
Consistent Optimal Transport with Empirical Conditional Measures

Piyushi Manupriya¹
IIT Hyderabad, INDIA

Rachit Keerti Das²
Microsoft, INDIA

Sayantana Biswas²
Amazon, INDIA

J. Saketha Nath
IIT Hyderabad, INDIA

Abstract

Given samples from two joint distributions, we consider the problem of Optimal Transportation (OT) between them when conditioned on a common variable. We focus on the general setting where the conditioned variable may be continuous, and the marginals of this variable in the two joint distributions may not be the same. In such settings, standard OT variants cannot be employed, and novel estimation techniques are necessary. Since the main challenge is that the conditional distributions are not explicitly available, the key idea in our OT formulation is to employ kernelized-least-squares terms computed over the joint samples, which implicitly match the transport plan's marginals with the empirical conditionals. Under mild conditions, we prove that our estimated transport plans, as a function of the conditioned variable, are asymptotically optimal. For finite samples, we show that the deviation in terms of our regularized objective is bounded by $O(1/m^{1/4})$, where m is the number of samples. We also discuss how the conditional transport plan could be modelled using explicit probabilistic models as well as using implicit generative ones. We empirically verify the consistency of our estimator on synthetic datasets, where the optimal plan is analytically known. When employed in applications like prompt learning for few-shot classification and conditional-generation in the context of predicting cell responses to treatment, our methodology improves upon state-of-the-art methods.

1 INTRODUCTION

Optimal Transport (OT) [Kantorovich, 1942] serves as a powerful tool for comparing distributions. OT has been instrumental in diverse ML applications [Peyré and Cuturi, 2019, Liu et al., 2020, Fatras et al., 2021, Cao et al., 2022, Chen et al., 2023] that involve matching distributions. The need for comparing conditional distributions also frequently arises in machine learning. For instance, in the supervised learning of (probabilistic) discriminative models, one needs to compare the model's label posterior with the label posterior of the training data. Learning implicit conditional-generative models is another such application. Typically, the observed input covariates in these applications are continuous rather than discrete. Consequently, one may only assume access to samples from the input-label joint distribution rather than having multiple samples for a given input. It is well known that estimating conditionals is a significantly more challenging problem than estimating joints (e.g. refer to Section (2) in [Li et al., 2022]). Hence, it is not straightforward to apply OT between the relevant conditionals, as the conditionals are implicitly given via samples from the joint distribution. This issue becomes more pronounced when the distributions of input covariates in the two joints are not the same, e.g. in medical applications [Hahn et al., 2019] where the distributions of treated and untreated patients differ. In such cases, merely performing an OT between the joint distributions of input and label is not the same as comparing the corresponding conditionals.

In this paper, we address this challenging problem of estimating OT plan between two conditionals, say $s_{Y|X}(\cdot|x)$ and $t_{Y'|X'}(\cdot|x)$, when samples from the joint distributions, $s_{X,Y}$, $t_{X',Y'}$, are given. As motivated above, we do not restrict the conditioned variable to be discrete, nor do we assume that the marginals of the common variable, s_X and $t_{X'}$, are the same. As we discuss in our work, the key challenge in estimating OT between conditionals comes in enforcing the

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

¹Corresponding author (cs18m20p100002@iith.ac.in).

²Work majorly done while at IIT Hyderabad.

marginal constraints involving the conditionals, because the samples provided are not from the conditionals, but from the joints $s_{X,Y}$ and $t_{X',Y'}$. Our formulation employs kernelized-least-squares terms, computed over the joint samples, to address this issue. These regularizer terms implicitly match the transport plan’s marginals with the empirical conditionals. Under mild assumptions, we prove that our conditional transport plan is indeed an optimal one, asymptotically. Hence, the corresponding transport cost will match the true Wasserstein between the conditionals. For finite samples, m , we show that the deviation in our regularized objective is upper bounded by $O(1/m^{1/4})$.

Few prior works have considered special cases of this problem and have focused on learning conditional optimal transport maps [Tabak et al., 2021, Bunne et al., 2022]. To the best of our knowledge, our work is the first to formulate OT between conditionals in a general setting that also leads to provably consistent estimators for the optimal transport cost as well as the transport plan as a function of the conditioned variable’s value, x . Further, instead of directly modelling the transport plan, $\pi_{Y,Y'|X}$, we instead propose modelling it’s factors: $\pi_{Y|Y',X}$, $\pi_{Y|X}$. This gives a three-fold advantage: (i) The models for the factors are much simpler than for the joint (ii) when dealing with discriminative/conditional-generative models we can directly choose $\pi_{Y|X}(\cdot|x)$ as the discriminative model being learnt. (ii) When implicit generative models are used for the factors, $\pi_{Y'|Y,X}(\cdot|y, x)$ can be readily be used for inference in applications like cell population dynamics (e.g., section 5.2).

We empirically show the utility of our approach in the conditional generative task for modelling cell population dynamics, where we consistently outperform the baselines. Furthermore, we pose the task of learning prompts for few-shot classification as a conditional optimal transport problem. We argue that this is advantageous than posing it as a classical optimal transport problem, which is the approach existing works employ. We test this novel approach on the benchmark EuroSAT [Helber et al., 2019] dataset and show improvements over [Chen et al., 2023], a state-of-the-art prompt learning method.

In Table 1, we highlight some of the key features of COT, comparing it with the related works. Our main contributions are summarized below.

Contributions

- We propose novel estimators for optimal transport between conditionals in a general setting where the conditioned variable may be continuous, and its marginals in the two joint distributions may

differ.

- We prove the consistency of the proposed estimators. To the best of our knowledge, we are the first to present a consistent estimator for conditional optimal transport in the general setting.
- While recent approaches model the optimal transport map [Tabak et al., 2021], [Bunne et al., 2022], we model the transport plan, which enables more general inferences.
- We empirically verify the correctness of the proposed estimator on synthetic datasets. We further evaluate the proposed approach on downstream applications of conditional generation for modelling cell population dynamics and prompt learning for few-shot classification, showing its utility over some of the state-of-the-art baselines.

2 PRELIMINARIES

Let \mathcal{X}, \mathcal{Y} be two sets (domains) that form compact Hausdorff spaces. Let $\mathcal{P}(\mathcal{X})$ be the set of all probabilistic measures over \mathcal{X} .

Optimal Transport (OT) Given a cost function, $c : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}$, OT compares two measures $s, t \in \mathcal{P}(\mathcal{Y})$ by finding a plan to transport mass from one to the other, that incurs the least expected cost. More formally, Kantorovich’s OT formulation [Kantorovich, 1942] is given by:

$$W_c(s, t) \equiv \min_{\pi \in \mathcal{P}(\mathcal{Y} \times \mathcal{Y})} \int c \, d\pi, \text{ s.t. } \pi_1 = s, \pi_2 = t, \quad (1)$$

where π_1, π_2 are the marginals of π . A valid cost metric over $\mathcal{Y} \times \mathcal{Y}$ defines the 1-Wasserstein metric, $W_c(s, t)$, over distributions $s, t \in \mathcal{P}(\mathcal{Y})$. The cost metric is referred to as the ground metric.

Maximum Mean Discrepancy (MMD)

Given a characteristic kernel function [Sriperumbudur et al., 2011], $k : \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}$, MMD defines a metric over probability measures given by: $\text{MMD}^2(s, t) \equiv \mathbb{E}_{X \sim s, X' \sim s}[k(X, X')] + \mathbb{E}_{Y \sim t, Y' \sim t}[k(Y, Y')] - 2\mathbb{E}_{X \sim s, Y \sim t}[k(X, Y)]$. With \mathcal{H}_k as the RKHS associated with the characteristic kernel k , the dual norm definition of MMD is given by $\text{MMD}(s, t) = \max_{f \in \mathcal{H}_k; \|f\| \leq 1} \mathbb{E}_s[f(X)] - \mathbb{E}_t[f(Y)]$.

3 RELATED WORK

Few prior works have attempted to solve the conditional OT problem in some special cases, which we discuss below. [Frogner et al., 2015] presents an estimator for the case when the marginals, s_X and $t_{X'}$, are

Table 1: Summary of related works and the proposed COT method.

	[Tabak et al., 2021]	[Luo and Ren, 2021]	[Bunne et al., 2022]	COT
Consistent estimator	N/A	N/A	N/A	✓
Models OT plan with flexibility of implicit modelling	✗	✗	✗	✓
Flexibility with the ground cost	✓	✗	✗	✓
Allows single sample per conditioned variable	✓	✓	✗	✓

the same and y takes discrete values. Their estimator does not generalize to the case where y is continuous. Further, they solve individual OT problems at each x rather than modelling the transport map/plan as a function of x . [Luo and Ren, 2021] characterizes the conditional distribution discrepancy using the Conditional Kernel Bures (CKB). With the assumption that the kernel embeddings for the source and target are jointly Gaussian, CKB defines a metric between conditionals. [Luo and Ren, 2021] does not discuss any (sufficient) conditions for this assumption to hold. Moreover, CKB only estimates the discrepancy between the two conditionals, and it is unclear how to retrieve an optimal transport plan/map with CKB, limiting its applications. [Bunne et al., 2022] considers special applications where multiple samples from $s_{Y|X}(\cdot|x)$, $t_{Y'|X'}(\cdot|x)$ are available at each x . They learn a transport map as a function of x by solving standard OT problems between $s_{Y|X}(\cdot|x)$, $t_{Y'|X'}(\cdot|x)$ individually for each sample x . Also, their approach additionally assumes the ground cost is squared Euclidean. In contrast, we neither assume access to multiple samples from $s_{Y|X}(\cdot|x)$, $t_{Y'|X'}(\cdot|x)$ at each x nor make restrictive assumptions on the ground cost. Further, we estimate the transport plan rather than the transport map. The work closest to ours is [Tabak et al., 2021]. However, there are critical differences between the two approaches, which we highlight below. [Tabak et al., 2021] formulates a min-max adversarial formulation with a KL divergence-based regularization to learn a transport map. Such adversarial formulations are often unstable, and [Tabak et al., 2021] does not present any convergence results. Their empirical evaluation is also limited to small-scale qualitative experiments. Moreover, unlike the estimation bounds we prove, [Tabak et al., 2021] does not discuss any learning theory bounds or consistency results. It is expected that such bounds would be cursed with dimensions [Séjourné et al., 2023b, Séjourné et al., 2023a]. Additionally, the proposed formulation allows us to learn transport plans using implicit models (§ 4.2.2). Such an approach may not be possible with KL-regularized formulation in [Tabak et al., 2021] due to non-overlapping support of the distributions. Owing to these differences, our proposed method is more

widely applicable.

