

IMPROVING MINI-BATCH OPTIMAL TRANSPORT VIA PARTIAL TRANSPORTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Mini-batch optimal transport (m-OT) has been widely used recently to deal with the *memory issue* of OT in large-scale applications. Despite their practicality, m-OT suffers from *misspecified mappings*, namely, mappings that are optimal on the mini-batch level but are partially wrong in the comparison with the optimal transportation plan between the original measures. To address the misspecified mappings issue, we propose a novel mini-batch method by using partial optimal transport (POT) between mini-batch empirical measures, which we refer to as *mini-batch partial optimal transport* (m-POT). Leveraging the insight from the partial transportation, we explain the source of misspecified mappings from the m-OT and motivate why limiting the amount of transported masses among mini-batches via POT can alleviate the incorrect mappings. Finally, we carry out extensive experiments on various applications to compare m-POT with m-OT and recently proposed mini-batch method, mini-batch unbalanced optimal transport (m-UOT). We observe that m-POT is better than m-OT in deep domain adaptation applications while having comparable performance with m-UOT. On other applications, such as deep generative model and color transfer, m-POT yields more favorable performance than m-OT while m-UOT is non-trivial to apply.

1 INTRODUCTION

From its origin in mathematics and economics, optimal transport (OT) has recently become a useful and popular tool in machine learning applications, such as (deep) generative models (Arjovsky et al., 2017; Tolstikhin et al., 2018), computer vision/ graphics (Solomon et al., 2016; Nguyen et al., 2021d), clustering problem (Ho et al., 2017; 2019), and domain adaptation (Courty et al., 2016; Damodaran et al., 2018; Le et al., 2021). An important impetus for such popularity is the recent advance in the computation of OT. In particular, the works of (Altschuler et al., 2017; Dvurechensky et al., 2018; Lin et al., 2019) demonstrate that we can approximate optimal transport with a computational complexity of the order $\mathcal{O}(n^2/\epsilon^2)$ where n is the maximum number of supports of probability measures and $\epsilon > 0$ is the desired tolerance. It is a major improvement over the standard computational complexity $\mathcal{O}(n^3 \log n)$ of computing OT via interior point methods (Pele & Werman, 2009). Another approach named sliced OT in (Bonneel et al., 2015; Nguyen et al., 2021a;c; Deshpande et al., 2019) is able to reduce greatly the computational cost to the order of $\mathcal{O}(n \log n)$ by exploiting the closed-form of one dimensional OT on the projected data. However, due to the projection step, sliced optimal transport is not able to keep all the main differences between high dimensional measures, which leads to a decline in practical performance.

Despite these major improvements on computation, it is still impractical to use optimal transport and its variants in large-scale applications when n can be as large as a few million, such as deep domain adaptation and deep generative model. The bottleneck mainly comes from the *memory issue*, namely, it is impossible for us to store $n \times n$ cost matrix when n is extremely large. To deal with this issue, practitioners often replace the original large-scale computation of OT with cheaper computation on subsets of the whole dataset, which widely refer to as mini-batch approaches (Arjovsky et al., 2017; Genevay et al., 2018; Sommerfeld et al., 2019). In details, instead of computing the original large-scale optimal transport problem, the mini-batch approach splits it into smaller transport problems (sub-problems). Each sub-problem is to solve optimal transport between subsets of supports (or mini-batches), which belong to the original measures. After solving all the sub-problems, the final transportation cost (plan) is obtained by aggregating evaluated mini-batch transportation costs (plans)

from these sub-problems. Due to the small number of samples in each mini-batch, computing OT between mini-batch empirical measures is doable when memory constraints and computational constraints exist. This mini-batch approach is formulated rigorously under the name of mini-batch OT (m-OT) and some of its statistical properties are investigated in (Fatras et al., 2020; 2021b).

Despite its computational practicality, the trade-off of m-OT approach is the *misspecified mappings* in comparison with the full OT transportation plan. In particular, a mini-batch is a sparse representation of the original supports; hence, solving an optimal transport problem between empirical mini-batch measures tends to create transport mappings that do not match the global optimal transport mapping between the original probability measures. The existence of misspecified mappings had been noticed in (Fatras et al., 2021a), however, it is not discussed thoroughly. (Fatras et al., 2021a) proposed to use unbalanced optimal transport (UOT) (Chizat et al., 2018b) as a replacement of OT for the transportation between mini-batches, which they referred to as mini-batch UOT (m-UOT), and showed that m-UOT can reduce the effect of misspecified mappings from the m-OT. They further provided an objective loss based on m-UOT, which achieves considerably better results on deep domain adaptation. However, it is intuitively hard to understand the transportation plan from m-UOT. Also, we observe that m-UOT suffers from an issue that its hyperparameter (τ) is not robust to the scale of the cost matrix, namely, the value of m-UOT’s hyperparameter needs to be chosen to be large if the distances between supports are large, and vice versa. Hence, in applications such as deep generative models where the supports of measures change significantly, we should not fix the regularized parameter of m-UOT.

Contributions. In this work, we develop a novel mini-batch framework to alleviate the misspecified matchings issue of m-OT, the aforementioned issue of m-UOT, and to be robust to the scale of the cost matrix. In short, our contributions can be summarized as follows:

1. We propose to use partial optimal transport (POT) between empirical measures formed by mini-batches that can reduce the effect of misspecified mappings and can be more convenient to use than m-UOT. The new mini-batch framework, named *mini-batch partial optimal transport* (m-POT), has two practical applications: (i) the first application is an efficient mini-batch transportation cost used for objective functions in deep learning problems; (ii) the second application is a meaningful mini-batch transportation plan used for barycentric mappings in color transfer problem. Finally, via some simple examples, we further argue why partial optimal transport (POT) can be a natural solution for the misspecified mappings.

2. We conduct extensive experiments on applications that benefit from using mini-batches, including deep domain adaption, deep generative model, gradient flow, and color transfer to compare m-POT with m-OT and its variant m-UOT (Fatras et al., 2021a). From experimental results, we observe that m-POT is better than m-OT on deep domain adaptation. In particular, m-POT gives at least 4.9 **higher** on the average accuracy than m-OT. Furthermore, it is comparable to m-UOT on that application, namely, m-POT yields to similar classification accuracy to m-UOT on digits datasets, Office-Home datasets, VisDA datasets while mini-batch entropic POT pushes the accuracy a little higher in some tasks. On the deep generative model, m-POT produces lower FID scores than m-OT while m-UOT is unstable and difficult to apply. Finally, experiments on gradient flow and color transfer also demonstrate the favorable performance of m-POT compared to m-OT and m-UOT.

As a result, the practitioners may consider using POT as a replacement for OT in practical problems that uses an OT loss with mini-batches.

Organization. The paper is organized as follows. In Section 2, we provide backgrounds on (unbalanced) optimal transport and their mini-batch versions, then we highlight the limitations of each mini-batch method. In Section 3, we propose mini-batch partial optimal transport to alleviate the limitations of previous mini-batch methods. We provide extensive experiment results of mini-batch POT in Section 4 and conclude the paper with a few discussions in Section 5. Extra experimental, theoretical results and settings are deferred to the Appendices.

Notation: For α and β are two discrete probability measures $\Pi(\alpha, \beta) := \{\pi \in \mathbb{R}_+^{|\alpha| \times |\beta|} : \pi 1_{|\beta|} = \alpha, \pi^\top 1_{|\alpha|} = \beta\}$ is the set of transportation plans between μ and ν , where $|\alpha|$ denotes the number of supports of α , $||\alpha||$ denotes total masses of α . Also, we denote \mathbf{u}_n is the uniform distribution over n supports (similar definition with \mathbf{u}_m). For a set of m samples $X^m := \{x_1, \dots, x_m\}$, P_{X^m} denotes the empirical measures $\frac{1}{m} \sum_{i=1}^m \delta_{x_i}$. For any $x \in \mathbb{R}$, we denote by $\lfloor x \rfloor$ the greatest integer

less than or equal to x . For any probability measure μ on the Polish measurable space (\mathcal{X}, Σ) , we denote $\mu^{\otimes m} (m \geq 2)$ as the product measure on the product measurable space $(\mathcal{X}^m, \Sigma^m)$.

2 BACKGROUND

In this section, we first restate the definition of Kantorovich optimal transport (OT) and unbalanced Optimal transport (UOT) between two empirical measures. After that, we review the definition of the mini-batch optimal transport (m-OT), mini-batch unbalanced optimal transport (m-UOT), and discuss misspecified matchings issue of m-OT and interpretability issue of m-UOT.

2.1 OPTIMAL TRANSPORT AND UNBALANCED OPTIMAL TRANSPORT

Let $\mathcal{X} := \{x_i\}_{i=1}^n$, $\mathcal{Y} := \{y_j\}_{j=1}^n$ be two interested samples. The corresponding empirical measures are denoted by $\mu_n := \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ and $\nu_n := \frac{1}{n} \sum_{j=1}^n \delta_{y_j}$.

Optimal Transport: The Kantorovich optimal transport (Villani, 2009; Peyré & Cuturi, 2019) between μ_n and ν_n is defined as follows:

$$\text{OT}(\mu_n, \nu_n) := \min_{\pi \in \Pi(\mu_n, \nu_n)} \langle C, \pi \rangle, \quad (1)$$

where C is the distance matrix (or equivalently cost matrix) between \mathcal{X} and \mathcal{Y} that is produced by a ground metric (e.g., Euclidean distance or other designed distances).

Unbalanced Optimal Transport: The unbalanced optimal transport (Chizat et al., 2018b) between μ_n and ν_n is defined as follows:

$$\text{UOT}_\phi^\tau(\mu_n, \nu_n) := \min_{\pi} \langle C, \pi \rangle + \tau D_\phi(\pi_1, \mu_n) + \tau D_\phi(\pi_2, \nu_n), \quad (2)$$

where C is the distance matrix, $\tau > 0$ is a regularized parameter, D_ϕ is a certain probability divergence (e.g., KL divergence and total variational distance), and π_1, π_2 are respectively the marginal distributions of non-negative measure π and correspond to μ_n, ν_n .

The computational cost of both OT and UOT problems are of order $\mathcal{O}(n^2/\varepsilon^2)$ and $\mathcal{O}(n^2/\varepsilon)$, respectively. Its solutions are obtained by running the Sinkhorn algorithm (Cuturi, 2013), which updates directly the whole $n \times n$ matrix. It means that storing an $n \times n$ matrix is unavoidable in this approach, thus it is important that the memory capacity is able to match the matrix size.

2.2 MINI-BATCH OPTIMAL TRANSPORT

In real applications, the number of samples n is usually very large (e.g., millions). It is due to the large-scale empirical measures or detailed discretization of continuous measures. Therefore, solving directly OT between μ_n and ν_n is generally impractical due to the limitation of computational devices, namely, memory constraints, and vast computation. As a solution, the original n samples of two measures are divided (via sampling with or without replacement) into subsets of m samples, which we refer to as mini-batches. The mini-batch size m is often chosen to be the largest number that the computational device can process. Then, a mini-batch framework is developed to aggregate the optimal transport between pairs of the corresponding mini-batches into a useful result.

Motivating examples: We now provide some motivating examples to further illustrate the practical importance of mini-batch methods. The first example is regarding training a deep learning model. In practice, it is trained by a loss that requires computing a large-scale OT, e.g., deep generative models (Genevay et al., 2018) and deep domain adaptation (Damodaran et al., 2018). For example, a GPU that has 8GB memory can only store a 1024×1024 float cost matrix. However, the size of the cost matrix is much smaller in practice since the memory also needs to store models and data. The second example is color transfer application when the numbers of pixels in both source and target images are very large (e.g., millions). The mini-batch approach is used to transport a small number of pixels from source images to a small number of pixels from target images (Fattras et al., 2020). That process is repeated until the histogram of the source images is similar to that of the target.

Related methods: As another option, stochastic optimization can be utilized to solve the Kantorovich dual form with parametric functions, i.e., Wasserstein GAN (Arjovsky et al., 2017; Leygonie et al.,




Figure 1: The illustration of Example 1 for m-OT. The green points and blue points are respectively the supports of the empirical measures μ_n and ν_n . Black solid arrows represent the optimal mappings between μ_n and ν_n . Red solid arrows represent misspecified mappings. The 5×5 matrix is the incomplete transportation matrix $\pi_{P_{X^m}, P_{Y^m}}^{OT}$ which is created from solving OT between P_{X^m} and P_{Y^m} . The boldness of color of arrows and entries of the transportation matrix represents their mass values.

2019). Due to the limitation of parametrization, this approach has been shown that provides a very different type of discrepancy from the original Wasserstein distance (Mallasto et al., 2019; Stanczuk et al., 2021). Recently, input convex neural networks are developed to approximate the Brenier potential (Makkuva et al., 2020). Nevertheless, due to limited power in approximating Brenier potential (Korotin et al., 2021), recent work has indicated that input convex neural networks are not sufficient for computing OT. Finally, both approaches require special choices of the ground metric of OT. Namely, Wasserstein GAN has to use the \mathcal{L}_1 norm to make the constraint of dual form into the Lipchitz constraint and Brenier potential exists only when the ground metric is \mathcal{L}_2 .

Next, we revise the definition of mini-batch optimal transport (m-OT) (Fratras et al., 2020; 2021b). To ease the ensuing presentation, some notations in that paper are adapted into our paper. To build a mini-batch of $1 \leq m \leq n$ points, we sample $X^m := \{x_1, \dots, x_m\}$ with or without replacement from \mathcal{X}^m (similarly, Y^m are drawn from \mathcal{Y}^m) where m is mini-batch size.

