L. C. 2019. Gromov-wasserstein learning for graph matching and node embedding. In *International conference on machine learning*, 6932–6941. PMLR.
- Zhang, Y.; Kang, B.; Hooi, B.; Yan, S.; and Feng, J. 2021. Deep long-tailed learning: A survey. *arXiv preprint arXiv:2110.04596*.
- Zhang, Y.; Kang, B.; Hooi, B.; Yan, S.; and Feng, J. 2023. Deep long-tailed learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.