4 PROBLEM FORMULATION

This section formally defines the Conditional Optimal Transport (COT) problem and presents a consistent estimator for it in the general setting. We begin by recalling the definition of OT between two given measures $s_{Y|X}(\cdot|x)$ and $t_{Y'|X'}(\cdot|x)$ for a given x . $W_c(s_{Y|X}(\cdot|x), t_{Y'|X'}(\cdot|x))$ is defined as follows.

$$\begin{aligned} \min_{\pi_{Y,Y'|X}(\cdot,\cdot|x) \in \mathcal{P}(\mathcal{Y} \times \mathcal{Y})} \int_{\mathcal{Y} \times \mathcal{Y}} c \, d\pi_{Y,Y'}(\cdot,\cdot|x), \quad (2) \\ \text{s.t. } \pi_{Y|X}(\cdot|x) = s_{Y|X}(\cdot|x), \pi_{Y'|X'}(\cdot|x) = t_{Y'|X'}(\cdot|x), \end{aligned}$$

where $\pi_{Y|X}(\cdot|x)$ and $\pi_{Y'|X'}(\cdot|x)$ denotes the marginals of $\pi_{Y,Y'|X}(\cdot,\cdot|x)$. If the cost is a valid metric, then $W_c(s_{Y|X}(\cdot|x), t_{Y'|X'}(\cdot|x))$ is nothing but the Wasserstein distance between $s_{Y|X}(\cdot|x)$ and $t_{Y'|X'}(\cdot|x)$. While $W_c(s_{Y|X}(\cdot|x), t_{Y'|X'}(\cdot|x))$ helps comparing/transporting measures given a specific $x \in \mathcal{X}$, in typical learning applications, one needs a comparison in an expected sense rather than at a specific $x \in \mathcal{X}$. Accordingly, we consider $\mathbb{E}_{X'' \sim a} [W_c(s_{Y|X}(\cdot|X''), t_{Y'|X'}(\cdot|X''))]$, where a is a given auxiliary measure:

$$\begin{aligned} \int_{\mathcal{X}} \min_{\substack{\pi_{Y,Y'|X}(\cdot,\cdot|x) \in \mathcal{P}(\mathcal{Y} \times \mathcal{Y}) \\ \forall x \in \mathcal{X}}} \int_{\mathcal{Y} \times \mathcal{Y}} c \, d\pi_{Y,Y'|X}(\cdot,\cdot|x) \, da(x), \\ \text{s.t. } \pi_{Y|X}(\cdot|x) = s_{Y|X}(\cdot|x), \\ \pi_{Y'|X'}(\cdot|x) = t_{Y'|X'}(\cdot|x) \quad \forall x \in \mathcal{X} \\ \equiv \min_{\pi_{Y,Y'|X}: \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y} \times \mathcal{Y})} \int_{\mathcal{X}} \int_{\mathcal{Y} \times \mathcal{Y}} c \, d\pi_{Y,Y'|X}(\cdot,\cdot|x) \, da(x), \\ \text{s.t. } \pi_{Y|X}(\cdot|x) = s_{Y|X}(\cdot|x), \\ \pi_{Y'|X'}(\cdot|x) = t_{Y'|X'}(\cdot|x) \quad \forall x \in \mathcal{X}. \quad (3) \end{aligned}$$

In the special case where the auxiliary measure, a , is degenerate, (3) gives back (2). Henceforth, we analyze the proposed COT formulation defined in (3).

Now, in typical machine learning applications, the conditionals are not explicitly given, and only samples from the joints are available. Estimation of COT

from samples seems challenging because the problem of estimating conditional densities itself has been acknowledged to be a significantly difficult one with known impossibility results (e.g., refer to Section 2 in [Li et al., 2022]). Hence, some regularity assumptions are necessary for consistent estimation. Further, even after making appropriate assumptions, the typical estimation errors are cursed with dimensions (e.g., Theorem 2.1 in [Graham et al., 2020]).

On the other hand, estimation of RKHS embeddings of conditional measures can be performed at rates $O(1/m^{1/4})$, where m is the number of samples [Song et al., 2009] [Grünewälder et al., 2012]. This motivates us to enforce the constraints in COT (3) by penalizing the distance between their RKHS embeddings. More specifically, we exploit the equivalence: $\pi_{Y|X}(\cdot|x) = s_{Y|X}(\cdot|x) \forall x \in \mathcal{X} \iff \int_{\mathcal{X}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), s_{Y|X}(\cdot|x)) ds_X(x) = 0$. This is true because MMD is a valid metric and we assume $s_X(x) > 0, t_{X'}(x) > 0 \forall x \in \mathcal{X}$. Using this, COT (3) can be relaxed as:

$$\begin{aligned} & \min_{\pi_{Y,Y'|X} : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y} \times \mathcal{Y})} \int_{\mathcal{X}} \int_{\mathcal{Y} \times \mathcal{Y}} c \, d\pi_{Y,Y'|X}(\cdot, \cdot|x) \, da(x), \\ & + \lambda_1 \int_{\mathcal{X}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), s_{Y|X}(\cdot|x)) \, ds_X(x) \\ & + \lambda_2 \int_{\mathcal{X}} \text{MMD}^2(\pi_{Y'|X}(\cdot|x), t_{Y'|X}(\cdot|x)) \, dt_{X'}(x), \end{aligned} \quad (4)$$

where $\lambda_1, \lambda_2 > 0$ are regularization hyperparameters. Note that (4) is exactly the same as (3) if $\lambda_1, \lambda_2 \rightarrow \infty$.

Now, we use a standard result, $\mathbb{E}[\|G - h(X)\|^2] = \mathbb{E}[\|G - \mathbb{E}[G|X]\|^2] + \mathbb{E}[\|\mathbb{E}[G|X] - h(X)\|^2]$ with G taken as the kernel mean embedding of δ_Y and $h(X)$ taken as the kernel mean embedding of $\pi_{Y|X}(\cdot|X)$ [Muandet et al., 2017]. This gives us $\int_{\mathcal{X} \times \mathcal{Y}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), \delta_y) \, ds_{X,Y}(x, y) = \int_{\mathcal{X}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), s_{Y|X}(\cdot|x)) \, ds_X(x) + v(s)$, where $v(s) \geq 0$. Here, ϕ is the feature map corresponding to the kernel defining the MMD. This leads to the following formulation:

$$\begin{aligned} & \min_{\pi_{Y,Y'|X} : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y} \times \mathcal{Y})} \int_{\mathcal{X}} \int_{\mathcal{Y} \times \mathcal{Y}} c \, d\pi_{Y,Y'|X}(\cdot, \cdot|x) \, da(x), \\ & + \lambda_1 \int_{\mathcal{X} \times \mathcal{Y}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), \delta_y) \, ds_{X,Y}(x, y) \\ & + \lambda_2 \int_{\mathcal{X} \times \mathcal{Y}} \text{MMD}^2(\pi_{Y'|X}(\cdot|x), \delta_y) \, dt_{X',Y'}(x, y). \end{aligned} \quad (5)$$

Since $v(s), v(t)$ are independent of π , the solutions of (5) are exactly the same as those of COT (3) as $\lambda_1, \lambda_2 \rightarrow \infty$. The advantage of this reformulation is

that it can be efficiently estimated using samples from the joints, as we detail below.

4.1 Sample-Based Estimation

In our set-up, in order to solve (5) and perform estimation, we are only provided with samples $\mathcal{D}_m^s = \{(x_1, y_1), \dots, (x_m, y_m)\}$ and $\mathcal{D}_m^t = \{(x'_1, y'_1), \dots, (x'_m, y'_m)\}$ from $s_{X,Y}$ and $t_{X',Y'}$, respectively. Hence, we employ a sample-based estimator for the regularizer terms: $\int_{\mathcal{X} \times \mathcal{Y}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), \delta_y) \, ds_{X,Y}(x, y) \approx \frac{1}{m} \sum_{i=1}^m \text{MMD}^2(\pi_{Y|X}(\cdot|x_i), \delta_{y_i})$. The following lemma shows that this regularizer estimator is statistically consistent.

Lemma 1. *Assuming k is a normalized characteristic kernel, with probability at least $1 - \delta$, we have*

$$\begin{aligned} & \left| \int_{\mathcal{X} \times \mathcal{Y}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), \delta_y) \, ds_{X,Y}(x, y) \right. \\ & \left. - \frac{1}{m} \sum_{i=1}^m \text{MMD}^2(\pi_{Y|X}(\cdot|x_i), \delta_{y_i}) \right| \leq 2\sqrt{\frac{2}{m} \log\left(\frac{2}{\delta}\right)}. \end{aligned}$$

Using this result for the regularization terms, (5) can in-turn be estimated as:

$$\begin{aligned} & \min_{\pi_{Y,Y'|X} : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y} \times \mathcal{Y})} \int_{\mathcal{X}} \int_{\mathcal{Y} \times \mathcal{Y}} c \, d\pi_{Y,Y'|X}(\cdot, \cdot|x) \, da(x) \\ & + \lambda_1 \frac{1}{m} \sum_{i=1}^m \text{MMD}^2(\pi_{Y|X}(\cdot|x_i), \delta_{y_i}) \\ & + \lambda_2 \frac{1}{m} \sum_{i=1}^m \text{MMD}^2(\pi_{Y'|X}(\cdot|x'_i), \delta_{y'_i}). \end{aligned} \quad (6)$$

We choose not to estimate the first term with empirical average as a is a known distribution. In the following theorem, we prove the consistency of our COT estimator.