Definition 1. (Mini-batch Optimal Transport) For $1 \leq m \leq n$ and $k \geq 1$, X_1^m, \dots, X_k^m and Y_1^m, \dots, Y_k^m are sampled with or without replacement from \mathcal{X}^m and \mathcal{Y}^m respectively. The m-OT transportation cost and transportation plan between μ_n and ν_n are defined as follow:

$$m\text{-OT}_k(\mu_n, \nu_n) = \frac{1}{k} \sum_{i=1}^k OT(P_{X_i^m}, P_{Y_i^m}); \quad \pi^{m\text{-OT}_k} = \frac{1}{k} \sum_{i=1}^k \pi_{P_{X_i^m}, P_{Y_i^m}}^{OT}, \quad (3)$$

where $\pi_{P_{X_i^m}, P_{Y_i^m}}^{OT}$ is a transportation matrix that is returned by solving $OT(P_{X_i^m}, P_{Y_i^m})$. Note that, $\pi_{P_{X_i^m}, P_{Y_i^m}}^{OT}$ is expanded to a $n \times n$ matrix that has padded zero entries to indices which are different from those of X_i^m and Y_i^m .

We would like to recall that $k = 1$ is the choice that practitioners usually used in real applications.

Misspecified matchings issue of m-OT: m-OT suffers from the problem which we refer to as *misspecified mappings*. In particular, misspecified mappings are non-zero entries in $\pi_k^{m\text{-OT}}$ while they have values of zero in the optimal transport plan π between original measures μ_n and ν_n . To demonstrate this point, we consider the following simple example:

Example 1. Let μ_n, ν_n be two empirical distributions with 5 supports on 2D: $\{(0, 1), (0, 2), (0, 3), (0, 4), (0, 5)\}$ and $\{(1, 1), (1, 2), (1, 3), (1, 4), (1, 5)\}$. The optimal mappings between μ_n and ν_n , $\{(0, i) - (1, i)\}_{i=1}^5$ are shown in Figure 1. Assuming that we use mini-batches of size 3 for m-OT. We specifically consider a pair of mini-batches of samples $X^m = \{(0, 1), (0, 2), (0, 3)\}$ and $Y^m = \{(1, 3), (1, 4), (1, 5)\}$. Solving OT between X^m and Y^m turns into 3 misspecified mappings $(0, 1) - (1, 3)$, $(0, 2) - (1, 4)$, and $(0, 3) - (1, 5)$ that have masses $1/3$ (see Figure 1).

2.3 MINI-BATCH UNBALANCED OPTIMAL TRANSPORT

Currently, (Fratras et al., 2021b) had shown that misspecified mappings are caused by matching two mini-batches which are not the optimal matched. To mitigate the misspecified matching issue,

they propose to use unbalanced optimal transport as the transportation type between samples of mini-batches. The mini-batch unbalanced optimal transport is defined as follow:

Definition 2. (*Mini-batch Unbalanced Optimal Transport*) For $1 \leq m \leq n$, $k \geq 1$, $\tau > 0$, a given divergence D_ϕ , X_1^m, \dots, X_k^m and Y_1^m, \dots, Y_k^m are sampled with or without replacement from \mathcal{X}^m and \mathcal{Y}^m respectively. The m -UOT transportation cost and transportation plan between μ_n and ν_n are defined as follow:

$$m\text{-UOT}_k^{\phi, \tau}(\mu_n, \nu_n) = \frac{1}{k} \sum_{i=1}^k \text{UOT}_\phi^\tau(P_{X_i^m}, P_{Y_i^m}); \quad \pi^{m\text{-UOT}_k^{\phi, \tau}} = \frac{1}{k} \sum_{i=1}^k \pi_{P_{X_i^m}, P_{Y_i^m}}^{\text{UOT}_\phi^\tau}, \quad (4)$$

where $\pi_{P_{X_i^m}, P_{Y_i^m}}^{\text{UOT}_\phi^\tau}$ is a transportation matrix that is returned by solving $\text{UOT}_\phi^\tau(P_{X_i^m}, P_{Y_i^m})$. Note that $\pi_{P_{X_i^m}, P_{Y_i^m}}^{\text{UOT}_\phi^\tau}$ is expanded to a $n \times n$ matrix that has padded zero entries to indices which are different from of X_i^m and Y_i^m .

Example: Going back to Example 1, UOT can reduce masses on misspecified matchings by relaxing the marginals of the transportation plan. Using D_ϕ as KL divergence, we show the illustration of m-UOT results in Figure 4 in Appendix C.1. We would like to recall that the regularized coefficient τ controls the degree of the marginal relaxation in m-UOT.

Some issues of m-UOT: The m-UOT has some issues which originally come from the nature of UOT. First, the ‘‘transport plan’’ for the UOT is hard to interpret since the UOT is developed for measures with different total masses. Second, the magnitude of regularized parameter τ depends on the cost matrix in order to make the regularization effective. Hence we need to search for τ in the wide range of \mathbb{R}^+ , which is a problem when the cost matrix changes its magnitude. We illustrate a simulation to demonstrate that the transportation plan of UOT for a fixed parameter τ changes after scaling supports by a constant in Figure 5 in Appendix C.1. As a result, choosing τ is challenging to some extent. In addition, the transportation plan of UOT tends to have non-zero entries due to the effect of KL divergence (a similar phenomenon tends to occur with other popular divergences that need to use a density ratio). Because of these weaknesses, it is not very convenient to apply m-UOT into practical applications. Also, it is non-trivial to explain m-UOT’s behaviors when we try to alleviate misspecified matchings.

3 MINI-BATCH PARTIAL OPTIMAL TRANSPORT

To address the misspecified mappings issue of m-OT and the aforementioned issue of m-UOT, in this section we propose a novel mini-batch approach, named *mini-batch partial optimal transport* (m-POT), that uses *partial optimal transport* (POT) as the transportation at the mini-batch level. We first review the definition of partial optimal transport in Section 3.1. Then, we define mini-batch partial optimal transport and discuss its properties in Section 3.2. Moreover, we illustrate that POT can be a natural choice of transportation among samples of mini-batches via simple simulations.

3.1 PARTIAL OPTIMAL TRANSPORT

Now, we restate the definition of partial optimal transport (POT) that is defined in (Figalli, 2010). Similar to the definition of transportation plans, we define the notion of partial transportation plans. Let $0 < s \leq 1$ to be transportation fraction. Partial transportation plan between two discrete probability measures α and β is $\Pi_s(\alpha, \beta) := \{\pi \in \mathbb{R}_+^{|\alpha| \times |\beta|} : \pi 1_{|\beta|} \leq \alpha, \pi^\top 1_{|\alpha|} \leq \beta, 1^\top \pi 1 = s\}$. With previous notations, the partial optimal transport between μ_n and ν_n is defined as follow:

$$\text{POT}_s(\mu_n, \nu_n) = \min_{\pi \in \Pi_s(\mathbf{u}_n, \mathbf{u}_n)} \langle C, \pi \rangle, \quad (5)$$

where C is the distance matrix. Equation (5) can be solved by adding dummy points (according to (Chapel et al., 2020)) to expand the cost matrix $\bar{C} = \begin{bmatrix} C & 0 \\ 0 & A \end{bmatrix}$, where $A > \|C\|_\infty$. In this case, solving the POT turns into solving the following OT problem:

$$\min_{\pi \in \Pi(\bar{\alpha}, \bar{\alpha})} \langle \bar{C}, \pi \rangle, \quad (6)$$




Figure 2: The illustration of Example 1 for m-POT. The meanings of green, blue points and arrows are similar to those in Figure 1. The parameter s is a fraction of masses of POT. The 5×5 matrix is the incomplete transportation matrix $\pi_{P_{X^m}, P_{Y^m}}^{POT_s}$ which is created from solving POT between P_{X^m} and P_{Y^m} . The boldness of color of arrows and entries of the transportation matrix represent their mass values.

with $\bar{\alpha} = [\mathbf{u}_n, 1 - s]$. Furthermore, the optimal partial transportation plan in equation (5) can be derived from removing the last row and column of the optimal transportation plan in equation (6).

3.2 MINI-BATCH PARTIAL OPTIMAL TRANSPORT

The partial transportation naturally fits the mini-batch setting since it can decrease the transportation masses of misspecified mappings (cf. the illustration in Figure 2 when two mini-batches contain optimal matching of the original transportation plan). Specifically, reducing the number of masses to be transported, i.e., reducing s , (from right images to left images in Figure 2) returns globally better mappings. With the right choice of the transport fraction, we can select mappings between samples that are as optimal as doing full optimal transportation. Moreover, POT is also stable to compute since it boils down to OT. Therefore, there are several solvers that can be utilized to compute POT.

Now, we define *mini-batch partial optimal transport* (m-POT) between μ_n and ν_n as follow:

Definition 3. (Mini-batch Partial Optimal Transport) For $1 \leq m \leq n$, $k \geq 1$, $0 < s \leq 1$, X_1^m, \dots, X_k^m and Y_1^m, \dots, Y_k^m are sampled with or without replacement from \mathcal{X}^m and \mathcal{Y}^m respectively. The m-POT transportation cost and transportation plan between μ_n and ν_n are defined as follow:

$$m\text{-POT}_k^s(\mu_n, \nu_n) = \frac{1}{k} \sum_{i=1}^k POT_s(P_{X_i^m}, P_{Y_i^m}); \quad \pi^{m\text{-POT}_k^s} = \frac{1}{k} \sum_{i=1}^k \pi_{P_{X_i^m}, P_{Y_i^m}}^{POT_s}, \quad (7)$$

where $\pi_{P_{X_i^m}, P_{Y_i^m}}^{POT_s}$ is a transportation matrix that is returned by solving $POT_s(P_{X_i^m}, P_{Y_i^m})$. Note that $\pi_{P_{X_i^m}, P_{Y_i^m}}^{POT_s}$ is expanded to a $n \times n$ matrix that has padded zero entries to indices which are different from of X_i^m and Y_i^m .

Computational complexity of m-POT: From the equivalence form of POT in equation (6), we have an equivalent form of m-POT in Definition 3 as follows:

$$m\text{-POT}_k^s(\mu_n, \nu_n) = \frac{1}{k} \sum_{i=1}^k \min_{\pi \in \Pi(\bar{\alpha}_i, \bar{\alpha}_i)} \langle \bar{C}_i, \pi \rangle. \quad (8)$$

Here, $\bar{C}_i = \begin{bmatrix} C_i & 0 \\ 0 & A_i \end{bmatrix} \in \mathbb{R}_+^{(m+1) \times (m+1)}$ and $\bar{\alpha}_i = [\mathbf{u}_m, 1 - s]$ where C_i is a cost matrix formed by the differences of elements of X_i^m and Y_i^m and $A_i > \|C_i\|_\infty$ for all $i \in [k]$. The computational complexity of approximating each OT problem in equation (8) is at the order of $\mathcal{O}\left(\frac{(m+1)^2}{\varepsilon^2}\right)$ (Altschuler et al., 2017; Lin et al., 2019) where $\varepsilon > 0$ is the tolerance. Therefore, the total computational complexity of approximating mini-batch POT is at the order of $\mathcal{O}\left(\frac{k(m+1)^2}{\varepsilon^2}\right)$. It is comparable to the computational complexity of m-OT, which is of the order of $\mathcal{O}\left(\frac{k \cdot m^2}{\varepsilon^2}\right)$ and slightly larger than that of m-UOT, which is $\mathcal{O}\left(\frac{k \cdot m^2}{\varepsilon}\right)$ (Pham et al., 2020), in terms of ε .

Table 1: DA results in classification accuracy on digits datasets (higher is better).

Methods	SVHN to MNIST	USPS to MNIST	MNIST to USPS	Avg
DANN	95.59	94.87	92.03	94.16
ALDA	98.73	98.38	95.47	97.53
m-OT	93.78	96.97	87.59	92.78
m-UOT	99.06	98.75	95.76	97.86
m-POT	98.71	98.58	95.76	97.68
em-POT	99.08	98.72	96.06	97.95

Table 2: DA results in classification accuracy on VisDA dataset (higher is better).

Methods	Accuracy
DANN	68.09
ALDA	71.38
m-OT	55.21
m-UOT	71.52
m-POT	72.46
em-POT	72.41

Concentration of m-POT: We first provide a guarantee on the concentration of the m-POT’s value for any given mini-batch size m and given the number of mini-batches k .

Theorem 1. *For any given number of minibatches $k \geq 1$ and minibatch size $1 \leq m \leq n$, there exists universal constant C such that with probability $1 - \delta$ we have*

$$\mathbb{P} \left(|m\text{-POT}_k^s(\mu_n, \nu_n) - m\text{-POT}^s(\mu, \nu)| \geq C \left(\sqrt{\frac{1}{\lfloor \frac{n}{m} \rfloor}} + \frac{1}{\sqrt{k}} \right) \sqrt{\log(2/\delta)} \right) \leq \delta,$$

where $m\text{-POT}^s(\mu, \nu) := \mathbb{E}_{X \sim \mu^{\otimes m}, Y \sim \nu^{\otimes m}} [POT_s(P_X^m, P_Y^m)]$.

The proof of Theorem 1 is in Appendix B.1. Due to space constraints, we further provide a concentration bound for the transportation plan of the m-POT in Theorem 2 in Appendix B.2. These results indicate that m-POT has good concentration behaviors around its expectation and its full mini-batch version (see Definition 4 in Appendix B).

Practical consideration for m-POT: First of all, as indicated in equation (8), m-POT can be converted to m-OT with mini-batch size $m + 1$. Therefore, it is slightly more expensive than m-OT and m-UOT in terms of memory and computation. The second issue of m-POT is the dependence on the choice of fraction of masses s because s plays a vital role in alleviating misspecified mappings from m-OT. At the first glance, choosing s may seem as inconvenient as choosing τ in m-UOT; however, it appears that searching for good s is actually easier than τ . Specifically, the range of s , which is $(0, 1]$, is small compared to that of τ , which is \mathbb{R}^+ . For example, we show that the transportation plan of POT for a fixed parameter s is the same when we scale the supports of two measures by a constant in Figure 6 in Appendix C.1

Discussion on the fraction of masses s : First, as indicated in toy examples in Figure 2, choosing small s is a way to mitigate misspecified mappings of m-OT; however, it may also cut off other good mappings when they exist. Therefore, small s tends to require a larger number of mini-batches to retrieve enough optimal mappings of the full-OT between original measures. Due to that issue, the fraction of masses s should be chosen adaptively based on the number of good mappings that exist in a pair of mini-batches. In particular, if a pair of mini-batches contains several optimal pairs of samples, s will be set to a high value. Nevertheless, knowing the number of optimal matchings in each pair of mini-batches requires additional information of original large-scale measures. Since designing an adaptive algorithm for s is a non-trivial open question, we leave this to future work. In the paper, we will only use grid searching in a set of chosen values of s in our experiments and use it for all mini-batches.