Theorem 1. *Let Π be a given model for the conditional transport plans, $\pi_{Y,Y'|X} : \mathcal{X} \mapsto \mathcal{P}(\mathcal{Y} \times \mathcal{Y})$. Assume $\lambda_1 = \lambda_2 = \lambda$. Let $\hat{\pi}_m, \pi^*$ denote optimal solutions over the restricted model Π corresponding to (6), (5) respectively. Let $\mathcal{U}_m[\pi], \mathcal{U}[\pi]$ denote the objectives as a function of $\pi \in \Pi$ in (6), (5) respectively. Then, we prove the following:*

1. *With probability at least $1 - \delta$, $\mathcal{U}[\hat{\pi}_m] - \mathcal{U}[\pi^*] \leq 2\lambda_1 \mathcal{R}_m(\Pi) + 2\lambda_2 \mathcal{R}'_m(\Pi) + 6(\lambda_1 + \lambda_2) \sqrt{\frac{2}{m} \log \frac{3}{\delta}}$, where the Rademacher based complexity term, $\mathcal{R}_m(\Pi)$, is defined as: $\frac{1}{m} \mathbb{E} \left[\max_{\pi \in \Pi} \sum_{i=1}^m \epsilon_i \text{MMD}^2(\pi_{Y|X}(\cdot|X_i), \delta_{Y_i}) \right]$; (X_i, Y_i) are IID samples from $s_{X,Y}$ and ϵ_i denotes the Rademacher random variable. $\mathcal{R}'_m(\Pi)$, is analogously defined as: $\frac{1}{m} \mathbb{E} \left[\max_{\pi \in \Pi} \sum_{i=1}^m \epsilon_i \text{MMD}^2(\pi_{Y'|X}(\cdot|X'_i), \delta_{Y'_i}) \right]$, where (X'_i, Y'_i) are IID samples from $t_{X',Y'}$ and*

ϵ_i denotes the Rademacher random variable. Recall that $\pi_{Y|X}(\cdot|x)$ and $\pi_{Y'|X}(\cdot|x)$ denote the marginals of $\pi_{Y,Y'|X}(\cdot, \cdot|x)$.

- In the special case Π is a neural network based conditional generative model, the kernel employed is universal, normalized, and non-expansive [Waarde and Sepulchre, 2022], and $\lambda = O(m^{1/4})$, with high probability we have that $\mathcal{U}[\hat{\pi}_m] - \mathcal{U}[\pi^*] \leq O(1/m^{1/4})$. More importantly, when $m \rightarrow \infty$, $\hat{\pi}_m$ is an optimal solution to the original COT problem (3) whenever Π is rich enough such that $\exists \pi^* \in \Pi \ni \pi_{Y|X}^*(\cdot|x) = s_{Y|X}(\cdot|x)$ and $\pi_{Y'|X}^*(\cdot|x) = t_{Y'|X}(\cdot|x) \forall x \in \mathcal{X}$.

The proof is presented in Supplementary § (S1.2). The conditions for consistency are indeed mild because (i) neural conditional generators are known to be universal (Lemma 2.1 in [Liu et al., 2021], [Kidger and Lyons, 2020]) (ii) the popularly used Gaussian kernel is indeed universal, normalized, and non-expansive (for a large range of hyperparameters). The proof for the first part of the theorem is an adaptation of classical uniform convergence based arguments; however, further bounding the complexity terms in the case of neural conditional generative models is novel and we derive this using vector contraction inequalities along with various properties of the kernel.

4.2 Modelling the Transport Plan

We now provide details of modelling the transport plan function, i.e., choices for Π , from a pragmatic perspective. Firstly, we model the transport plan $\pi_{Y,Y'|X}(y, y'|x)$ by modelling its factors: $\pi_{Y'|Y,X}(y'|y, x)$ and $\pi_{Y|X}(y|x)$. Since the factors can be modelled using simpler models, this brings us computational benefits, among other advantages that we discuss. Secondly, employing COT with such a factorization enables us to directly choose $\pi_{Y|X}(\cdot|x)$ as the label posterior of the model to be learnt in discriminative modelling applications. Moreover, the other factor $\pi_{Y'|Y,X}(\cdot|y, x)$ can be readily used for inference (see § 5.1.2, § 5.2).

4.2.1 Transport Plan with Explicit Models

Here, we discuss our modelling choice with explicit probabilistic models when $\mathcal{Y} = \{l_1, \dots, l_n\}$ is a finite set. Accordingly, we model the factors $\pi_{Y'|Y,X}(y'|y, x)$, $\pi_{Y|X}(y|x)$ with fixed-architecture neural networks, parameterized by ψ and ϕ respectively, with the output layer as softmax over $|\mathcal{Y}|$ labels.

The COT estimator 6 in this case simplifies as:

$$\begin{aligned} & \min_{\psi, \theta} \int_{\mathcal{X}} \sum_{i=1}^n \sum_{j=1}^n c(l_i, l_j) \pi_{\psi}(l_i|l_j, x) \pi_{\theta}(l_j|x) da(x) \\ & + \lambda_1 \frac{1}{m} \sum_{i=1}^m \text{MMD}^2 \left(\sum_{j=1}^n \pi_{\psi}(\cdot|l_j, x_i) \pi_{\theta}(l_j|x_i), \delta_{y_i} \right) \\ & + \lambda_2 \frac{1}{m} \sum_{i=1}^m \text{MMD}^2 \left(\pi_{\theta}(\cdot|x'_i), \delta_{y'_i} \right), \end{aligned} \quad (7)$$

where ψ, θ are the network parameters we wish to learn. In discriminative learning applications, the factor $\pi_{\theta}(\cdot|x)$ can be readily used as a probabilistic classifier (e.g., section 5.3).

4.2.2 Transport Plan with Implicit Models

As mentioned earlier, in applications such as § 5.1.2, § 5.2, it is required to generate samples from $\pi_{Y'|Y,X}(\cdot|y, x)$ for inference. In such applications, one would prefer modelling these transport plan factors using implicit generative models.

Since the MMD metric, unlike KL-divergence, can be employed to compare measures with non-overlapping support, implicit generative models can be readily employed for modelling our transport plan. More specifically, we model the factors $\pi_{Y'|Y,X}(y'|y, x)$, $\pi_{Y|X}(y|x)$ with fixed-architecture generative neural networks, π_{ψ} and π_{θ} , respectively. We use $\eta, \eta' \sim \mathcal{N}(0, 1)$ to denote the noise random variables. The π_{θ} network takes as input x and random η' to produce (random) y , to be distributed as $\pi_{Y|X}(\cdot|x)$. Like-wise, the π_{ψ} network takes as input y, x and random η to produce (random) y' , to be distributed as $\pi_{Y'|Y,X}(\cdot|y, x)$. We denote the outputs of π_{θ} by $y(x, \eta'_i; \theta)$ $i = 1, \dots, m$ (i.e., samples from $\pi_{Y|X}(\cdot|x)$). And, we denote outputs of π_{ψ} by $y(x, \eta_i, \eta'_i; \theta, \psi)$ $i = 1, \dots, m$, when inputs are $y(x, \eta'_i; \theta), x, \eta_i$. We illustrate the overall model in figure 1. Then, the COT estimator, with implicit modelling, reads as:

$$\begin{aligned} & \min_{\theta, \psi} \int_{\mathcal{X}} \frac{1}{m} \sum_{i=1}^m c(y(x, \eta'_i; \theta), y(x, \eta_i, \eta'_i; \theta, \psi)) da(x) \\ & + \lambda_1 \frac{1}{m} \sum_{i=1}^m \text{MMD}^2 \left(\frac{1}{m} \sum_{j=1}^m \delta_{y(x, \eta_j, \eta'_j; \theta, \psi)}, \delta_{y_i} \right) \\ & + \lambda_2 \frac{1}{m} \sum_{i=1}^m \text{MMD}^2 \left(\frac{1}{m} \sum_{j=1}^m \delta_{y(x'_i, \eta'_j; \theta)}, \delta_{y'_i} \right). \end{aligned} \quad (8)$$

We note that solving the COT problem, then readily provides us with the factors $\pi_{Y'|Y,X}(y'|y, x)$ and $\pi_{Y|X}(y|x)$, which can be used for inference purposes. This is in contrast to a typical implicit modelling approach, where one would require samples of (x, y, y') for learning such a model. The unavailability of such triplets (as in § 5.2) often limits such typical approaches. However, as we can see, COT now allows us to learn such a model without the availability of such triplets, only using samples from $s_{X,Y}$ and $t_{X',Y'}$. This clearly shows the benefits of the proposed approach.




Figure 1: Illustration of the proposed factorization and implicit modelling for learning the transport plan $\pi_{Y, Y'|X}(y, y'|x)$ through the factors $\pi_\theta(y|x)\pi_\psi(y'|y, x)$, parameterized by fixed-architecture neural networks 4.2.2. $\eta, \eta' \sim \mathcal{N}(0, 1)$ denotes the noise input to the implicit models.

5 EXPERIMENTS

In this section, we showcase the utility of the proposed estimator 5 in various applications. We choose the auxiliary distribution a as the empirical distribution over the training covariates and use $\lambda_1 = \lambda_2 = \lambda$ in all our experiments. More experimental details and results are in Supplementary § (S2.2).¹

5.1 Verifying Correctness of Estimator

We empirically verify the correctness of the proposed estimator in synthetically constructed settings where the closed-form solutions are known.

5.1.1 Convergence to the True Wasserstein

We learn the implicit networks with the proposed COT loss 8, keeping λ high enough. With the learnt networks, we draw samples $y(x, \eta'_i; \theta) \sim \pi_\theta(\cdot|x)$ and $y(x, \eta_i, \eta'_i; \theta, \psi) \sim \pi_\psi(\cdot|y(x, \eta'_i; \theta), x)$, for $i = 1, \dots, m$, and compute the transport cost (first term in 8) and compare it with $W_c(s_{Y|X}(\cdot|x), t_{Y'|X'}(\cdot|x))$. In order to verify that our estimate converges to the true Wasserstein, we consider a case where the analytical solution for the Wasserstein distance W_c is known and compare it with our estimate.

Experimental Setup We consider two distributions $y \sim \mathcal{N}(4(x - 0.5), 1)$ and $y' \sim \mathcal{N}(-2(x' - 0.5), 8x' + 1)$ where $x \sim \beta(2, 4)$ and $x' \sim \beta(4, 2)$ generate m samples from each them. The true Wasserstein distance between them at x turns out to be $(6(x - 0.5))^2 + (\sqrt{8x + 1} - 1)^2$ (see Equation (2.39) in [Peyré and Cuturi, 2019]), which we compare against. We use the RBF kernel and squared Euclidean distance as our ground cost. The factors $\pi_\theta(\cdot|x)$ and $\pi_\psi(\cdot|y, x)$ are modelled using two 2-layer MLP neural networks.

¹The code for reproducing our experiments is publicly available at <https://github.com/atmlr-lab/COT>.




Figure 2: As $m \in \{100, 200, 400, 800\}$ increases from left to right, we plot the true Wasserstein distance in red and mark the means (in orange) and medians (in green) of the distances estimated using [Tabak et al., 2021] and the proposed COT estimator. The statistics are obtained from runs over multiple seeds. The corresponding MSEs are $\{245.530, 290.458, 89.715, 27.687\}$ and $\{22.711, 6.725, 8.052, 1.580\}$ respectively. It can be seen that the proposed COT objective converges to the true Wasserstein faster than [Tabak et al., 2021].

Results Figure 2 shows the convergence to the true Wasserstein as m increases. The variance of the estimated values and the MSEs decrease as the number of samples increases. The quadratic nature of the function is also captured with our estimator.

5.1.2 Convergence to the True Barycenter

For further verification of our estimator, we show that the barycenter estimated using our transport plan and the true barycenter converge in Wasserstein distance.

Experimental Setup Two independent Gaussian distributions are taken $y \sim \mathcal{N}(2(x - 0.5), 1)$ and $y' \sim \mathcal{N}(-4(x' - 0.5), 4)$ where $x \sim \beta(2, 4)$ and $x' \sim \beta(4, 2)$. The analytical solution of the barycenter is calculated as $y_c \sim \mathcal{N}(-x + 0.5, 2.5)$ [Peyré and Cuturi, 2019]. Recall that the barycenter can also be computed using the optimal transport map (Remark 3.1 in [Gordaliza et al., 2019]) using the expression: $B_x = \rho S_x + (1 - \rho)T_x$, where $\rho \in [0, 1]$ and B_x, S_x, T_x denote the random variables corresponding to the barycenter, source measure and the transported sample, conditioned on x , respectively. Accordingly, samples from the barycenter, B_{x_i} , are obtained using: $\rho y_i + (1 - \rho)y$, where $y \sim \pi_\psi(\cdot|y_i, x_i)$.




Figure 3: Barycenters shown on varying $\rho \in [0, 1]$ with colors interpolated between red and blue. Left: Conditional barycenter learnt by the proposed COT method. Right: Analytical barycenter.




Figure 4: For increasing values of m , we show box plots of the Wasserstein distance between the learnt barycenter, B_x , and the analytical barycenter. The corresponding MSEs are $\{22.399, 3.408, 3.964, 2.534, 1.687\}$ for [Tabak et al., 2021] and $\{4.441, 0.654, 0.353, 0.099, 0.058\}$ for the proposed COT estimator. It can be seen that the proposed COT-based barycenter converges to the true solution faster than [Tabak et al., 2021].

Results For evaluation, we generate 500 samples from our transport plan based barycenter and the true barycenter. We use kernel density estimation to plot the barycenters. Figures 3 and 4 show that the proposed estimate of barycenter closely resembles the analytical barycenter and converges on increasing m .

5.2 Cell Population Dynamics

The study of single-cell molecular responses to treatment drugs is a major problem in biology. Existing single-cell-sequencing methods allow one to observe gene expressions of the cells, but do so by destroying them. As a result, one ends up with cells from control (unperturbed) and target (perturbed) distributions without a correspondence between them. Optimal transport has emerged as a natural method [Bunne et al., 2021] to obtain a mapping between the source and the target cells, which can then be used for predictions on unseen cells. As the drug dosage is highly correlated with the predicted cell populations, [Bunne et al., 2022] learns such optimal trans-

port maps conditioned on the drug dosage. We apply the proposed COT formulation to generate samples from the distributions over perturbed cells conditioned on the drug dosage given to an unperturbed cell.

Dataset We consider the dataset used by [Bunne et al., 2022] and [Bunne et al., 2021] corresponding to the cancer drug Givinostat applied at different dosage levels, $\{x_1 = 10nM, x_2 = 100nM, x_3 = 1000nM, x_4 = 10000nM\}$. At each dosage level, x_i , samples of perturbed cells are given: y_{i1}, \dots, y_{im_i} . The total perturbed cells are 3541. Samples of unperturbed cells are also provided: $y'_1, \dots, y'_m, m = 17, 565$. Each of these cells is described by gene-expression levels of $n = 1000$ highly variable genes, i.e., $y_{ij}, y'_i \in \mathbb{R}^{1000}$. Following [Bunne et al., 2022], the representations of cells are brought down to 50 dimensions with PCA.

COT-Based Generative Modelling Our goal is to perform OT between the distribution of the unperturbed cells and the distribution of the perturbed cell conditioned on the drug dosage. As the representations of the cells lie in $\mathcal{Y} = \mathbb{R}^{50}$, we choose implicit modelling (§ 4.2.2) for learning the conditional transport plans. The factor π_θ is taken as the empirical distribution over the unperturbed cells. With this notation, our COT estimator, (8), simplifies as follows.

$$\min_{\psi} \frac{1}{4} \sum_{q=1}^4 \frac{1}{m} \sum_{i=1}^m c(y'_i, y(x_q, \eta_i; \psi)) + \lambda_1 \frac{1}{4} \sum_{i=1}^4 \text{MMD}^2 \left(\frac{1}{m} \sum_{j=1}^m \delta_{y(x_i, \eta_j; \psi)}, \frac{1}{m_i} \sum_{j=1}^{m_i} \delta_{y_{ij}} \right),$$

where $y(x, \eta_i; \psi)$ $i = 1, \dots, m$ are samples from the network $\pi_\psi(\cdot | y'_i, x)$.

Experimental Setup Similar to [Bunne et al., 2022], we take the cost function, c , as squared Euclidean. For the MMD regularization, we use the characteristic inverse multi-quadratic (IMQ) kernel.

Results Following [Bunne et al., 2022], we evaluate the performance of COT by comparing samples from the predicted and ground truth perturbed distributions. We report the l_2 norm between the Perturbation Signatures [Stathias et al., 2018], for 50 marker genes for various dosage levels. We also report the MMD distances between the predicted and target distributions on various dosage levels. The distances are reported for in-sample settings, i.e. the dosage levels are seen during training. We compare our performance to the reproduced CellOT [Bunne et al., 2021] and CondOT [Bunne et al., 2022] baselines.

We summarize our results in Tables 2 and 3. We observe that COT consistently outperforms state-of-the-

Table 2: l_2 (PS) distances (lower is better) between predicted and ground truth distributions

Dosage	CellOT	CondOT	Proposed
10nM	1.2282	0.3789	0.3046
100nM	1.2708	0.2515	0.2421
1000nM	0.8653	0.7290	0.3647
10000nM	4.9035	0.3819	0.2607
Average	2.067	0.4353	0.2930

Table 3: MMD distances (lower is better) between predicted and ground truth distributions

Dosage	CellOT	CondOT	Proposed
10nM	0.01811	0.00654	0.00577
100nM	0.0170	0.00555	0.00464
1000nM	0.0154	0.01290	0.00647
10000nM	0.1602	0.01034	0.00840
Average	0.0526	0.00883	0.00632

art baselines CondOT [Bunne et al., 2022] and CellOT [Bunne et al., 2021] in terms of l_2 (PS) as well as the MMD distances.

5.3 Prompt Learning

In order to show the versatility of our framework, we adapt our estimator for learning prompts for large-scale vision-language models and evaluate the performance in a limited supervision setting.

The success of vision-language models in open-world visual understanding has motivated efforts which aim to learn prompts [Zhou et al., 2022a, Zhang et al., 2022, Zhou et al., 2022b, Chen et al., 2023] to adapt the knowledge from pre-trained models like CLIP [Radford et al., 2021] for downstream tasks since it is infeasible to fine-tune such models due to a large number of parameters. Typically, these approaches rely on learning class-specific prompts for each category to better adapt the vision-language model for downstream tasks without the need for fine-tuning. A recent approach, PLOT [Chen et al., 2023], achieved state-of-the-art results by incorporating an OT-based loss between distributions over the set of local visual features and the set of textual prompt features, each of 1024 dimensions, to learn the downstream classifier. For each image, PLOT computes an OT-based loss between $M(49)$ visual features of the image and $N(4)$ textual prompt features per class.

As prompts are shared across images of a class [Chen et al., 2023], learning optimal transport plans conditioned on class-level information is expected to improve the downstream performance




Figure 5: We pose learning prompts in few-shot classification as the conditional optimal transport problem. The figure shows our neural network diagram for learning conditional optimal transport plans.

Table 4: AUC on test data (higher is better). We compare the performance of COT against other OT-based losses ϵ -OT ([Frogner et al., 2015]) and CKB ([Luo and Ren, 2021]).

Dataset	ϵ -OT	CKB	Proposed
MNIST	0.89	0.99	0.99
CIFAR10	0.66	0.73	0.79
Animals with Attribute	0.68	0.64	0.86

compared to solving an OT problem separately for each (image, class) pair. Hence, we pose this prompt learning task as a COT problem, where the conditional transport plans are modelled explicitly (§ 4.2.1).