4 EXPERIMENTS

In this section, we focus on discussing experimental results on two applications, namely, deep domain adaption (deep DA) and deep generative model (DGM). In particular, we compare the performance of m-POT with previous mini-batch optimal transport approaches, m-OT and m-UOT, and some baseline methods on deep domain adaptation. On deep generative models, we learn generators using the m-OT loss and the m-POT loss while we have not been able to make m-UOT loss to work in this model. In the supplementary, we also conduct simulations on mini-batch transportation matrices and experiments on color transfer application to show that m-POT can provide a good barycentric mapping. A simple experiment on gradient flow is also carried out to illustrate the benefit of m-POT compared to m-OT and m-UOT. The details of applications and their algorithms are given in Appendix A. In Appendix C, we report detail of results on deep DA, DGM, and results of other applications. The detailed experimental settings of all applications are deferred to Appendix D.

Table 3: DA results in classification accuracy on Office-Home dataset (higher is better).

Methods	A2C	A2P	A2R	C2A	C2P	C2R	P2A	P2C	P2R	R2A	R2C	R2P	Avg
DANN	47.81	66.86	74.96	53.11	62.51	65.69	53.19	44.08	74.50	64.85	53.17	79.39	61.68
ALDA	54.07	74.81	77.28	61.97	71.64	72.96	59.54	51.34	76.57	68.15	56.40	82.11	67.24
m-OT	48.68	65.24	74.41	56.94	64.99	67.20	54.47	46.94	73.03	65.02	54.41	76.73	62.34
m-UOT	54.76	74.23	80.33	65.39	75.40	74.78	65.93	53.79	80.17	73.80	59.47	83.89	70.16
m-POT	54.85	73.51	80.7	65.76	73.62	76.02	64.28	53.5	79.76	74.08	59.22	83.17	69.87
em-POT	54.98	73.82	80.81	66.5	74.25	75.88	63.58	52.97	80.4	74.58	59.59	83.91	70.11

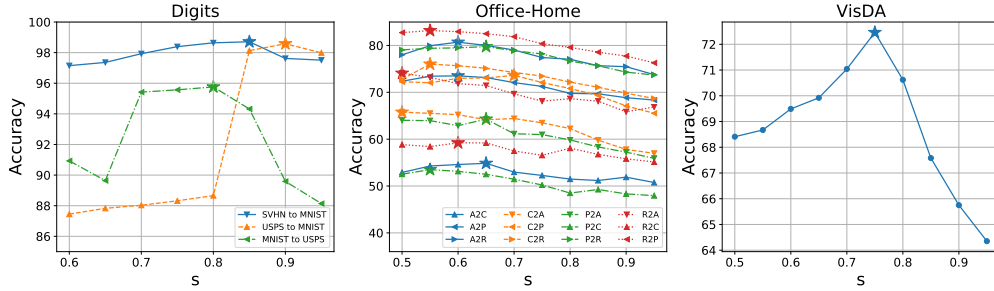


Figure 3: Performance of m-POT on deep DA when changing the fraction of masses s . The optimal values of s , which achieve the best accuracy, are marked by \star symbol. In the left figure, the optimal ratios for digits datasets lie between 0.8 and 0.9. In the middle figure, the best-performing values are smaller, from 0.5 to 0.7. On the VisDA dataset in the right figure, the optimal fraction of masses is 0.75.

4.1 DEEP DOMAIN ADAPTATION

In the following experiments, we conduct deep domain adaptation on a classification task on three types of datasets, namely, digits, Office-Home, and VisDA. To evaluate final models that are obtained from different methods including mini-batch methods, we use the classification accuracy on the adapted domain. Details about datasets, architectures of neural networks, and hyper-parameters settings are given in Appendix D.1. We compare our method against other DA methods: DANN (Ganin et al., 2016), m-OT (DeepJDOT) (Damodaran et al., 2018), ALDA (Chen et al., 2020), and m-UOT (JUMBOT) (Fratras et al., 2021a). The same pre-processing techniques as in m-UOT are applied. We report the best-performing checkpoint for all methods in Tables 1-3. The detailed results for each set of parameters, e.g., learning rate, the number of mini-batches k , and the fraction of masses s , in all datasets are provided in Tables 5-7 in Appendix C.3.

Classification results: According to Tables 1-3, m-POT is better than m-OT in all datasets. In greater detail, m-POT results in significant increases of accuracy 4.9, 7.53, and 17.25 on digits, Office-home, and VisDA datasets, respectively. This phenomenon is expected since m-POT can mitigate the issue of misspecified matchings while m-OT cannot. For two baselines DANN and ALDA, m-POT also produces more favorable performance in all experiments. Compared to m-UOT, m-POT is comparable in digits datasets and Office-Home datasets, namely, m-POT gives only 0.18 lower accuracy (average) on digits datasets and 0.29 lower accuracy (average) on Office-Home datasets. In contrast, on VisDA dataset, m-POT gives higher accuracy than m-UOT (72.46 compared to 71.52). The entropic regularization of m-POT (em-POT) yields higher performance than m-POT on digits and Office-Home datasets and has comparable performance with m-POT on VisDA dataset.

Performance of baselines: We reproduce baseline methods instead of using the reported performance. Although our results do not match their numbers, we still manage to have some better results than those in the original papers. For digits datasets, m-OT achieves higher accuracy on the adaptation from USPS to MNIST while m-UOT gets better performance on both SVHN to MNIST and USPS to MNIST. For the Office-Home dataset, m-OT has higher accuracies in 12 out of 12 scenarios. ALDA and m-UOT achieved better performance in 8 and 6 scenarios, respectively.

Computational speed: We report the speed of m-OT, m-UOT, and m-POT on digits datasets in Appendix C.8. m-OT is the fastest method and its running time is nearly half of m-UOT’s running time. m-POT takes less time per iteration than m-UOT for all choices of s while still preserving the adaptation performance.

Table 4: FID scores of generators that are trained with m-OT and m-POT on CIFAR10 and CelebA datasets (smaller is better).

Dataset	m	m-OT	m-POT(0.7)	m-POT(0.75)	m-POT(0.8)	m-POT(0.85)	m-POT(0.9)	m-POT(0.95)	m-POT(0.98)	m-POT(0.99)
CIFAR10	100	67.90	74.97	71.34	74.36	69.34	69.50	68.30	66.77	69.85
	200	66.42	70.49	67.95	67.53	69.72	70.31	66.29	67.98	66.34
CelebA	100	54.16	60.92	58.86	67.71	65.67	53.76	61.53	59.3	55.13
	200	56.85	49.25	53.95	58.34	53.36	52.52	55.67	49.76	59.32

The role of s : We demonstrate the effect of choosing s in m-POT. We observe a general pattern in all DA experiments as follows. When s comes closer to the optimal mass, the accuracy increases, then dropping as s moves away. As discussed in Section 3, the right value of s is important in making m-POT perform well. In Figure 3, we show the accuracy of different values of s on all datasets. From that figure, the best value of s is the value that is not too big and not too small. The reason is that a large value of s creates more misspecified matchings while a small value of s might drop off too many mappings. The latter leads to a lazy gradient, thus slowing the learning process.

4.2 DEEP GENERATIVE MODEL

In this section, we compare m-OT and m-POT on the deep generative model application. We also show the results for different values of the mini-batch size $m \in \{100, 200\}$ and the fraction of masses $s \in \{0.7, 0.75, 0.8, 0.85, 0.9, 0.95, 0.98, 0.99\}$. We describe the experiment settings including datasets, neural network architectures, and the used parameters in Appendix D.2.

Comparison between m-OT and m-POT: By evaluating the FID score, we observe that m-POT yields better generators than m-OT on both CIFAR10 and CelebA. The quantitative results are described in Table 4 where the number in bracket indicates the fraction of masses s . When $m = 100$, there is only one choice of s that leads to a lower FID score. As the mini-batch size increases, we have more options to choose a good value of s . Specifically, m-POT yields better performance in 2 and 6 different values of s for CIFAR10 and CelebA datasets, respectively. The qualitative results (randomly generated images) of m-OT and m-POT are given in Figures 12, 13 in Appendix C.4. From these images, we can conclude that m-POT produces more realistic generated images than m-OT. Thus, m-POT should be utilized as an alternative training loss for deep generative models such as those in (Tolstikhin et al., 2018; Genevay et al., 2018; Salimans et al., 2018).

Applications of m-UOT: We have not been able to achieve a good enough setting for m-UOT (a good enough setting means a setting that can provide a generative model that can generate reasonable images). To our best knowledge, there are not any generative models that are implemented with m-UOT, hence, we leave this comparison to future works.

Computational speed: The running speed of m-OT and m-POT in the deep generative model experiments are described in Table 15 in Appendix C.8. Two mini-batch approaches have the same number of iterations per second when using the mini-batch size of 100. As the mini-batch size increases to 200, m-OT runs faster than m-POT in both datasets at the cost of higher FID scores.

5 DISCUSSION

In this paper, we have introduced a novel mini-batch approach that is referred to as mini-batch partial optimal transport (m-POT). The new mini-batch approach can address the issue of misspecified mappings in the conventional mini-batch optimal transport approach and also provides an interpretable framework to understand such types of wrong matchings. Moreover, we also observe that m-POT is easier than mini-batch unbalanced optimal transport in choosing hyper-parameter. Finally, via extensive experiment studies, we demonstrate that m-POT is comparable to m-UOT in deep domain adaptation while it is better than m-OT in this application. Furthermore, in other applications, including deep generative model, color transfer, and gradient flow, m-POT is better than both m-OT and m-UOT. There are a few natural future directions arising from our work: (i) first, we will develop efficient algorithms to choose the fraction of masses s of m-POT adaptively; (ii) second, we would like to explore further the dependence structure between mini-batches (Nguyen et al., 2021b) and implement m-POT in more large-scale experiments.

Reproducibility Statement: We provide our source codes for all experiments in the paper in the supplementary of the paper. We also provide careful instruction for running our codes in an independent "Readme" file. The details of experimental settings, computational infrastructure, and other used public libraries are given in Appendix D. All datasets in the paper are public, hence, they are easy to be downloaded.

Ethics Statement: Given the methodological nature of the work, we do not foresee any negative societal and ethical impacts of the work.

REFERENCES

- Jason Altschuler, Jonathan Niles-Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *Advances in neural information processing systems*, pp. 1964–1974, 2017.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pp. 214–223, 2017.
- Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and Radon Wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 1(51):22–45, 2015.
- Laetitia Chapel, Mokhtar Alaya, and Gilles Gasso. Partial optimal transport with applications on positive-unlabeled learning. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, 2020.
- Minghao Chen, Shuai Zhao, Haifeng Liu, and Deng Cai. Adversarial-learned loss for domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 3521–3528, 2020.
- Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of Computation*, 87(314):2563–2609, 2018a.
- Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal transport: Dynamic and Kantorovich formulations. *Journal of Functional Analysis*, 274(11): 3090–3123, 2018b.
- Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 39(9):1853–1865, 2016.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pp. 2292–2300, 2013.
- Bharath Bhushan Damodaran, Benjamin Kellenberger, Rémi Flamary, Devis Tuia, and Nicolas Courty. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 447–463, 2018.
- Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen Zhao, David Forsyth, and Alexander G Schwing. Max-sliced Wasserstein distance and its use for GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10648–10656, 2019.
- P. Dvurechensky, A. Gasnikov, and A. Kroshnin. Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn’s algorithm. In *ICML*, pp. 1367–1376, 2018.
- Kilian Fatras, Younes Zine, Rémi Flamary, Rémi Gribonval, and Nicolas Courty. Learning with minibatch Wasserstein: asymptotic and gradient properties. In *AISTATS 2020-23rd International Conference on Artificial Intelligence and Statistics*, volume 108, pp. 1–20, 2020.

- Kilian Fatras, Thibault Sejourne, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3186–3197. PMLR, 18–24 Jul 2021a. URL <http://proceedings.mlr.press/v139/fatras21a.html>.
- Kilian Fatras, Younes Zine, Szymon Majewski, Rémi Flamary, Rémi Gribonval, and Nicolas Courty. Minibatch optimal transport distances; analysis and applications. *arXiv preprint arXiv:2101.01792*, 2021b.
- Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trounev, and Gabriel Peyré. Interpolating between optimal transport and MMD using Sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2681–2690, 2019.
- Alessio Figalli. The optimal partial transport problem. *Archive for rational mechanics and analysis*, 195(2):533–560, 2010.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- Aude Genevay, Gabriel Peyré, and Marco Cuturi. Learning generative models with Sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pp. 1608–1617. PMLR, 2018.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, pp. 6626–6637, 2017.
- Nhat Ho, XuanLong Nguyen, Mikhail Yurochkin, Hung Hai Bui, Viet Huynh, and Dinh Phung. Multilevel clustering via Wasserstein means. In *International Conference on Machine Learning*, pp. 1501–1509, 2017.
- Nhat Ho, Viet Huynh, Dinh Phung, and Michael I. Jordan. Probabilistic multilevel clustering via composite transportation distance. In *AISTATS*, 2019.
- Jonathan J. Hull. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5):550–554, 1994.
- Alexander Korotin, Lingxiao Li, Justin Solomon, and Evgeny Burnaev. Continuous Wasserstein-2 barycenter estimation without minimax optimization. *arXiv preprint arXiv:2102.01752*, 2021.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- Trung Le, Tuan Nguyen, Nhat Ho, Hung Bui, and Dinh Phung. LAMDA: Label matching deep domain adaptation. In *ICML*, 2021.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Jacob Leygonie, Jennifer She, Amjad Almahairi, Sai Rajeswar, and Aaron Courville. Adversarial computation of optimal transport maps. *arXiv preprint arXiv:1906.09691*, 2019.
- Tianyi Lin, Nhat Ho, and Michael Jordan. On efficient optimal transport: An analysis of greedy and accelerated mirror descent algorithms. In *International Conference on Machine Learning*, pp. 3982–3991, 2019.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Ashok Makkuva, Amirhossein Taghvaei, Sewoong Oh, and Jason Lee. Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning*, pp. 6672–6681. PMLR, 2020.