Validating the Proposed Explicit Modelling

Before working on the challenging few-shot classification task, we evaluate the proposed explicit modelling-based COT estimator on a simpler multi-class classification task. Let the discriminative model to be learnt be f_θ . The idea is to match this conditional to that in the training data using COT. We choose the transport plan factor $\pi_\theta \equiv f_\theta$ and a as the marginal of input covariates in the training data, simplifying our COT estimator, (7), as:

$$\min_{\psi, \theta} \frac{1}{m} \sum_{q=1}^m \sum_{i=1, j=1}^{i=n, j=n} c(l_i, l_j) \pi_\psi(l_i | l_j, x_q) f_\theta(l_j | x_q) + \lambda_1 \frac{1}{m} \sum_{i=1}^m \text{MMD}^2(\sum_{j=1}^n \pi_\psi(\cdot | l_j, x_i) f_\theta(l_j | x_i), \delta_{y_i}), \quad (9)$$

where ψ, θ are the network parameters we wish to learn. Table 4 validates the performance with the proposed explicit modelling.

COT Formulation for Prompt Learning

We learn an explicit model $\pi_{\psi_r}(\cdot | l_{jqr}, x_{qr})$ over the N textual prompt features l_{1r}, \dots, l_{Nr} for each class. Here, x_{qr} is the q^{th} image from class r and l_{jqr} is the j^{th} visual feature for image x_{qr} . Following PLOT, the distribution over image features given an image is considered uniform and, hence, not modelled as the other

Table 5: Prompt Learning experiment: Average accuracy (higher is better) on EuroSAT dataset. The class-level context brought by the proposed COT method allows it to outperform the state-of-the-art PLOT baseline, especially in the challenging case of lesser K .

	CoOp	PLOT	Proposed
$K = 1$	52.12 ± 5.46	54.05 ± 5.95	61.20 ± 3.65
$K = 2$	59.00 ± 3.48	64.21 ± 1.90	64.67 ± 2.37
$K = 4$	68.61 ± 3.54	72.36 ± 2.29	72.53 ± 2.60
$K = 8$	77.08 ± 2.42	78.15 ± 2.65	78.57 ± 2.38

factor in the transport plan. Figure (5) depicts the proposed setup. Our formulation for prompt learning for K -shot classification (only K training images per class) is as follows.

$$\min_{\psi_r} \frac{1}{K} \sum_{q=1}^K \sum_{i=1, j=1}^{i=N, j=M} c(l_{ir}, l_{jq_r}) \pi_{\psi_r}(l_{jq_r}, x_{qr}) \mathbf{v}_j + \lambda_1 \text{MMD}^2(\sum_{q=1}^K \sum_{j=1}^M \pi_{\psi_r}(\cdot | l_{jq_r}, x_{qr}) \mathbf{v}_j, \mathbf{u}). \quad (10)$$

Following the PLOT setup, we take \mathbf{v}, \mathbf{u} as uniform distributions over the $M(49)$ visual features and the $N(4)$ prompt features, respectively. As the prompts are shared across the images of a class, our MMD-regularization term matches the cumulative marginals to the distribution over prompt features.

Experimental Setup We take the same experimental setup used in CoOp[Zhou et al., 2022b] and PLOT[Chen et al., 2023] for learning prompts and only change the training loss to 10. The kernel employed is the characteristic inverse multi-quadratic, and the ground cost is the cosine cost. We follow the common training/evaluation protocol used in CoOp and PLOT and report the mean and standard deviation of the accuracies obtained with 3 seeds.

Results In Table 5, we report the accuracies on the EuroSAT benchmark dataset [Helber et al., 2019] for the number of shots K as 1, 2, 4 and 8. As the number of shots represents the number of training images per class, learning with lesser K is more difficult. The advantage of class-level context brought by the proposed COT formulation is evident in this setting.

6 CONCLUSION

Often, machine learning applications need to compare conditional distributions. Remarkably, our framework enables such a comparison solely using samples from (observational) joint distributions. To the best of our knowledge, the proposed method is the first work that consistently estimates the conditional transport plan in the general setting. The cornerstone of our work lies in the theoretical analysis of its convergence properties, demonstrating different modelling choices for

learning and empirically validating its correctness. We further showcase the utility of the proposed method in downstream applications of cell population dynamics and prompt learning for few-shot classification. A possible future work would be to extend the proposed approach of generating conditional barycenters (§ 5.1.2) to work with more than two conditionals.

Acknowledgements

The first author is supported by the Google PhD Fellowship. JSN would like to thank Fujitsu Limited, Japan, for the generous research grant. We thank Charlotte Bunne for the clarifying discussions on reproducing the CondOT method. We also thank Dr Pratik Jawanpuria, Kusampudi Venkata Datta Sri Harsha, Shivam Chandhok, Aditya Saibewar, Amit Chandhak and the anonymous reviewers who helped us improve our work. PM thanks Suvodip Dey and Sai Srinivas Kancheti for the support.

References

[Bojanowski et al., 2017] Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.

[Bunne et al., 2022] Bunne, C., Krause, A., and Cuturi, M. (2022). Supervised training of conditional monge maps. In *NeurIPS*.

[Bunne et al., 2021] Bunne, C., Stark, S. G., Gut, G., del Castillo, J. S., Lehmann, K.-V., Pelkmans, L., Krause, A., and Rättsch, G. (2021). Learning single-cell perturbation responses using neural optimal transport. *bioRxiv*.

[Bunne et al., 2023] Bunne, C., Stark, S. G., Gut, G., del Castillo, J. S., Levesque, M., Lehmann, K.-V., Pelkmans, L., Krause, A., and Ratsch, G. (2023). Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*.

[Cao et al., 2022] Cao, Z., Xu, Q., Yang, Z., He, Y., Cao, X., and Huang, Q. (2022). Otkge: Multi-modal knowledge graph embeddings via optimal transport. In *NeurIPS*.

[Chen et al., 2023] Chen, G., Yao, W., Song, X., Li, X., Rao, Y., and Zhang, K. (2023). Prompt learning with optimal transport for vision-language models. In *ICLR*.

[FAtlas et al., 2021] Atlas, K., Séjourné, T., Courty, N., and Flamary, R. (2021). Unbalanced minibatch optimal transport; applications to domain adaptation. In *ICML*.

- [Fatras et al., 2020] Fatras, K., Zine, Y., Flamary, R., Gribonval, R., and Courty, N. (2020). Learning with minibatch wasserstein: asymptotic and gradient properties. In *AISTATS*.
- [Frogner et al., 2015] Frogner, C., Zhang, C., Mobahi, H., Araya, M., and Poggio, T. A. (2015). Learning with a wasserstein loss. In *NIPS*.
- [Gordaliza et al., 2019] Gordaliza, P., Barrio, E. D., Fabrice, G., and Loubes, J.-M. (2019). Obtaining fairness using optimal transport theory. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2357–2365. PMLR.
- [Graham et al., 2020] Graham, B. S., Niu, F., and Powell, J. L. (2020). Minimax risk and uniform convergence rates for nonparametric dyadic regression. *NBER Working Paper Series*.
- [Grünewälder et al., 2012] Grünewälder, S., Lever, G., Gretton, A., Baldassarre, L., Patterson, S., and Pontil, M. (2012). Conditional mean embeddings as regressors. In *ICML*.
- [Hahn et al., 2019] Hahn, P. R., Dorie, V., and Murray, J. S. (2019). Atlantic causal inference conference (ACIC) data analysis challenge 2017.
- [Helber et al., 2019] Helber, P., Bischke, B., Dengel, A., and Borth, D. (2019). Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226.
- [Jawanpuria et al., 2021] Jawanpuria, P., Satyadev, N., and Mishra, B. (2021). Efficient robust optimal transport with application to multi-label classification. In *IEEE Conference on Decision and Control (CDC)*.
- [Kantorovich, 1942] Kantorovich, L. (1942). On the transfer of masses (in russian). *Doklady Akademii Nauk*, 37(2):227–229.
- [Kidger and Lyons, 2020] Kidger, P. and Lyons, T. (2020). Universal Approximation with Deep Narrow Networks. In *ICML*.
- [Krizhevsky et al., 2009] Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.
- [Lampert et al., 2009] Lampert, C. H., Nickisch, H., and Harmeling, S. (2009). Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*.
- [LeCun and Cortes, 2010] LeCun, Y. and Cortes, C. (2010). MNIST handwritten digit database.
- [Li et al., 2022] Li, M., Neykov, M., and Balakrishnan, S. (2022). Minimax optimal conditional density estimation under total variation smoothness. *Electronic Journal of Statistics*, 16(2):3937 – 3972.
- [Liu et al., 2021] Liu, S., Zhou, X., Jiao, Y., and Huang, J. (2021). Wasserstein generative learning of conditional distribution. *ArXiv*.
- [Liu et al., 2020] Liu, Y., Zhu, L., Yamada, M., and Yang, Y. (2020). Semantic correspondence as an optimal transport problem. In *CVPR*.
- [Luo and Ren, 2021] Luo, Y.-W. and Ren, C.-X. (2021). Conditional bures metric for domain adaptation. In *CVPR*.
- [Maurer, 2016] Maurer, A. (2016). A vector-contraction inequality for rademacher complexities. In *ALT*.
- [Muandet et al., 2017] Muandet, K., Fukumizu, K., Sriperumbudur, B., and Schölkopf, B. (2017). Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1-2):1–141.
- [Neyshabur, 2017] Neyshabur, B. (2017). Implicit regularization in deep learning.
- [Peyré and Cuturi, 2019] Peyré, G. and Cuturi, M. (2019). Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607.
- [Radford et al., 2021] Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. (2021). Learning transferable visual models from natural language supervision. In *ICML*.
- [Song et al., 2009] Song, L., Huang, J., Smola, A., and Fukumizu, K. (2009). Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *ICML*.
- [Sriperumbudur et al., 2011] Sriperumbudur, B. K., Fukumizu, K., and Lanckriet, G. R. G. (2011). Universality, characteristic kernels and RKHS embedding of measures. *Journal of Machine Learning Research*, 12:2389–2410.
- [Stathias et al., 2018] Stathias, V., Jermakowicz, A. M., Maloof, M. E., Forlin, M., Walters, W. M., Suter, R. K., Durante, M. A., Williams, S. L., Harbour, J. W., Volmar, C.-H., Lyons, N. J.,

Wahlestedt, C., Graham, R. M., Ivan, M. E., Komotar, R. J., Sarkaria, J. N., Subramanian, A., Golub, T. R., Schürer, S. C., and Ayad, N. G. (2018). Drug and disease signature integration identifies synergistic combinations in glioblastoma. *Nature Communications*, 9.

[Séjourné et al., 2023a] Séjourné, T., Bonet, C., Fatras, K., Nadjahi, K., and Courty, N. (2023a). Unbalanced optimal transport meets sliced-wasserstein.

[Séjourné et al., 2023b] Séjourné, T., Feydy, J., Vialard, F.-X., Trounev, A., and Peyré, G. (2023b). Sinkhorn divergences for unbalanced optimal transport.

[Tabak et al., 2021] Tabak, E. G., Trigila, G., and Zhao, W. (2021). Data driven conditional optimal transport. *Machine Learning*, 110(11):3135–3155.

[Waarde and Sepulchre, 2022] Waarde, H. v. and Sepulchre, R. (2022). Training lipschitz continuous operators using reproducing kernels. In *Annual Learning for Dynamics and Control Conference*.

[Wolf et al., 2018] Wolf, F. A., Angerer, P., and Theis, F. J. (2018). Scanpy: large-scale single-cell gene expression data analysis. *Genome Biology*, 19(1):15.

[Zhang et al., 2022] Zhang, R., Zhang, W., Fang, R., Gao, P., Li, K., Dai, J., Qiao, Y., and Li, H. (2022). Tip-Adapter: Training-free adaption of clip for few-shot classification. In *ECCV*.

[Zhou et al., 2022a] Zhou, K., Yang, J., Loy, C. C., and Liu, Z. (2022a). Conditional prompt learning for vision-language models. In *CVPR*.

[Zhou et al., 2022b] Zhou, K., Yang, J., Loy, C. C., and Liu, Z. (2022b). Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Consistent Optimal Transport with Empirical Conditional Measures: Supplementary Materials

In continuation to the main paper, we present theoretical proofs, more details on the experiments and some additional experimental results. Our key sections are listed as follows.

- Theoretical proofs **S1**.
- Visualizing predictions of our conditional generator **S2.1**.
- More experimental details and additional results **S2.2**.

S1 THEORETICAL PROOFS

S1.1 Proof of Lemma 1

Lemma 1. Assuming k is a normalized characteristic kernel, with probability at least $1 - \delta$, we have:

$$\left| \int_{\mathcal{X} \times \mathcal{Y}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), \delta_y) \, ds_{X,Y}(x, y) - \frac{1}{m} \sum_{i=1}^m \text{MMD}^2(\pi_{Y|X}(\cdot|x_i), \delta_{y_i}) \right| \leq 2\sqrt{\frac{2}{m} \log\left(\frac{2}{\delta}\right)}.$$

Proof. Recall that MMD is nothing but the RKHS norm-induced distance between the corresponding kernel embeddings i.e., $\text{MMD}(s, t) = \|\mu_k(s) - \mu_k(t)\|$, where $\mu_k(s) \equiv \int \phi_k(x) \, ds_X$, is the kernel mean embedding of s [Muandet et al., 2017], ϕ_k is the canonical feature map associated with the characteristic kernel k . Let \mathcal{H}_k denote the RKHS associated with the kernel k . Since our kernel is normalized we have that $\|\mu_k(b)\| \leq 1 \, \forall b \in \mathcal{P}(\mathcal{Y})$. Hence, $0 \leq \text{MMD}^2(\pi_{Y|X}(\cdot|x), s_{Y|X}(\cdot|x)) = \|\mu_k(\pi_{Y|X}(\cdot|x)) - \mu_k(s_{Y|X}(\cdot|x))\|^2 \leq \|\mu_k(\pi_{Y|X}(\cdot|x)) + \mu_k(s_{Y|X}(\cdot|x))\|^2 \leq (\|\mu_k(\pi_{Y|X}(\cdot|x))\| + \|\mu_k(s_{Y|X}(\cdot|x))\|)^2 \leq 4$, where the second last step uses the triangle inequality. From Chernoff-Hoeffding bound, we have that: with probability at least $1 - \delta$, $\left| \int_{\mathcal{X} \times \mathcal{Y}} \text{MMD}^2(\pi_{Y|X}(\cdot|x), \delta_y) \, ds_{X,Y}(x, y) - \frac{1}{m} \sum_{i=1}^m \text{MMD}^2(\pi_{Y|X}(\cdot|x_i), \delta_{y_i}) \right| \leq 2\sqrt{\frac{2}{m} \log\left(\frac{2}{\delta}\right)}$. \square

S1.2 Proof of Theorem 1

We first restate Corollary (4) from the result of vector-contraction inequality for Rademacher in [Maurer, 2016], which we later use in our proof.

Corollary (Restated from [Maurer, 2016]). Let \mathcal{H} denote a Hilbert space and let f be a class of functions $f : \mathcal{X} \mapsto \mathcal{H}$, let $h_i : \mathcal{H} \mapsto \mathbb{R}$ have Lipschitz norm L . Then

$$\mathbb{E} \sup_{f \in F} \sum_i \epsilon_i h_i(f(x_i)) \leq \sqrt{2}L \sum_{i,k} \epsilon_{i,k} f_k(x_i),$$

where ϵ_{ik} is an independent doubly indexed Rademacher sequence, and $f_k(x_i)$ is the k -th component of $f(x_i)$.

Our consistency theorem from the main paper is presented below, followed by its proof.

Proof of Theorem 1.

Proof. From the definition of $\mathcal{U}[\hat{\pi}_m]$ and $\mathcal{U}[\pi^*]$, it follows that $0 \leq \mathcal{U}[\hat{\pi}_m] - \mathcal{U}[\pi^*]$.

$$\begin{aligned} 0 \leq \mathcal{U}[\hat{\pi}_m] - \mathcal{U}[\pi^*] &= \mathcal{U}[\hat{\pi}_m] - \hat{\mathcal{U}}_m[\hat{\pi}_m] + \hat{\mathcal{U}}_m[\hat{\pi}_m] - \hat{\mathcal{U}}_m[\pi^*] + \hat{\mathcal{U}}_m[\pi^*] - \mathcal{U}[\pi^*] \\ &\leq \mathcal{U}[\hat{\pi}_m] - \hat{\mathcal{U}}_m[\hat{\pi}_m] + \hat{\mathcal{U}}_m[\pi^*] - \mathcal{U}[\pi^*] \quad (\because \hat{\pi}_m \text{ is the solution of 6}) \\ &\leq \max_{\pi \in \Pi} (\mathcal{U}[\pi] - \hat{\mathcal{U}}_m[\pi]) + \hat{\mathcal{U}}_m[\pi^*] - \mathcal{U}[\pi^*] \end{aligned} \quad (\text{S11})$$

We now separately upper bound the two terms in S11 : $(\hat{\mathcal{U}}_m[\pi^*] - \mathcal{U}[\pi^*])$ and $\max_{\pi \in \Pi} (\mathcal{U}[\pi] - \hat{\mathcal{U}}_m[\pi])$. From Lemma 1, with probability at least $1 - \delta$,

$$\hat{\mathcal{U}}_m[\pi^*] - \mathcal{U}[\pi^*] \leq 2(\lambda_1 + \lambda_2) \sqrt{\frac{2}{m} \log \frac{2}{\delta}} \quad (\text{S12})$$

We now turn to the second term. We show that $\max_{\pi \in \Pi} \mathcal{U}[\pi] - \hat{\mathcal{U}}_m[\pi]$ satisfies the bounded difference property. Let Z_i denote the random variable (X_i, Y_i) . Let $Z = \{Z_1, \dots, Z_i, \dots, Z_m\}$ be a set of independent random variables. Consider another such set that differs only at the i^{th} position : $Z' = \{Z_1, \dots, Z_i', \dots, Z_m\}$. Let $\hat{\mathcal{U}}_m[\pi]$ and $\hat{\mathcal{U}}_m'[\pi]$ be the corresponding objectives in 6.

$$\begin{aligned} &\left| \max_{\pi \in \Pi} (\mathcal{U}[\pi] - \hat{\mathcal{U}}_m[\pi]) - \max_{\pi \in \Pi} (\mathcal{U}[\pi] - \hat{\mathcal{U}}_m'[\pi]) \right| \\ &\leq \left| \max_{\pi \in \Pi} -\hat{\mathcal{U}}_m[\pi] + \hat{\mathcal{U}}_m'[\pi] \right| \\ &\leq \frac{\lambda_1}{m} \left| \max_{\pi \in \Pi} \text{MMD}^2(\pi_{Y|X}(\cdot|x_i), \delta_{y_i}) - \text{MMD}^2(\pi_{Y|X}(\cdot|x_i'), \delta_{y_i'}) \right| \\ &+ \frac{\lambda_2}{m} \left| \max_{\pi \in \Pi} \text{MMD}^2(\pi_{Y'|X}(\cdot|x_i), \delta_{y_i}) - \text{MMD}^2(\pi_{Y'|X}(\cdot|x_i'), \delta_{y_i'}) \right| \quad (\text{Using triangle inequality}) \\ &\leq \frac{8(\lambda_1 + \lambda_2)}{m}, \end{aligned} \quad (\text{S13})$$

where for the last step, we use that, with a normalized kernel, $(\text{MMD}(\pi_{\bar{Y}}(\cdot|x_i), \delta_{y_i}) + \text{MMD}(\pi_{\bar{Y}}(\cdot|x_i'), \delta_{y_i'})) \leq 4$ and $(\text{MMD}(\pi_{\bar{Y}}(\cdot|x_i), \delta_{y_i}) - \text{MMD}(\pi_{\bar{Y}}(\cdot|x_i'), \delta_{y_i'})) \leq 2$ for $\bar{Y} \in \{Y, Y'\}$.