- Anton Mallasto, Guido Montúfar, and Augusto Gerolin. How well do WGANs estimate the Wasserstein metric? *arXiv preprint arXiv:1910.03875*, 2019.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- Khai Nguyen, Nhat Ho, Tung Pham, and Hung Bui. Distributional sliced-Wasserstein and applications to generative modeling. In *International Conference on Learning Representations*, 2021a. URL <https://openreview.net/forum?id=QYj070ACDK>.
- Khai Nguyen, Quoc Nguyen, Nhat Ho, Tung Pham, Hung Bui, Dinh Phung, and Trung Le. BoMb-OT: On batch of mini-batches optimal transport. *arXiv preprint arXiv:2102.05912*, 2021b.
- Khai Nguyen, Son Nguyen, Nhat Ho, Tung Pham, and Hung Bui. Improving relational regularized autoencoders with spherical sliced fused Gromov-Wasserstein. In *International Conference on Learning Representations*, 2021c. URL <https://openreview.net/forum?id=DiQD7FWL233>.
- T. Nguyen, Q.-H. Pham, T. Le, T. Pham, N. Ho, and B.-S. Hua. Point-set distances for learning representations of 3D point clouds. In *ICCV*, 2021d.
- O. Pele and M. Werman. Fast and robust earth mover’s distance. In *ICCV*. IEEE, 2009.
- Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- K. Pham, K. Le, N. Ho, T. Pham, and H. Bui. On unbalanced optimal transport: An analysis of sinkhorn algorithm. In *ICML*, 2020.
- Tim Salimans, Han Zhang, Alec Radford, and Dimitris Metaxas. Improving GANs using optimal transport. In *International Conference on Learning Representations*, 2018.
- Justin Solomon, Gabriel Peyré, Vladimir G Kim, and Suvrit Sra. Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG)*, 35(4):72, 2016.
- Max Sommerfeld, Jörn Schrieber, Yoav Zemel, and Axel Munk. Optimal transport: Fast probabilistic approximation with exact solvers. *Journal of Machine Learning Research*, 20:105–1, 2019.
- Jan Stanczuk, Christian Etmann, Lisa Maria Kreusser, and Carola-Bibiane Schönlieb. Wasserstein GANs work because they fail (to approximate the Wasserstein distance). *arXiv preprint arXiv:2103.01678*, 2021.
- Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018.
- Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.

Supplement to "Improving Mini-batch Optimal Transport via Partial Transportation"

In this supplement, we first discuss applications of OT that have been used in mini-batch fashion and their corresponding algorithms in Appendix A. Next, we derive concentration bounds for m-POT's value and m-POT's transportation plan in Appendix B. Moreover, we provide additional experimental results in Appendix C. In particular, we demonstrate the illustration of transportation plans from m-OT, m-UOT, and m-POT in Appendix C.1 and Appendix C.2. Moreover, we present detailed results on deep domain adaptation in Appendix C.3, deep generative model in Appendix C.4. Furthermore, we conduct experiments on color transfer application and gradient flow application to show the favorable performance of m-POT in Appendix C.5 and Appendix C.6 respectively. The comparison between m-POT and m-OT (m-UOT) with one additional sample in a mini-batch is given in Appendix C.7. The computational time of mini-batch methods in applications is reported in Appendix C.8. Finally, we report the experimental settings including neural network architectures, hyper-parameter choices in Appendix D.

A APPLICATIONS TO DEEP UNSUPERVISED DOMAIN ADAPTION, DEEP GENERATIVE MODEL, COLOR TRANSFER, AND GRADIENT FLOW

In this section, we first state two popular applications that benefit from using mini-batches, namely, deep domain adaption and deep generative model in Appendix A.1 and Appendix A.2. We also include detailed algorithms for these two applications and the way that we evaluate them. Next, we review the mini-batch color transfer algorithm in Appendix A.3. Finally, we discuss the usage of mini-batch losses in gradient flow application in Appendix A.4.

A.1 MINI-BATCH DEEP DOMAIN ADAPTION

We follow the setting of DeepJDOT (Damodaran et al., 2018) that is composed of two parts: an embedding function $G: \mathcal{X} \rightarrow \mathcal{Z}$ which maps data to the latent space; and a classifier $F: \mathcal{Z} \rightarrow \mathcal{Y}$ which maps the latent space to the label space on the target domain. The mini-batch version of DeepJDOT can be expressed as follow, for given the number of mini-batches k and the size of mini-batches m , the goal is to minimize the following objective function:

$$\min_{G,F} \frac{1}{k} \sum_{i=1}^k L_{OT_i}; \quad L_{OT_i} = \left(\frac{1}{m} \sum_{j=1}^m L_s(y_{ij}, F(G(s_{ij}))) + \min_{\pi \in \Pi(\mathbf{u}_m, \mathbf{u}_m)} \langle C_{S_i^m, Y_i^m, T_i^m}^{G,F}, \pi \rangle \right), \quad (9)$$

where L_s is the source loss function, S_1^m, \dots, S_k^m are source mini-batches that are sampled with or without replacement from the source domain $S^m \in \mathcal{X}^m$, Y_1^m, \dots, Y_k^m are corresponding labels of S_1^m, \dots, S_k^m , with $S_i^m = \{s_{i1}, \dots, s_{im}\}$ and $Y_1^m := \{y_{i1}, \dots, y_{im}\}$. Similarly, T_1^m, \dots, T_k^m ($T_1^m := \{t_{i1}, \dots, t_{im}\}$) are target mini-batches that are sampled with or without replacement from the target domain $\mathcal{T}^m \in \mathcal{X}^m$. The cost matrix $C_{S_i^m, Y_i^m, T_i^m}^{G,F}$ is defined as follows:

$$C_{1 \leq j, z \leq m} = \alpha \|G(s_{ij}) - G(t_{iz})\|^2 + \lambda_t L_t(y_{ij}, F(G(t_{iz}))), \quad (10)$$

where L_t is the target loss function, α and λ_t are hyper-parameters that control two terms.

JUMBOT: By replacing OT with UOT in DeepJDOT, authors in (Fattras et al., 2021a) introduce JUMBOT (joint unbalanced mini-batch optimal transport) which achieves the state-of-the-art domain adaptation results on various datasets. Therefore, the objective function in equation (9) turns into:

$$\min_{G,F} \frac{1}{k} \sum_{i=1}^k L_{UOT_i}; \quad L_{UOT_i} = \left(\frac{1}{m} \sum_{j=1}^m L_s(y_{ij}, F(G(s_{ij}))) + \min_{\pi} \langle C_{S_i^m, Y_i^m, T_i^m}^{G,F}, \pi \rangle + \tau (D_{\phi}(\pi_1, S_i^m) + D_{\phi}(\pi_2, T_i^m)) \right), \quad (11)$$

where π_1 and π_2 are marginals of non-negative measure π and respectively correspond to S_i^m and T_i^m .

Algorithm 1 Mini-batch Deep Domain Adaptation

Input: k, m , source domain (S, Y) , target domain T , chosen $L_{DA} \in \{L_{OT}, L_{UOT}, L_{POT}\}$
Initialize G_θ (parametrized by θ), F_ϕ (parametrized by ϕ)
while (θ, ϕ) do not converge **do**
 $\text{grad}_\theta \leftarrow \mathbf{0}; \text{grad}_\phi \leftarrow \mathbf{0}$
 for $i = 1$ **to** k **do**
 Sample $(s_1, y_i), \dots, (s_1, s_m)$ from (S, Y)
 Sample t_1, \dots, t_m from T
 $S^m \leftarrow \{s_1, \dots, s_m\}; Y^m \leftarrow \{y_1, \dots, y_m\}; T^m \leftarrow \{t_1, \dots, t_m\}$
 Compute $L_{DA} \leftarrow L_{DA}(S^m, Y^m, T^m, G_\theta, F_\phi)$
 $\text{grad}_\theta \leftarrow \text{grad}_\theta + \nabla_\theta L_{DA}$
 $\text{grad}_\phi \leftarrow \text{grad}_\phi + \nabla_\phi L_{DA}$
 end for
 $\theta \leftarrow \text{Adam}(\theta, \text{grad}_\theta)$
 $\phi \leftarrow \text{Adam}(\phi, \text{grad}_\phi)$
end while

Deep domain adaptation with m-POT: Similar to JUMBOT, by changing OT into POT with the fraction of masses s , we obtain the following objective loss:

$$\min_{G, F} \frac{1}{k} \sum_{i=1}^k L_{\text{POT}_i}; \quad L_{\text{POT}_i} = \left(\frac{1}{m} \sum_{j=1}^m L_s(y_{ij}, F(G(s_{ij}))) + \min_{\pi \in \Pi_s(\mathbf{u}_m, \mathbf{u}_m)} \langle C_{S_i^m, Y_i^m, T_i^m}^{G, F}, \pi \rangle \right), \quad (12)$$

Training Algorithms: We present the algorithm for training domain adaption with m-OT (mini-batch DeepJDOT), m-UOT (JUMBOT), and m-POT in a generalized algorithm which is stated in Algorithm 1.

We utilize Algorithm 1 to compare the performance of m-OT, m-UOT, and m-POT. The evaluation criterion is chosen based on the task on domains e.g. classification accuracy on the target domain (classification problems).

A.2 MINI-BATCH DEEP GENERATIVE MODEL

We first recall the setting of deep generative models. Given the data distribution $\mu_n := \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ with $x_i \in \mathcal{X}$, a prior distribution $p(z) \in \mathcal{P}(\mathcal{Z})$ e.g. $p(z) = \mathcal{N}(\mathbf{0}, \mathbf{I})$, and a generator (generative function) $G_\theta : \mathcal{Z} \rightarrow \mathcal{X}$ (where $\theta \in \Theta$ is a neural net). The goal of the deep generative model is to find the parameter θ^* that minimizes the discrepancy (e.g. KL divergence, optimal transport distance, etc) between μ_n and $G_\theta \# p(z)$ where the $\#$ symbol denotes the push-forward operator.

Due to the intractable computation of optimal transport distance (n is very large, implicit density of $G_\theta \# p(z)$), mini-batch losses (m-OT, m-UOT, m-POT) are used as surrogate losses to train the generator G_θ . The idea is to estimate the gradient of the mini-batch losses to update the neural network θ . In practice, the real metric space of data samples is unknown, hence, adversarial training is used as unsupervised metric learning (Genevay et al., 2018; Salimans et al., 2018). In partial, a discriminator function $F_\phi : \mathcal{X} \rightarrow \mathcal{H}$ where \mathcal{H} is a chosen feature space. The function F_ϕ is trained by maximizing the mini-batch OT loss. For greater detail, the training procedure is described in Algorithm 2.

Evaluation: To evaluate the quality of the generator G_θ , we use the FID score (Heusel et al., 2017) to compare the closeness of $G_\theta \# p(z)$ to μ_n .

On debiased mini-batch energy: Both m-OT and m-POT can be extended to debiased energy versions based on the work (Salimans et al., 2018). However, in this paper, we want to focus on the effect of misspecified matchings on applications, hence, the original mini-batch losses are used.

A.3 MINI-BATCH COLOR TRANSFER

In this section, we first state the mini-batch color transfer algorithm that we used to compare mini-batch methods in Algorithm 3. The idea is to transfer color from a part of a source image to a part

Algorithm 2 Mini-batch Deep Generative Model

Input: k, m , data distribution μ_n , prior distribution $p(z)$, chosen mini-batch loss $L_{\text{DGM}} \in \{\text{OT}, \text{UOT}, \text{POT}\}$
Initialize $G_\theta; F_\phi$
while θ does not converge **do**
 $\text{grad}_\theta \leftarrow \mathbf{0}$
 $\text{grad}_\phi \leftarrow \mathbf{0}$
 for $i = 1$ **to** k **do**
 Sample x_1, \dots, x_m from μ_n
 Sample z_1, \dots, z_m from $p(z)$
 Compute $y_1, \dots, y_m \leftarrow G_\theta(z_1), \dots, G_\theta(z_m)$
 $X^m \leftarrow \{x_1, \dots, x_m\}; Y^m \leftarrow \{y_1, \dots, y_m\}$
 Compute $L_{\text{DGM}}(F_\phi(X^m), F_\phi(Y^m))$
 $\text{grad}_\theta \leftarrow \text{grad}_\theta + \nabla_\theta L_{\text{DGM}}$
 $\text{grad}_\phi \leftarrow \text{grad}_\phi - \nabla_\phi L_{\text{DGM}}$
 end for
 $\theta \leftarrow \text{Adam}(\theta, \text{grad}_\theta)$
 $\phi \leftarrow \text{Adam}(\phi, \text{grad}_\phi)$
end while

Algorithm 3 Color Transfer with mini-batches

Input: k, m , source domain $X_s \in \mathbb{R}^{n \times d}$, target domain $X_t \in \mathbb{R}^{n \times d}$, $T \in \{\text{OT}, \text{UOT}, \text{POT}\}$
Initialize $Y_s \in \mathbb{R}^{n \times d}$
for $i = 1$ **to** k **do**
 Select set A of m samples in X_s
 Select set B of m samples in X_t
 Compute the cost matrix $C_{A,B}$
 $\pi \leftarrow T(C_{A,B}, \mathbf{u}_m, \mathbf{u}_m)$
 $Y_s|_A \leftarrow Y_s|_A + \pi \cdot X_t|_B$
end for
Output: Y_s

of a target image via a barycentric map that is obtained from two mini-batch measures (the source mini-batch and the target mini-batch). This process is repeated multiple times until reaching the chosen number of steps. The algorithm is introduced in (Fattras et al., 2020).

A.4 MINI-BATCH GRADIENT FLOW

In this appendix, we describe the gradient flow application and the usage of mini-batch methods in this case.

Gradient flow is a non-parametric method to learn a generative model. The goal is to find a distribution μ that is close to the data distribution ν in sense of a discrepancy D . So, a gradient flow can be defined as:

$$\partial_t \mu_t = -\nabla_{\mu_t} D(\mu_t, \nu) \quad (13)$$

We follow the Euler scheme to solve this equation as in (Feydy et al., 2019), starting from an initial distribution at time $t = 0$. In this paper, we consider using mini-batch losses such as m-OT, m-UOT, and m-POT as the discrepancy D .

B CONCENTRATION OF THE M-POT

In this appendix, we provide concentration bounds for the m-POT’s value and m-POT’s transportation plan. To ease the presentation, we first define the *full mini-batch* version of the m-POT’s value and m-POT’s transportation plan, namely, when we draw all the possible mini-batches from the set of all m elements of the data. We denote by $\binom{X^n}{m}$ and $\binom{Y^n}{m}$ the set of all m elements of

$\{X_1, X_2, \dots, X_n\}$ and $\{Y_1, Y_2, \dots, Y_n\}$ respectively. We define $\binom{X^n}{m} = \{X_{(1)}^m, X_{(2)}^m, \dots, X_{(S)}^m\}$ and $\binom{Y^n}{m} = \{Y_{(1)}^m, Y_{(2)}^m, \dots, Y_{(S)}^m\}$ where $S := \binom{n}{m}$.

Definition 4. (Full Mini-batch Partial Optimal Transport) For $1 \leq m \leq n$, $0 < s \leq 1$, the full mini-batch POT transportation cost and transportation plan between μ_n and ν_n are defined as follow:

$$m\text{-POT}^s(\mu_n, \nu_n) = \frac{1}{S^2} \sum_{i=1}^S \sum_{j=1}^S \text{POT}_s(P_{X_{(i)}^m}, P_{Y_{(j)}^m}); \quad \pi^{m\text{-POT}^s} = \frac{1}{S^2} \sum_{i=1}^S \sum_{j=1}^S \pi_{P_{X_{(i)}^m}, P_{Y_{(j)}^m}}^{\text{POT}_s}, \quad (14)$$

where $\pi_{P_{X_{(i)}^m}, P_{Y_{(j)}^m}}^{\text{POT}_s}$ is a transportation matrix that is returned by solving $\text{POT}_s(P_{X_{(i)}^m}, P_{Y_{(j)}^m})$. Note that $\pi_{P_{X_{(i)}^m}, P_{Y_{(j)}^m}}^{\text{POT}_s}$ is expanded to a $n \times n$ matrix that has padded zero entries to indices which are different from $X_{(i)}^m$ and $Y_{(j)}^m$.

B.1 CONCENTRATION OF THE M-POT'S VALUE

We provide the proof for Theorem 1. Recall that in the main text, we define $m\text{-POT}^s(\mu, \nu) := \mathbb{E}_{X \sim \mu^{\otimes m}, Y \sim \nu^{\otimes m}} [\text{POT}_s(P_{X^m}, P_{Y^m})]$.

Proof of Theorem 1. An application of triangle inequality leads to

$$\begin{aligned} |\text{m-POT}_k^s(\mu_n, \nu_n) - \text{m-POT}^s(\mu, \nu)| &\leq |\text{m-POT}_k^s(\mu_n, \nu_n) - \text{m-POT}^s(\mu_n, \nu_n)| \\ &\quad + |\text{m-POT}^s(\mu_n, \nu_n) - \text{m-POT}^s(\mu, \nu)| \\ &= T_1 + T_2. \end{aligned} \quad (15)$$

Bound for T_2 : A direct application of Hoeffding's inequality for U-statistics leads to

$$\mathbb{P}\left(T_2 \geq B \sqrt{\frac{\log(2/\delta)}{2 \lfloor \frac{n}{m} \rfloor}}\right) \leq \delta,$$

where $B = \max_{1 \leq i, j \leq S} \text{POT}_s(P_{X_{(i)}^m}, P_{Y_{(j)}^m})$. We now prove that for any X^m and Y^m , the $\text{POT}_s(P_{X^m}, P_{Y^m})$ is always bounded by a constant with high probability. In fact, from the equivalence between POT and OT, we can rewrite $\text{POT}_s(P_{X^m}, P_{Y^m})$ as follows:

$$\text{POT}_s(P_{X^m}, P_{Y^m}) = \min_{\pi \in \Pi(\bar{\alpha}, \bar{\alpha})} \langle \bar{C}, \pi \rangle,$$

where $\bar{\alpha} = [\mathbf{u}_m, 1 - s]$ and $\bar{C} = \begin{bmatrix} C & 0 \\ 0 & A \end{bmatrix}$. Here, we choose $A = \|C\|_\infty + 1$ and C is the cost matrix of the differences of the supports of P_{X^m} and P_{Y^m} . From here, we have

$$\min_{\pi \in \Pi(\bar{\alpha}_i, \bar{\alpha}_i)} \langle \bar{C}, \pi \rangle_F \leq \langle \bar{C}, \bar{\alpha} \bar{\alpha}^\top \rangle_F \leq (\|C\|_\infty + 1)(2 - s)^2.$$

Since $\|C\|_\infty < \infty$ with high probability, these bounds indicate that we can find universal constant C such that $\text{POT}_s(P_{X^m}, P_{Y^m}) \leq C$ with high probability. Therefore, we have $B \leq C$ and it indicates that

$$\mathbb{P}\left(T_2 \geq C \sqrt{\frac{\log(2/\delta)}{2 \lfloor \frac{n}{m} \rfloor}}\right) \leq \delta. \quad (16)$$

Bound for T_1 : Direct calculation shows that

$$T_1 = \left| \frac{1}{k} \sum_{i=1}^k \sum_{1 \leq j, l \leq S} \left(1_{\{(X_i^m, Y_i^m) \equiv (X_{(j)}^m, Y_{(l)}^m)\}} - \frac{1}{S^2} \right) \text{POT}(X_{(j)}^m, Y_{(l)}^m) \right|$$

We denote $Z_i = \sum_{1 \leq j, l \leq S} \left(1_{\{(X_i^m, Y_i^m) \equiv (X_{(j)}^m, Y_{(l)}^m)\}} - \frac{1}{S^2} \right) \text{POT}(X_{(j)}^m, Y_{(l)}^m)$ for $1 \leq i \leq k$. Then, given the data X_1, \dots, X_n and Y_1, \dots, Y_n , it is clear that Z_1, Z_2, \dots, Z_k are independent variables and bounded above by $2C$. A direct application of Hoeffding's inequality shows that

$$\mathbb{P}\left(T_1 \geq C \sqrt{\frac{\log(2/\delta)}{k}}\right) \leq \delta. \quad (17)$$

Plugging the results in equations (16) and (17) into equation (15), we obtain the conclusion of the theorem. \square

B.2 CONCENTRATION OF THE M-POT’S TRANSPORTATION PLAN

Now, we study the concentration of the m-POT’s transportation plan. We demonstrate that the row/ column sum of the m-POT’s transportation plan concentrates around the row/ column sum of the full m-POT’s transportation plan in Definition 4.

Theorem 2. *For any given number of batches $k \geq 1$ and minibatch size $1 \leq m \leq n$, there exists universal constant C such that with probability $1 - \delta$ we have*

$$\mathbb{P} \left(\left| \pi^{\text{m-POT}_k^s} 1_n - \pi^{\text{m-POT}^s} 1_n \right| \geq C \sqrt{\frac{\log(2/\delta)}{k}} \right) \leq \delta, \quad (18)$$

$$\mathbb{P} \left(\left| \left(\pi^{\text{m-POT}_k^s} \right)^\top 1_n - \left(\pi^{\text{m-POT}^s} \right)^\top 1_n \right| \geq C \sqrt{\frac{\log(2/\delta)}{k}} \right) \leq \delta, \quad (19)$$

where 1_n is the vector with all 1 values of its elements.

Proof. It is sufficient to prove equation (18). Direct calculation shows that

$$\pi^{\text{m-POT}_k^s} 1_n - \pi^{\text{m-POT}^s} 1_n = \frac{1}{k} \sum_{i=1}^k \sum_{1 \leq j, l \leq S} \left(1_{\{(X_i^m, Y_i^m) \equiv (X_{(j)}^m, Y_{(l)}^m)\}} - \frac{1}{S^2} \right) \pi_{P_{X_{(j)}^m}, P_{Y_{(l)}^m}}^{\text{POT}_s} 1_n.$$

We define $Z_i = \sum_{1 \leq j, l \leq S} \left(1_{\{(X_i^m, Y_i^m) \equiv (X_{(j)}^m, Y_{(l)}^m)\}} - \frac{1}{S^2} \right) \pi_{P_{X_{(j)}^m}, P_{Y_{(l)}^m}}^{\text{POT}_s} 1_n$ for $1 \leq i \leq k$. Conditioning on the data X_1, X_2, \dots, X_n and Y_1, Y_2, \dots, Y_n , the random variables Z_1, Z_2, \dots, Z_k are independent and upper bounded by $2s$. Therefore, a direct application of Hoeffding’s inequality shows that

$$\mathbb{P} \left(\left| \pi^{\text{m-POT}_k^s} 1_n - \pi^{\text{m-POT}^s} 1_n \right| \geq 2s \sqrt{\frac{\log(2/\delta)}{k}} \right) \leq \delta.$$

As a consequence, we obtain the conclusion of Theorem 2. \square

C ADDITIONAL EXPERIMENTS

We first visualize the transportation of the UOT and the POT in the context of mini-batch in Appendix C.1. Next, we run a toy example of 10×10 transportation problem to compare m-OT, m-UOT, m-POT, and full-OT in Appendix C.2. In the same section, we also investigate the role of transportation fraction s in m-POT. We then show tables that contain results of all run settings in domain adaptation in Appendix C.3. Then, we report detailed results of deep generative model experiments and show some randomly generated images in Appendix C.4. Moreover, we discuss experimental results on color transfer and gradient flow in Appendices C.5 and C.6, respectively. Next, we discuss the comparison between m-POT and m-OT (m-UOT) with the mini-batch size is raised by 1 for m-OT (m-UOT) in Appendix C.7. Finally, we discuss the computational time of m-UOT, m-POT, and m-OT in applications in Appendix C.8.

C.1 TRANSPORTATION VISUALIZATION

Transportation of UOT: We use Example 1 and vary the parameter τ to see the changing of the transportation plan of UOT. We show in Figure 4 the behavior of UOT that we observed. According to the figure, UOT’s transportation plan is always non-zero for every value of τ . Increasing τ leads to a higher value for entries of the transportation matrix. In alleviating misspecified matchings, UOT can cure the problem to some extent with the right choice of τ (e.g. 0.2, 0.5). In particular, the mass of misspecified matchings is small compared to the right matchings. However, it is not easy to know how many masses are transported in total in UOT.

Sensitivity to the scale of cost matrix: We consider an extended version of Example 1 where all the supports are scaled by a scalar 10. Again, we use the same set of $\tau \in \{0.2, 0.5, 1, 10\}$ then show the transportation matrices of UOT in Figure 5. From the figure, we can see that the good choice of τ

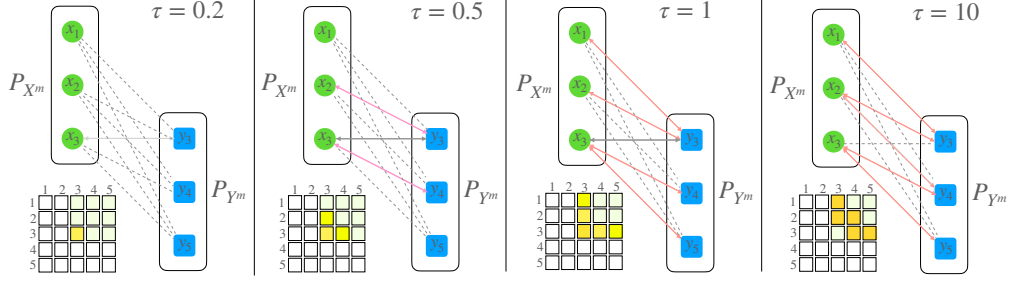


Figure 4: The illustration of Example 1 for m-UOT. As in Figure 1, the green points and blue points are respectively the supports of the empirical measures μ_n and ν_n . Black solid arrows represent the optimal mappings between μ_n and ν_n . Red solid arrows represent misspecified mappings. Dashed arrows are mappings that have very small masses. The parameter τ is the coefficient of the marginal relaxation term of UOT. The 5×5 matrix is the incomplete transportation matrix $\pi_{P_{X^m}, P_{Y^m}}^{\text{UOT}_\tau^\phi}$ which is created from solving UOT between P_{X^m} and P_{Y^m} . The boldness of color of arrows and entries of the transportation matrix represent their mass values.