Using the above in McDiarmid's inequality,

$$\max_{\pi \in \Pi} \mathcal{U}[\pi] - \hat{\mathcal{U}}_m[\pi] \leq \mathbb{E} \left[\max_{\pi \in \Pi} \mathcal{U}[\pi] - \hat{\mathcal{U}}_m[\pi] \right] + 4(\lambda_1 + \lambda_2) \sqrt{\frac{2}{m} \log \frac{1}{\delta}}. \quad (\text{S14})$$

Let $Z_i \equiv (X_i, Y_i) \sim s_{X,Y}$ and $Z = \{Z_1, \dots, Z_m\}$. Let $Z_i' \equiv (X_i', Y_i') \sim t_{X,Y}$ and $Z' = \{Z_1', \dots, Z_m'\}$. Let $(\epsilon_i)_{i \in \{1, \dots, m\}}$ be IID Rademacher random variables. We now follow the standard symmetrization trick and introduce the Rademacher random variables to get the following.

$$\mathbb{E}[\max_{\pi \in \Pi} \mathcal{U}[\pi] - \hat{\mathcal{U}}_m[\pi]] \leq 2\lambda_1 \underbrace{\frac{1}{m} \mathbb{E}_{Z, \epsilon} [\max_{\pi \in \Pi} \sum_{i=1}^m \epsilon_i \|\mu_k(\pi_{Y|X}(\cdot|x_i)) - \phi(Y_i)\|^2]}_{\mathcal{R}_m(\Pi)} + 2\lambda_2 \underbrace{\frac{1}{m} \mathbb{E}_{Z', \epsilon} [\max_{\pi \in \Pi} \sum_{i=1}^m \epsilon_i \|\mu_k(\pi_{Y'|X}(\cdot|x_i)) - \phi(Y_i)\|^2]}_{\mathcal{R}'_m(\Pi)}. \quad (\text{S15})$$

Recall that $\mu_k(s)$ is the kernel mean embedding of the measure s . Hence, using S12, S14 and S15, we prove that with probability at least $1 - \delta$,

$$\mathcal{U}[\hat{\pi}_m] - \mathcal{U}[\pi^*] \leq 2\lambda_1 \mathcal{R}_m(\Pi) + 2\lambda_2 \mathcal{R}'_m(\Pi) + 6(\lambda_1 + \lambda_2) \sqrt{\frac{2}{m} \log \frac{3}{\delta}}. \quad (\text{S16})$$

Bounding Rademacher in the Special Case: We now upper-bound $\mathcal{R}_m(\Pi)$ for the special case where $\pi(\cdot|x)$ is implicitly defined using neural conditional generative models. More specifically, let d be the dimensionality

of \mathcal{Y} and let $g_{\mathbf{w}}(x, N) \in \mathbb{R}^{2d} \sim \pi(\cdot|x)$, where $g_{\mathbf{w}}$ is a neural network function parameterized by \mathbf{w} , N denotes the noise random variable. We make a mild assumption on the weights of the neural network to be bounded. The first d outputs, denoted by $g_{\mathbf{w},1}(x, N)$ will be distributed as $\pi_{Y|X}(\cdot|x)$ and the last d outputs, denoted by $g_{\mathbf{w},2}(x, N)$ will be distributed as $\pi_{Y'|X}(\cdot|x)$. Let $\zeta_i(\pi_{Y|X}) \equiv \|\mu_k(\pi_{Y|X}(\cdot|x_i)) - \phi(y_i)\|^2$. We now compute the Lipschitz constant for ζ_i , used in our bound next.

$$\begin{aligned}
 \zeta_i(\pi_{Y|X}) - \zeta_i(\pi'_{Y|X}) &\leq 4 \left(\|\mu_k(\pi_{Y|X}(\cdot|x_i)) - \phi(y_i)\| - \|\mu_k(\pi'_{Y|X}(\cdot|x_i)) - \phi(y_i)\| \right) \quad (\text{With a normalized kernel}) \\
 &\leq 4 \|\mu_k(\pi_{Y|X}(\cdot|x_i)) - \mu_k(\pi'_{Y|X}(\cdot|x_i))\| \quad (\text{Using triangle inequality}) \\
 &= 4 \|\mathbb{E}[\phi(g_{\mathbf{w},1}(x_i, N))] - \mathbb{E}[\phi(g_{\mathbf{w}',1}(x_i, N))]\| \\
 &\leq 4\mathbb{E}[\|\phi(g_{\mathbf{w},1}(x_i, N)) - \phi(g_{\mathbf{w}',1}(x_i, N))\|] \quad \because (\text{Jensen's inequality}) \\
 &\leq 4\mathbb{E}[\|g_{\mathbf{w},1}(x_i, N) - g_{\mathbf{w}',1}(x_i, N)\|] \quad \because (\text{non-expansive kernel}) \\
 &\leq 4 \|\|g_{\mathbf{w},1}(x_i, n_{i,1}) - g_{\mathbf{w}',1}(x_i, n_{i,1})\| \| \quad \because n_{i,j} \equiv \arg \max_n \|\|g_{\mathbf{w},j}(x_i, n) - g_{\mathbf{w}',j}(x_i, n)\|\|.
 \end{aligned} \tag{S17}$$

We next use a vector-contraction inequality for Rademacher given in Corollary (4) from [Maurer, 2016]. This gives $\mathcal{R}_m(\Pi) \leq \frac{4\sqrt{2}}{m} \mathbb{E}_{Z,\epsilon} \max_{\mathbf{w}} \sum_{i=1}^m \sum_{j=1}^d r_{ij} g_{\mathbf{w},1}^j(x_i, n_{i,1})$ and $\mathcal{R}'_m(\Pi) \leq \frac{4\sqrt{2}}{m} \mathbb{E}_{Z',\epsilon} \max_{\mathbf{w}} \sum_{i=1}^m \sum_{j=1}^d r_{ij} g_{\mathbf{w},2}^j(x_i, n_{i,2})$. Here, $g_{\mathbf{w},1}^j, g_{\mathbf{w},2}^j$ denote the j^{th} output in the first and the second blocks; r_{ij} denotes an independent doubly indexed Rademacher variable. Thus, we have upper bounded the complexity of Π in terms of that of the neural networks.

Now, applying standard bounds (e.g. refer to §5 in [Neysshabur, 2017]) on Rademacher complexity of neural networks, we obtain $\mathcal{R}_m(\Pi) \leq O(1/\sqrt{m})$ and $\mathcal{R}'_m(\Pi) \leq O(1/\sqrt{m})$. If λ_1, λ_2 are chosen to be $O(m^{1/4})$, then from (S16), we have: $\mathcal{U}[\hat{\pi}_m] - \mathcal{U}[\pi^*] \leq O(1/m^{1/4})$. When $m \rightarrow \infty$, this shows that $\hat{\pi}_m$ is also an optimal solution of (6), in which case it is also an optimal solution of the original COT problem (when restricted to Π) because $\lambda \rightarrow \infty$ too. \square

S2 MORE ON EXPERIMENTS

This section contains more experimental details along with some additional results.

S2.1 Visualizing Predictions of the Conditional Generator

We visualize the predictions learnt by the implicit conditional generator trained with the COT loss 8 and the alternate formulation S18 described below. The COT formulation 4 employs a clever choice of MMD regularization over the conditionals, which is then computed using the samples from the joints 5. One may think of alternatively employing an MMD regularization over joints as follows.

$$\begin{aligned}
 \min_{\pi_{Y,Y'|X}: \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y} \times \mathcal{Y})} &\int_{\mathcal{X}} \int_{\mathcal{Y} \times \mathcal{Y}} c \, d\pi_{Y,Y'|X}(\cdot, \cdot|x) da(x) + \lambda_1 \text{MMD}^2(\pi_{Y|X}(\cdot|x)s(x), s(x, y)) \\
 &+ \lambda_2 \text{MMD}^2(\pi_{Y'|X}(\cdot|x)t(y), t(x, y)).
 \end{aligned} \tag{S18}$$

We argue that this choice is sub-optimal. We first note that as we only have samples from the joints and not the marginal distributions (s_X and t_X), matching conditionals through the above formulation is not straightforward. Computing the above formulation also incurs more memory because for computing the Gram matrix over (x, y) , we need to keep Gram matrices over the samples of x, y separately. Further, in this case, each of the Gram matrices is larger than the ones needed with the proposed formulation 5. We compared the performances of the two formulations in a toy regression case and found the proposed COT formulation better.

The training algorithm for learning with the proposed COT loss is presented in Algorithm S1. The per-epoch computational complexity is $O(m^2)$, where m . We fix λ to 500, noise dimension to 10. We use Adam optimizer with a learning rate of $5e-3$ and train for 1000 epochs. We use squared Euclidean distance and RBF kernel. Figure 6 shows we obtain a good fit for $\sigma^2 = 10, 100$.

In Table 6, we also show the per-epoch computation time taken (on an RTX 4090 GPU) by the COT loss as a function of the size of the minibatch, which shows the computational efficiency of the COT loss. On the other

Algorithm S1 Algorithm for learning with implicit models for a simple regression case.

Require: Implicit neural networks $\pi_\theta: \mathcal{X} \mapsto \mathcal{Y}$ and $\pi_\psi: \mathcal{X}, \mathcal{Y} \mapsto \mathcal{Y}$, training samples $(x_i, y_i)_{i=1}^m$, noise distribution η , cost function $c: \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{R}^+$, kernel, λ .