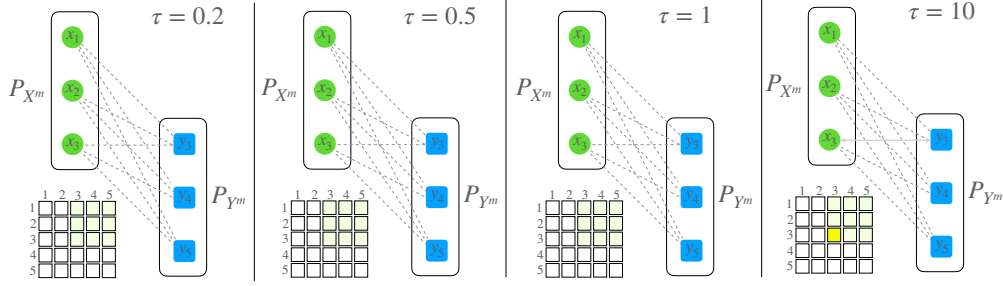


Figure 5: The UOT's illustration of Example 1 with supports of measures are scaled by 10. As in Figure 1, the green points and blue points are respectively the supports of the empirical measures μ_n and ν_n . Black solid arrows represent the optimal mappings between μ_n and ν_n . Red solid arrows represent misspecified mappings. Dashed arrows are mappings that have very small masses. The parameter τ is the coefficient of the marginal relaxation term of UOT. The 5×5 matrix is the incomplete transportation matrix $\pi_{P_{X^m}, P_{Y^m}}^{\text{UOT}_\tau^\phi}$ which is created from solving UOT between P_{X^m} and P_{Y^m} . The boldness of color of arrows and entries of the transportation matrix represent their mass values.




Figure 6: The POT's illustration of Example 1 supports of measures are scaled by 10. As in Figure 1, the green points and blue points are respectively the supports of the empirical measures μ_n and ν_n . Black solid arrows represent the optimal mappings between μ_n and ν_n . Red solid arrows represent misspecified mappings. Dashed arrows are mappings that have very small masses. The parameter s is the fraction of masses POT. The 5×5 matrix is the incomplete transportation matrix $\pi_{P_{X^m}, P_{Y^m}}^{\text{POT}_s}$ which is created from solving UOT between P_{X^m} and P_{Y^m} . The boldness of color of arrows and entries of the transportation matrix represent their mass values.

depends significantly on the scale of the cost matrix. So, UOT could not perform well in applications that change the supports of measures frequently such as deep generative models. Moreover, this sensitivity also leads to big efforts to search for a good value of τ in applications. In contrast, with the same set of the fraction of masses, s in Figure 2, the obtained transportation from POT is unchanged when we scale the supports of measures. We show the results in Figure 6 that suggests POT is a stabler choice than UOT.




Figure 7: The illustration of Example 1 in the non-overlapping case for both UOT, POT, and OT (POT $s = 1$). The appearance of the graphs are similar to previous figures.

Non-overlapping mini-batches: We now demonstrate the behavior of OT, UOT, and POT when a pair of mini-batches does not contain any global optimal mappings. According to the result of UOT, POT, and OT (POT $s=1$) in Figure 7, we can see that although UOT and POT cannot avoid misspecified matchings in this case, their solution is still better than OT in the sense that they can map a sample to the nearest one to it.

C.2 AGGREGATED TRANSPORTATION MATRIX

Comparison to m-OT and m-UOT: We demonstrate a simulation with 10 supports points drawn from bi-modal distributions $\frac{1}{2}\mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right) + \frac{1}{2}\mathcal{N}\left(\begin{bmatrix} 0 \\ 20 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\right)$ and $\frac{1}{2}\mathcal{N}\left(\begin{bmatrix} 10 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & -0.8 \\ -0.8 & 1 \end{bmatrix}\right) + \frac{1}{2}\mathcal{N}\left(\begin{bmatrix} 10 \\ 20 \end{bmatrix}, \begin{bmatrix} 1 & -0.8 \\ -0.8 & 1 \end{bmatrix}\right)$ respectively for the empirical measures μ_n and ν_n . We run m-OT, m-UOT (with KL divergence), and m-POT. All methods are run with $k = 32, m = 6$. Then, we visualize the mappings and transportation matrices in Figure 8. For illustration purposes, there are three numbers shown near the names of mini-batch approaches. They are the total number of mappings, the number of misspecified mappings, and the number of optimal mappings in turn. The total number of mappings is the number of non-zero entries in a transportation matrix. The number of optimal mappings is the number of mappings that exist in the full-OT’s transportation plan while the number of misspecified mappings is for mappings that do not exist. We can observe that m-POT and m-UOT provide more meaningful transportation plans than m-OT, namely, masses are put to connections between correct clusters. Importantly, we also observe that m-POT has lower misspecified matchings than m-OT (37 compared to 55). As mentioned in the background section, UOT tends to provide a transportation plan that contains non-zero entries; therefore, m-UOT still has many misspecified matchings through the weights of those matchings can be very small that can be ignored in practice. Note that, in the simulation results we select the best value of $\tau \in \{0.1, 0.5, 1, 2, 5, 10, 15, 20, 30, 40, 50\}$ for m-UOT and best value of $s \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$ for m-POT.

The role of the fraction of masses s : Here, we repeat the simulation in the main text with two empirical measures of 10 samples. The difference is that we run only m-POT with the value of the fraction of masses $s \in \{0.2, 0.4, 0.6, 0.8\}$. The result is shown in Figure 9. It is easy to see that a smaller value of s leads to a smaller number of misspecified mappings. However, a small value of s also removes good mappings (e.g. $s = 0.2$ only has 5 correct mappings). With the right choice of s (e.g. $s = 0.4$), we can obtain enough good mappings while having a reasonable number of misspecified matchings.

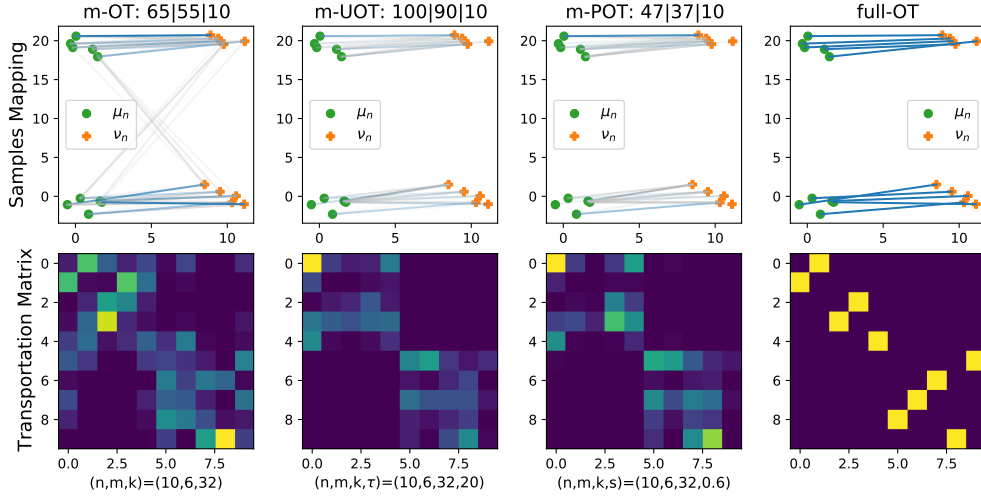


Figure 8: The first row presents sample mappings between μ_n and ν_n from different methods. These mappings are extracted from the transportation matrices which are shown in the second row. Blue lines represent optimal mappings while silver lines represent misspecified mappings. The boldness of the lines and entries of matrices are subject to their masses. There are three numbers near the name of mini-batch methods that are the total number of mappings, the number of misspecified mappings, and the number of optimal mappings respectively.

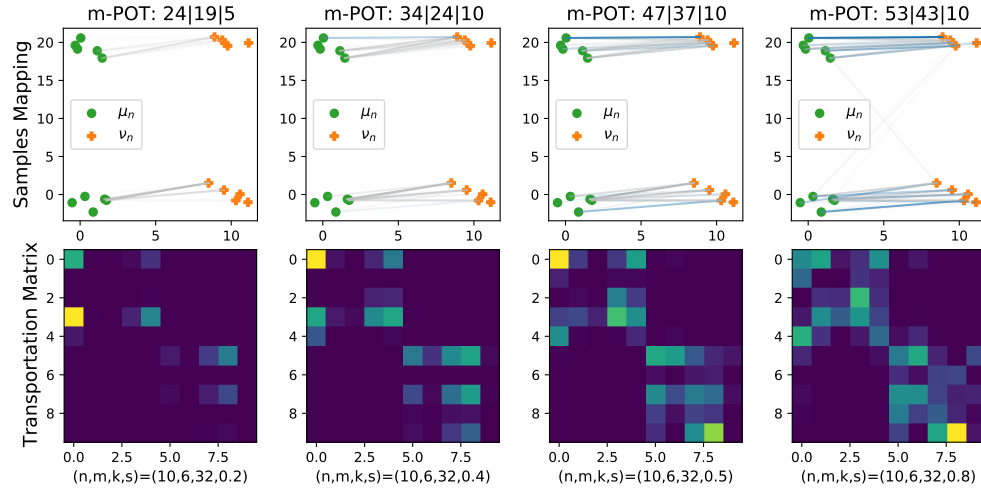


Figure 9: The first row presents sample mappings between μ_n and ν_n from different methods. These mappings are extracted from the transportation matrices which are shown in the second row. Blue lines represent optimal mappings while silver lines represent misspecified mappings. The boldness of the lines and entries of matrices are subject to their masses. There are three numbers near the name of mini-batch methods that are the total number of mappings, the number of misspecified mappings, and the number of optimal mappings respectively.

C.3 DEEP DOMAIN ADAPTATION EXPERIMENTS

We first want to emphasize that we have used the best parameters settings for m-UOT that are reported in (Fratras et al., 2021a) in each experiment.

Digits datasets: We vary the learning rate and the number of mini-batches. As can be seen from Table 5, m-POT outperforms m-OT in all combinations of hyper-parameters. Compared to m-UOT, m-UOT is slightly better than m-POT in adaptation from SVHN to MNIST and USPS to MNIST. On the other hand, m-POT is slightly better than m-UOT on MNIST to USPS adaptation. In general, we can conclude that m-UOT and m-POT are similar in terms of classification accuracy. In addition, a higher learning rate usually leads to higher accuracy for all methods.

Office-Home dataset: According to Table 6, m-POT is still much better than m-OT on Office-Home dataset. Similar to digits datasets, m-POT and m-UOT provide comparable classification accuracy on

Table 5: Details of DA results on digit datasets. The table involves multiple settings including the number of mini-batches k , the learning rate for m-OT, m-UOT, and m-POT. For m-POT, the fraction of mass s is the number in the bracket next to the name. Entries of the table are classification accuracy on target domain.

	k	LR	m-OT	m-UOT	m-POT(0.6)	m-POT(0.65)	m-POT(0.7)	m-POT(0.75)	m-POT(0.8)	m-POT(0.85)	m-POT(0.9)	m-POT(0.95)
S \rightarrow M	1	0.0002	93.19	98.72	96.33	97.36	97.57	97.87	97.61	97.24	97.45	96.17
		0.0003	93.23	98.95	96.52	97.07	97.42	98.26	98.4	97.88	96.86	97.51
		0.0004	93.84	99.06	96.86	97.09	97.80	98.14	97.69	98.71	97.28	95.96
	2	0.0002	93.41	98.83	97.02	97.25	97.93	97.01	98.46	97.55	97.13	93.48
		0.0003	93.66	98.89	97.15	97.14	97.80	98.39	98.63	98.66	97.42	93.71
		0.0004	93.78	99.03	96.99	97.29	97.70	97.09	98.64	98.09	97.62	94.43
	1	0.0002	96.72	97.12	80.82	87.17	87.68	87.96	88.20	97.71	97.85	97.80
		0.0003	96.88	98.63	86.50	87.42	87.85	88.13	88.47	97.90	98.06	97.99
		0.0004	96.97	98.75	86.76	87.67	87.97	88.32	88.57	97.93	98.20	97.93
	2	0.0002	96.90	97.42	80.12	87.66	87.85	88.05	88.66	97.99	98.24	97.64
		0.0003	96.95	97.47	81.13	87.75	87.91	88.15	88.47	98.05	98.45	97.61
		0.0004	96.90	98.28	87.45	87.83	88.03	88.21	88.59	98.13	98.58	97.68
M \rightarrow U	1	0.0002	86.95	95.47	90.53	88.04	95.02	95.02	95.52	93.82	89.49	88.14
		0.0003	87.05	95.57	90.83	89.64	95.07	95.17	95.42	93.82	89.59	87.64
		0.0004	85.80	95.47	90.93	89.64	95.07	95.22	95.52	94.32	89.24	87.74
	2	0.0002	87.59	95.67	87.94	87.74	95.27	95.57	95.67	93.32	88.74	87.19
		0.0003	86.85	95.67	89.44	87.89	95.42	95.47	95.76	93.22	88.69	87.79
		0.0004	86.90	95.76	90.58	87.79	95.37	95.52	95.76	93.42	89.29	87.49
	1	0.0002	86.95	95.47	90.53	88.04	95.02	95.02	95.52	93.82	89.49	88.14
		0.0003	87.05	95.57	90.83	89.64	95.07	95.17	95.42	93.82	89.59	87.64
		0.0004	85.80	95.47	90.93	89.64	95.07	95.22	95.52	94.32	89.24	87.74

Table 6: Details of DA results on Office-Home dataset. The table involves multiple settings including the number of mini-batches k , the learning rate for m-OT, m-UOT, and m-POT. For m-POT, the fraction of mass s is the number in the bracket next to the name. Entries of the table are classification accuracy on target domain.

Methods	A2C	A2P	A2R	C2A	C2P	C2R	P2A	P2C	P2R	R2A	R2C	R2P
DANN	47.81	66.86	74.96	53.11	62.51	65.69	53.19	44.08	74.50	64.85	53.17	79.39
ALDA	54.07	74.81	77.28	61.97	71.64	72.96	59.54	51.34	76.57	68.15	56.40	82.11
m-OT	48.68	65.24	74.41	56.94	64.99	67.20	54.47	46.94	73.03	65.02	54.41	76.73
m-UOT	54.76	74.23	80.33	65.39	75.40	74.78	65.93	53.79	80.17	73.80	59.47	83.89
m-POT(0.5)	52.90	72.36	78.00	65.76	72.25	72.80	64.03	52.53	79.05	74.08	58.83	82.74
m-POT(0.55)	54.27	73.44	79.94	65.51	72.04	76.02	63.95	53.50	79.40	73.18	58.42	83.17
m-POT(0.6)	54.59	73.51	80.70	65.23	72.99	75.67	62.88	53.13	79.48	71.86	59.22	82.95
m-POT(0.65)	54.85	73.15	80.01	64.15	73.06	75.12	64.28	52.51	79.76	71.45	59.20	82.50
m-POT(0.7)	52.97	72.07	79.02	64.40	73.62	74.23	61.15	51.46	78.86	69.67	57.48	81.87
m-POT(0.75)	52.26	71.26	77.42	63.49	72.09	73.45	60.98	50.22	78.20	68.11	56.56	80.36
m-POT(0.8)	51.48	69.77	77.05	62.26	70.80	72.18	59.83	48.50	76.66	68.64	58.08	79.59
m-POT(0.85)	51.18	69.68	75.67	59.79	69.32	71.08	58.34	49.26	75.67	68.11	56.75	78.56
m-POT(0.9)	51.89	68.85	75.47	57.81	67.09	69.82	57.23	48.32	74.30	65.84	55.81	77.74
m-POT(0.95)	50.75	68.35	73.77	56.98	65.53	68.67	55.91	47.95	73.70	66.87	55.14	76.28

the target domain. When the source domain is synthetic (A and C), m-POT outperforms m-UOT in two-third of scenarios. On the other hand, m-UOT achieves more favorable results for the real source domain (P and R). m-UOT is always better than m-POT if either source or target domain is Product (P).

VisDA dataset: Table 7 illustrates that m-POT yields the best performance in 6 out of 12 categories. Knife is a difficult category in which other baselines only achieve less than 20% accurate prediction. However, m-POT has an accuracy of 88.55%, which creates a large margin of nearly 70%. In terms of average class precision, m-POT also leads to a significant difference of at least 4.46% compared to other baselines.

The entropic regularized version of m-POT: We further run the entropic regularization of m-POT (em-POT) with the best-performing mass s . For digits datasets, we choose the regularization coefficient $\epsilon \in \{0.05, 0.1, 0.2\}$. For Office-Home dataset, we find the ϵ in the set $\{0.005, 0.01, 0.05, 0.1, 0.15, 0.2\}$. Finally, we select the best value of $\epsilon \in \{0.003, 0.004, 0.005, 0.01, 0.05, 0.1, 0.2\}$ for VisDA dataset. Table 8-10 illustrates the results of em-POT on three adaptation datasets. em-POT improves the performance over m-POT on digits and Office-Home datasets. The accuracy of em-POT (97.95) is even higher than that of m-UOT (97.86) on digits datasets. For VisDA dataset, em-POT performs slightly worse than m-POT, still leaving a gap of 0.89 compared to m-UOT.

Table 7: Details of DA results on VisDA dataset. The table involves multiple settings including the number of mini-batches k , the learning rate for m-OT, m-UOT, and m-POT. For m-POT, the fraction of mass s is the number in the bracket next to the name. Entries of the table are classification accuracy on target domain.

Methods	All	plane	bicycl	bus	car	house	knife	mcycl	person	plant	sktbrd	train	truck	Avg
DANN	68.09	85.77	91.41	67.57	68.25	91.39	9.79	59.99	78.86	72.40	37.92	61.47	47.40	64.35
ALDA	71.38	95.15	96.61	68.28	70.77	91.39	9.21	61.26	78.77	75.50	39.38	70.57	69.34	68.85
m-OT	55.21	59.49	87.44	69.06	68.24	78.19	19.54	44.96	55.66	62.33	21.75	46.54	41.37	54.55
m-UOT	71.52	97.07	94.00	64.19	74.69	91.40	7.06	65.49	74.88	84.83	40.28	64.66	45.39	67.00
m-POT(0.5)	68.41	80.64	91.86	81.27	64.96	88.15	51.90	67.98	67.00	66.97	30.33	57.49	56.89	67.12
m-POT(0.55)	68.67	84.71	94.47	72.98	67.04	90.49	53.57	64.84	65.79	68.30	29.81	61.51	64.29	68.15
m-POT(0.6)	69.49	82.33	93.57	72.09	69.81	86.67	66.03	67.07	65.98	68.26	28.81	60.47	70.10	69.27
m-POT(0.65)	69.92	82.31	93.89	71.37	70.63	88.21	69.44	67.38	63.59	71.43	29.74	60.46	68.50	69.75
m-POT(0.7)	71.04	85.92	92.63	75.10	72.37	89.08	88.55	58.05	68.57	76.65	29.64	62.17	80.97	73.31
m-POT(0.75)	72.46	87.22	92.23	72.74	76.15	92.41	74.11	59.87	71.53	74.89	35.28	62.80	71.83	72.59
m-POT(0.8)	70.62	82.40	91.61	64.31	69.41	92.65	70.18	61.98	70.95	76.10	27.60	65.96	53.02	68.85
m-POT(0.85)	67.58	71.53	92.33	67.19	70.16	90.93	7.42	71.76	74.36	72.23	24.75	52.88	36.73	61.02
m-POT(0.9)	65.75	93.44	83.76	64.78	73.35	84.72	4.83	58.64	75.27	69.60	27.92	49.89	43.93	60.84
m-POT(0.95)	64.35	79.21	80.31	65.71	64.90	74.95	61.04	53.03	66.13	78.66	24.48	60.03	20.61	60.76

Table 8: Results of em-POT on digits datasets. The table involves multiple settings including the number of mini-batches k , the learning rate, the fraction of mass s , and the entropic regularization coefficient ϵ .

	k	LR	s	ϵ	Accuracy
S \rightarrow M	1	0.0004	0.85	0.1	99.08
U \rightarrow M	2	0.0004	0.90	0.2	98.72
M \rightarrow U	2	0.0004	0.80	0.1	96.06
Avg					97.95

Table 9: Results of em-POT on Office-Home dataset. The table involves settings of the fraction of mass s and the entropic regularization coefficient ϵ .

	A2C	A2P	A2R	C2A	C2P	C2R	P2A	P2C	P2R	R2A	R2C	R2P	Avg
s	0.65	0.60	0.60	0.50	0.70	0.55	0.65	0.55	0.65	0.50	0.60	0.55	
ϵ	0.01	0.2	0.1	0.2	0.2	0.2	0.15	0.005	0.2	0.1	0.1	0.1	
Accuracy	54.98	73.82	80.81	66.50	74.25	75.88	63.58	52.97	80.4	74.58	59.59	83.91	70.11

Table 10: Results of em-POT on VisDA dataset ($s = 0.75, \epsilon = 0.004$).

	All	plane	bicycl	bus	car	house	knife	mcycl	person	plant	sktbrd	train	truck	Avg
Accuracy	72.41	95.6	94.54	70.01	73.10	92.80	4.65	61.77	74.25	88.63	39.31	67.31	73.99	69.66

C.4 DEEP GENERATIVE MODEL EXPERIMENTS

OT has been widely utilized to measure the discrepancy between a parametric distribution and real data distribution in a generative model. In this experiment, we apply m-OT and m-POT into the deep generative model application. We report the best-performing checkpoint, which is measured by FID score, along with their generated images from random noise. The generated images and the corresponding FID scores of m-OT and m-POT on CIFAR10 and CelebA datasets are given in Figure 12 and Figure 13, respectively. In terms of FID score, m-POT leads to a slight improvement of 1.13 over m-OT on CIFAR10 dataset. While it yields a larger difference of 7.6 on CelebA dataset. Randomly generated images are shown in Figure 12 and Figure 13 also indicate that m-POT is better than m-OT. For m-UOT, we tried very hard to make it work in this application, however, we could not have any good enough results from it. To our best knowledge, m-UOT has never been applied successfully on natural images. So, we suggest that m-OT should be replaced by m-POT for deep generative models.




Figure 10: Color Transfer results from m-OT, m-POT with $k = 10000$, $m = 100$. The leftmost image is the source image and the rightmost image is the target image. In between images transferred images with the corresponding name of the method on top.

Table 11: DA performance of on digits datasets for m-OT and m-UOT with one more sample in a mini-batch ($m + 1$).

Methods	SVHN to MNIST	USPS to MNIST	MNIST to USPS	Avg
m-OT($m=501$)	94.12	97.03	86.45	92.53
m-UOT($m=501$)	98.94	98.70	95.86	97.83
m-POT($m=500$)	98.71	98.58	95.76	97.68

Table 12: DA performance of equivalent problems on VisDA dataset for m-OT and m-UOT with one more sample in a mini-batch ($m + 1$).

Methods	Accuracy
m-OT($m=73$)	57.23
m-UOT($m=73$)	72.08
m-POT($m=72$)	72.46

C.5 COLOR TRANSFER

In this application, we run m-OT and m-POT to transfer color of images on two domains: arts and natural images. We show some transferred images from the m-OT and m-POT with two values of s (0.9 and 0.99) in Figure 10. We observe that m-POT with the right choice of parameter s generates more beautiful images than the m-OT, namely, m-POT ignores some color is too different between two images. Also, m-POT can work well on both two types of image domains. For m-UOT, We have tried to run m-UOT in this application, however, we have not found the setting that UOT can provide reasonable results yet.

C.6 GRADIENT FLOW

We compare m-OT, m-UOT, and m-POT in a toy example as in (Feydy et al., 2019) and present our results in Figure 11. In short, we want to move the colorful empirical measure to the "S-shape" measure. Each measure has 1000 support points. Here, we choose $(k, m) = (4, 4)$ and learning rate is set to 0.001. We use entropic regularization (Chizat et al., 2018a) to solve m-UOT, we choose the best entropic regularization parameter $\epsilon \in \{0.1, 1, 2, 5, 10\}$ and marginal relaxation parameter $\tau \in \{0.001, 0.01, 0.1, 1, 2, 5, 10\}$. From the figure, we observe that m-POT yields better flows than both m-OT and m-UOT, namely, the final Wasserstein-2 score of m-POT is the lowest. Interestingly, we see that reducing the amount of masses s in m-POT from 0.9 to 0.8 improves considerably the result. Here, we observe that although lower values of s provide a better generative model at the end, it can make the generative model learning slower. It suggests that s should be changed adaptively in training processes. We will leave this question for future works.

C.7 EQUIVALENT OPTIMAL TRANSPORT PROBLEM

As discussed in Section 3.2, m-POT with mini-batch size m is equivalent to m-OT with mini-batch size $m + 1$. Therefore, we run m-OT and m-UOT with the increasing mini-batch size to compare the performance. The classification accuracies on target datasets in deep DA are reported in Table 11-13. The results for the deep generative model can be found in Table 14. It can be observed clearly that m-POT still outperforms the corresponding version of m-OT in both deep domain adaptation and deep generative model applications. While m-UOT yields a higher performance on digits datasets, it

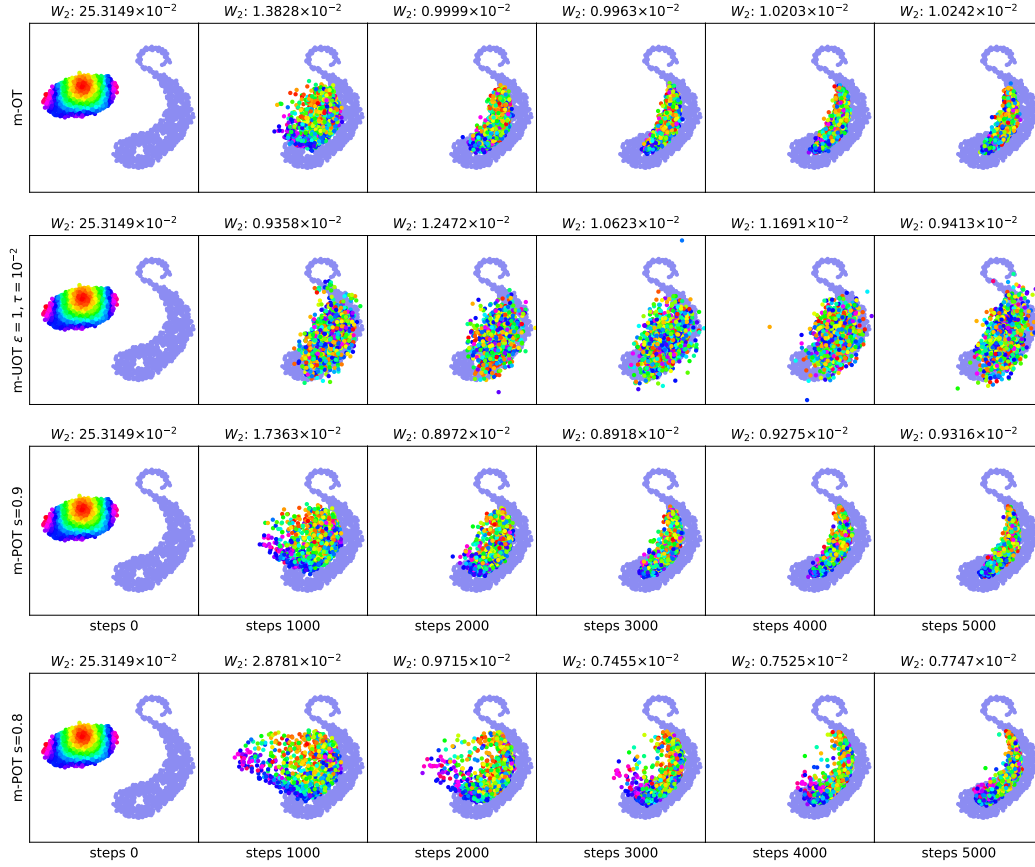


Figure 11: Gradient flows from different mini-batch methods. The Wasserstein-2 distance between two measures in each step is reported at the top of the corresponding image.

Table 13: DA performance of equivalent problems on Office-Home dataset for m-OT and m-UOT with one more sample in a mini-batch ($m + 1$).

Methods	A2C	A2P	A2R	C2A	C2P	C2R	P2A	P2C	P2R	R2A	R2C	R2P	Avg
m-OT($m=66$)	50.20	67.81	73.56	57.07	64.54	67.68	55.67	45.48	73.42	65.27	54.59	76.8	62.67
m-UOT($m=66$)	54.89	74.21	80.65	64.52	73.55	74.85	65.23	52.88	80.12	74.37	59.08	83.4	69.81
m-POT($m=65$)	54.85	73.51	80.7	65.76	73.62	76.02	64.28	53.5	79.76	74.08	59.22	83.17	69.87

Table 14: FID scores (smaller is better) of m-OT with one more samples in a mini-batch ($m + 1$) on CIFAR10 and CelebA datasets.

Dataset	CIFAR10			CelebA		
	m-OT ($m+1$)	m-POT(m)	Improvement	m-OT($m+1$)	m-POT(m)	Improvement
$m = 100$	69.06	66.77	2.29	65.55	53.76	11.79
$m = 200$	66.35	66.29	0.06	50.64	49.25	1.39

performs worse than m-POT on the other two adaptation datasets. For the deep generative model, m-POT leads to a significant drop of the FID score compared to m-OT when $m = 100$.

C.8 TIMING INFORMATION

The following timing benchmarks for deep domain adaptation and deep generative model experiments are done on an RTX 2080 Ti GPU. Figure 14 illustrates the running speed of different mini-batch methods on deep DA applications. Due to the simplicity of the problem, m-OT is the fastest method

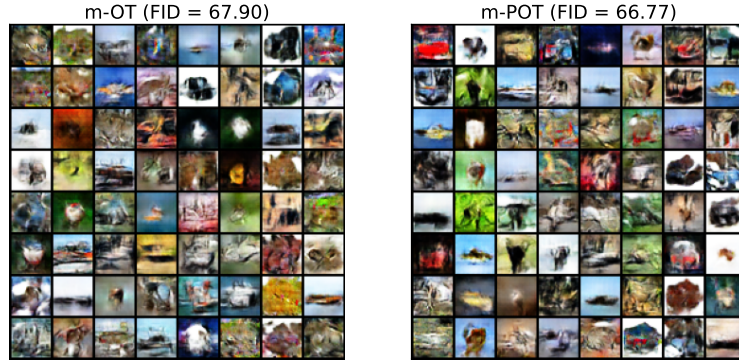


Figure 12: CIFAR10 generated images from m-OT and m-POT (best choice of s), $(k,m) = (2,100)$. The FID scores are reported in the title of images.

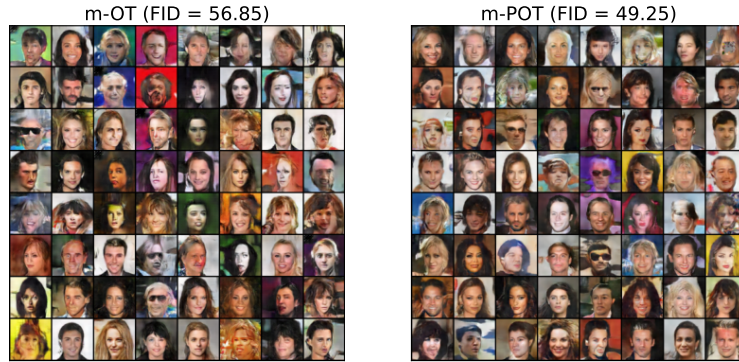


Figure 13: CelebA generated images from m-OT and m-POT (best choice of s), $(k,m) = (2,200)$. The FID scores are reported in the title of images.

Table 15: Number of iterations per second in the deep generative model experiment.

Dataset	k	m	Method	Iterations per epoch	Seconds per epoch	Iterations per second
CIFAR10	2	100	m-OT	250	25	10.00
			m-POT	250	25	10.00
		200	m-OT	125	23	5.43
			m-POT	125	30	4.17
CelebA	2	100	m-OT	814	295	2.76
			m-POT	814	295	2.76
		200	m-OT	407	277	1.47
			m-POT	407	287	1.42

in all experiments. While having comparable performance, m-POT consumes less time per iteration than m-UOT. An interesting result is that the running speed of m-POT decreases as the fraction of masses s increases. The computational speed of the deep generative model is illustrated in Table 15. Both m-OT and m-POT share the same speed when the mini-batch size m is set to 100. For $m = 200$, a training epoch of m-OT consumes 7 and 10 seconds fewer than that of m-POT on CIFAR10 and CelebA datasets, respectively. Different from the deep domain adaptation experiments, the speed does not change when varying the fraction of masses s .

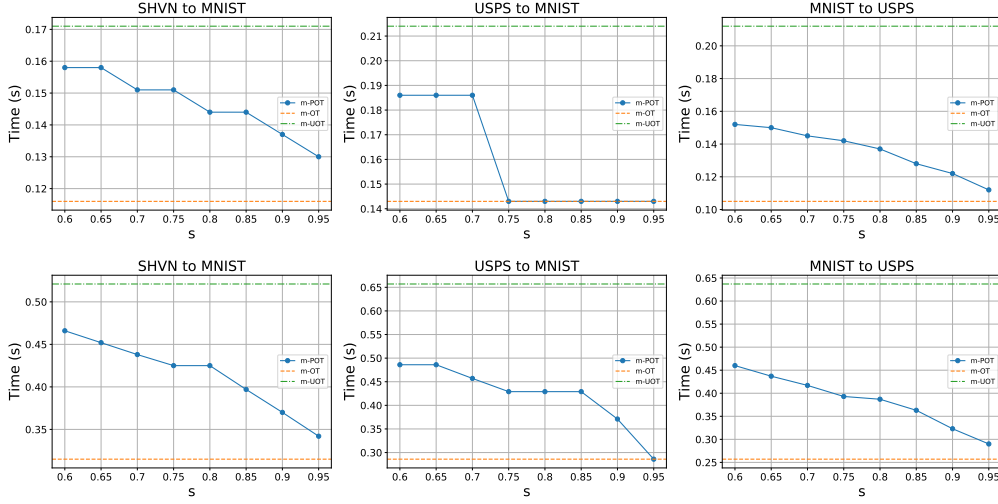


Figure 14: Time taken to run one iteration in seconds in deep DA on digits datasets. All experiments is run with a mini-batch size $m = 500$. k is set to 1 in the first row and 2 in the second row.

D EXPERIMENTAL SETTINGS

In this section, we provide datasets, architectures, and training procedure details for deep domain adaptation and deep generative model. Notice that we have used the best values of τ and ϵ that are reported in the original paper of m-UOT (Fattras et al., 2021a) for computing m-UOT in each experiment.

D.1 DEEP DOMAIN ADAPTATION

Datasets: We start with digits datasets. Following the evaluation protocol of DeepJDOT (Damodaran et al., 2018) we experiment on three adaptation scenarios: SHVN to MNIST ($S \rightarrow M$), USPS to MNIST ($U \rightarrow M$), MNIST to USPS ($M \rightarrow U$). The USPS dataset (Hull, 1994) consists of 7,291 training and 2,007 testing images, each one is a 16×16 grayscale handwritten image. The MNIST dataset (LeCun et al., 1998) contains 60,000 training and 10,000 testing grayscale images of size 28×28 . The SVHN dataset (Netzer et al., 2011) contains house numbers extracted from Google Street View images. This dataset has 73,212 training images, and 26,032 testing RGB images of size 32×32 . We evaluate all methods in 12 possible adaptation scenarios. Finally, VisDA-2017 (Peng et al., 2017) is a large-scale dataset for domain adaptation from synthetic to real images. VisDA consists of 152,397 source (synthetic) images and 55,388 target (real) images. Both source and target domains have the same 12 object categories. Following JUMBOT (Fattras et al., 2021a), we evaluate all methods on the validation set.

Parameter settings for digits datasets: We optimize using Adam with the initial learning rate $\eta_0 \in \{0.0002, 0.0003, 0.0004\}$. We train all algorithms with a batch size of 500 during 100 epochs. The hyperparameters in equation (10) follow the settings in m-UOT: $\alpha = 0.1, \lambda_t = 0.1$. For computing UOT, we also set τ to 1 and ϵ to 0.1. For computing m-POT, we select the best value of $s \in \{0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$.

Parameter settings for Office-Home dataset: We train all algorithms with a batch size of 65 during 10000 iterations. Following m-UOT, the hyperparameters for computing the cost matrix is as follows $\alpha = 0.01, \lambda_t = 0.5, \tau = 0.5, \epsilon = 0.01$. The value of s is chosen from the set $\{0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$.

Parameter settings for VisDA dataset: All methods are trained using a batch size of 72 during 10000 iterations. The coefficients in the cost formula is as follows $\alpha = 0.005, \lambda_t = 1, \tau = 0.3, \epsilon = 0.01$. The choice set of s for m-POT is $\{0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$.

Neural network architectures: Similar to DeepJDOT and JUMBOT, we use CNN for our generator and 1 FC layer for our classifier. For Office-Home and VisDA, our generator is a ResNet50 pre-trained on ImageNet except for the last FC layer, which is our classifier.

Generator architecture was used for SVHN dataset:

$$z \in \mathbb{R}^{32 \times 32 \times 3} \rightarrow \text{Conv}_{32} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{Conv}_{32} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{MaxPool2D} \rightarrow \text{Conv}_{64} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{Conv}_{64} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{MaxPool2D} \rightarrow \text{Conv}_{128} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{Conv}_{128} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{MaxPool2D} \rightarrow \text{Sigmoid} \rightarrow \text{FC}_{128}$$

Generator architecture was used for USPS dataset:

$$z \in \mathbb{R}^{28 \times 28 \times 3} \rightarrow \text{Conv}_{32} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{MaxPool2D} \rightarrow \text{Conv}_{64} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{Conv}_{128} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{MaxPool2D} \rightarrow \text{Sigmoid} \rightarrow \text{FC}_{128}$$

Classifier architecture was used for both SVHN and USPS datasets:

$$z \in \mathbb{R}^{128} \rightarrow \text{FC}_{10}$$

Classifier architecture was used for Office-Home dataset:

$$z \in \mathbb{R}^{512} \rightarrow \text{FC}_{65}$$

Classifier architecture was used for VisDA datasets:

$$z \in \mathbb{R}^{512} \rightarrow \text{FC}_{12}$$

Training details: Similar to both DeepJDOT and JUMBOT, we stratify the data loaders so that each class has the same number of samples in the mini-batches. For digits datasets, we also train our neural network on the source domain during 10 epochs before applying our method. For Office-home and VisDA datasets, because the classifiers are trained from scratch, their learning rates are set to be 10 times that of the generator. We optimize the models using an SGD optimizer with momentum = 0.9 and weight decay = 0.0005. We schedule the learning rate with the same strategy used in (Ganin et al., 2016). The learning rate at iteration p is $\eta_p = \frac{\eta_0}{(1+\mu q)^\nu}$, where q is the training progress linearly changing from 0 to 1, $\eta_0 = 0.01$, $\mu = 10$, $\nu = 0.75$.

Source code: Our source coded is based on <https://github.com/kilianFattras/JUMBOT>.

D.2 DEEP GENERATIVE MODEL

Datasets: We train generators on CIFAR10 (Krizhevsky et al., 2009) and CelebA (Liu et al., 2015) datasets. The CIFAR10 dataset contains 10 classes, with 50,000 training and 10,000 testing color images of size 32×32 . CelebA is a large-scale face attributes dataset with more than 200K celebrity images.

FID scores: We use 11000 samples from the generative model and all images from the training set to compute the FID score by the official code from authors in Heusel et al. (2017).

Parameter settings: We chose a learning rate equal to 0.0005. When the mini-batch size m is 100, the number of epochs is set to 200 for CIFAR10 dataset and 100 for CelebA dataset. When increasing the batch size to 200, we double the number of epochs for having the same number of gradient steps.

Neural network architectures: We used CNNs for both generators and discriminators on the CelebA and CIFAR10 datasets.

Generator architecture was used for CelebA:

$$z \in \mathbb{R}^{32} \rightarrow \text{TransposeConv}_{512} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{TransposeConv}_{256} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{TransposeConv}_{128} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{TransposeConv}_{64} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{TransposeConv}_3 \rightarrow \text{Tanh}$$

Discriminator architecture was used for CelebA:

$$x \in \mathbb{R}^{64 \times 64 \times 3} \rightarrow \text{Conv}_{64} \rightarrow \text{LeakyReLU}_{0.2} \rightarrow \text{Conv}_{128} \rightarrow \text{BatchNorm} \rightarrow \text{LeakyReLU}_{0.2} \rightarrow \text{Conv}_{256} \rightarrow \text{BatchNorm} \rightarrow \text{LeakyReLU}_{0.2} \rightarrow \text{Conv}_{512} \rightarrow \text{BatchNorm} \rightarrow \text{Tanh}$$

Generator architecture was used for CIFAR10:

$$z \in \mathbb{R}^{32} \rightarrow \text{TransposeConv}_{256} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{TransposeConv}_{128} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{TransposeConv}_{64} \rightarrow \text{BatchNorm} \rightarrow \text{ReLU} \rightarrow \text{TransposeConv}_3 \rightarrow \text{Tanh}$$

Discriminator architecture was used for CIFAR10:

$$x \in \mathbb{R}^{32 \times 32 \times 3} \rightarrow \text{Conv}_{64} \rightarrow \text{LeakyReLU}_{0.2} \rightarrow \text{Conv}_{128} \rightarrow \text{BatchNorm} \rightarrow \text{LeakyReLU}_{0.2} \rightarrow \text{Conv}_{256} \rightarrow \text{Tanh}$$

D.3 COMPUTATIONAL INFRASTRUCTURE

Deep generative models and deep domain adaptation experiments are done on a RTX 2080 Ti GPU. Color transfer and gradient flow applications are conducted on a MacBook Pro 11 inch M1.