- 1: **while** not converged or max epochs not reached **do**
- 2: Sample $z_i \sim \eta \forall i \in [m]$.
- 3: $y_i(x_i; \theta) = \pi_\theta(\cdot | x_i, z_i) \forall i \in [m]$.
- 4: Sample $z'_i \sim \eta \forall i \in [m]$.
- 5: $y_i(x_i; \theta, \psi) = \pi_\psi(\cdot | y_i(x_i; \theta), x_i, z'_i); \forall i \in [m]$.
- 6: Compute the COT loss (Simplified case of Equation 8)

$$\min_{\theta, \psi} \frac{1}{m} \sum_{i=1}^m c(y_i(x_i; \theta), y_i(x_i; \theta, \psi)) + \lambda \frac{1}{m} \sum_{i=1}^m \text{MMD}^2 \left(\frac{1}{m} \sum_{j=1}^m \delta_{y_j(x_i; \theta, \psi)}, \delta_{y_i} \right).$$

- 7: Update θ, ψ using gradient descent.
- 8: **end while**

hand, the computation time for the alternate formulation discussed in S18 (with MMD regularization over joints) is 0.245 ± 0.0012 s with minibatch-size 16 and resulted in the out-of-memory error for higher batchsizes.




Figure 7: The objective over the training epochs curve.

B	Time (s)
16	0.228 ± 0.0003
64	0.228 ± 0.0004
512	0.229 ± 0.0008
1024	0.231 ± 0.0016

Table 6: Time (in s) for COT loss 8 computation shown for increasing minibatch size (B). The computation time reported is based on 3 independent runs on the toy regression dataset.

S2.2 More Experimental Details

We provide more details for the experiments shown in § 5 of the main paper, along with some additional results.

Verifying the Correctness of Estimator We use Adam optimizer and jointly optimize π_θ and π_ψ . We choose λ from the set $\{1, 200, 500, 800, 1000\}$ and σ^2 used in the RBF kernel from the set $\{1e-2, 1e-1, 1, 10\}$. We found λ as 1000 and σ^2 as 1 to perform the best.

In Figure 8, we also show the OT plans. We draw 500 samples from the implicit maps learnt with the COT loss 8 and use kernel density estimation (KDE) to plot the distributions.

Cell Population Dynamics Dataset: We use the preprocessed dataset provided by [Bunne et al., 2023]. The dataset is publicly available for download using the following link

<https://polybox.ethz.ch/index.php/s/RAykIMfD10qCJaM>.

From this dataset, we extracted unperturbed cells and cells treated with Givinostat. This led to a total of 17565 control cells and a total of 3541 cells treated with Givinostat. We take the same data splits as in [Bunne et al., 2023].

More on evaluation: Following [Bunne et al., 2022], we use scanpy’s [Wolf et al., 2018] `rank_genes_groups` function for ranking and obtaining 50 marker genes for the drug, in this case Givinostat. The perturbed cells are grouped by drug, and the ranking is computed by keeping the unperturbed (i.e. control) cells as reference. We fix the architecture of our implicit model (ψ) as a 5-layer MLP and train it for 1000 epochs. Similar to

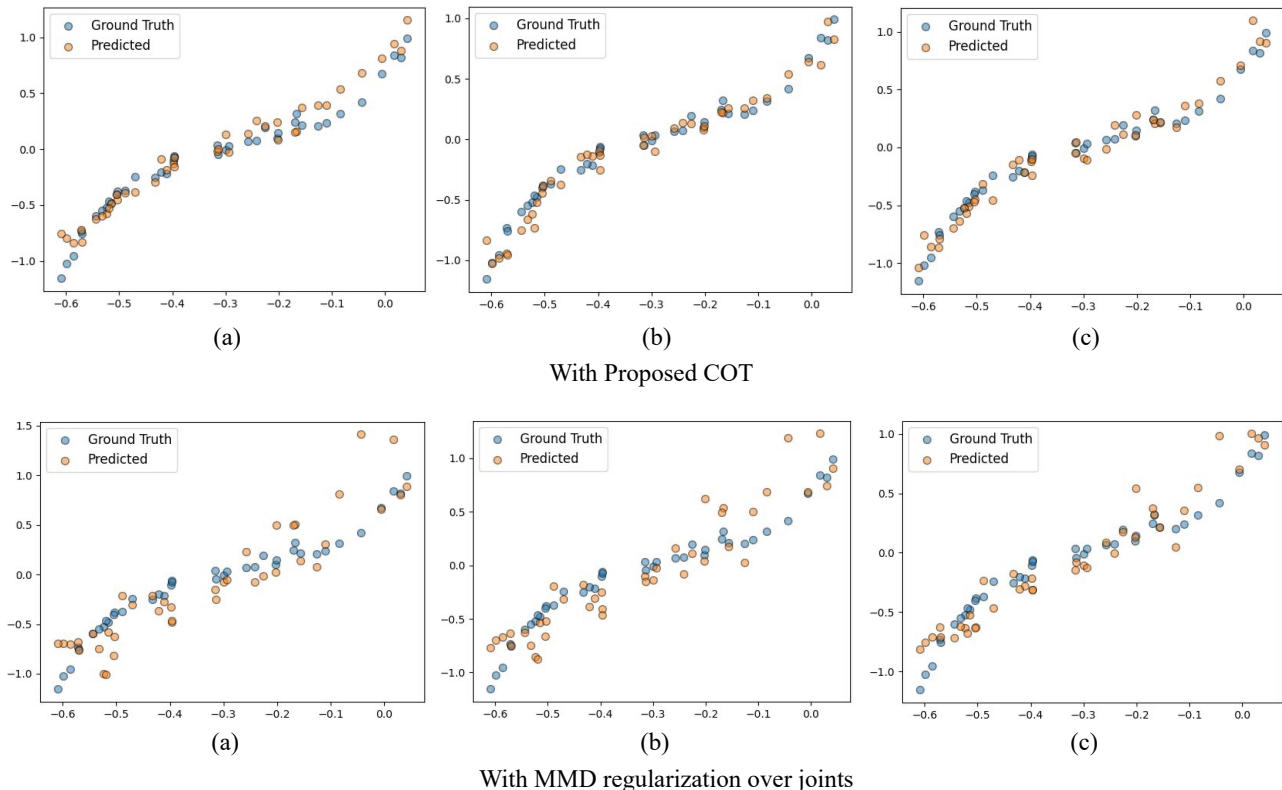


Figure 6: Predictions of the implicit conditional generator trained with the COT loss 8 and the alternate formulation S18 (with MMD regularization over joints). The plots show the effect of different σ^2 hyperparameters used in the RBF kernel as 1, 10 and 100, respectively. We quantitatively evaluate the methods using Explained Variance (between $-\infty$ and 1; higher is better). With the proposed COT loss, the explained variance scores are 0.94, 0.94 and 0.95, respectively. With the alternate formulation S18, the explained variance scores are 0.63, 0.73 and 0.85. This shows the superiority of the proposed COT formulation 8.

[Bunne et al., 2022], we train on the 50-dimensional representation after applying PCA on the 1000-dimensional original representation. It is worth noting that training our MLP models is much stabler than the Partial Input Convex Neural Networks (PICNN) used in [Bunne et al., 2022], which needs carefully chosen initialization schemes. Following the evaluation scheme in [Bunne et al., 2022], we get back to the original 1000 dimensions, and then 50 marker genes are computed for the evaluation metrics.

Following the in-sample experiment done in [Bunne et al., 2022], we tune our hyperparameters on the training data split. Based on the scale of terms in the COT objective, we chose λ from the set $\{400, 2000, 10000\}$ and found $\lambda = 400$ to be the optimal choice. For the IMQ kernel, we chose the hyperparameter from the set $\{1, 10, 50, 100\}$ and found 100 to be the optimal choice. Since we model the transport plan and not the transport map, the following procedure is followed for inference. We generate one sample corresponding to each pair of (source sample, condition) through our implicit model, and measure the required metrics on the generated distributions. This procedure is repeated $n = 50$ times, and the average metric is reported.

Following [Bunne et al., 2022], we quantitatively evaluate our performance using the MMD distance and the l_2 distance between the perturbation signatures, $l_2(\text{PS})$ metric. Let μ be the set of observed unperturbed cell population, ν be the set of the observed perturbed cell population (of size m_1), and ν' be the set of predicted perturbed state of population μ (of size m_2). The perturbation signature $\text{PS}(\nu, \mu)$ is then defined as $\frac{1}{m_1} \sum_{y_i \in \nu} y_i - \frac{1}{m_2} \sum_{y_i \in \mu} y_i'$. The $l_2(\text{PS})$ metric is the l_2 distance between $\text{PS}(\nu, \mu)$ and $\text{PS}(\nu', \mu)$. Following [Bunne et al., 2022], we report MMD (§ 2) with RBF kernel averaged over the kernel widths: $\{2, 1, 0.5, 0.1, 0.01, 0.005\}$.

Additional Results: In addition to the results reported in Tables 2 and 3 where the marker genes are computed




Figure 8: The OT plans computed with the COT formulation 8 for the case of source and target as the conditional Gaussian distributions. For each value of the conditioned variable, we show the corresponding source, target and the obtained OT plan in a given color.

Table 7: Insample setting: l_2 (PS) distances (lower is better) between predicted and ground truth distributions where the marker genes are computed at a per-dose level.

Dosage	CellOT	CondOT	Proposed
10nM	0.7164	0.4718	0.3682
100nM	0.5198	0.3267	0.3051
1000nM	0.7075	0.6982	0.3917
10000nM	4.8131	0.3457	0.2488
Average	1.6892	0.4606	0.3284

on a per-drug level, in Tables 7 and 8, we show results where marker genes are computed on a per-dosage level. Further, we present results for the out-of-sample setting, i.e., the dosage levels we predict are not seen during training. In Tables 9 and 10, we show the results when marker genes are computed on a per-drug level and in Tables 11 and 12. we show the results when marker genes are computed on a per-dose level. In Figures 9 and 10, we also show how closely the marginals of the proposed conditional optimal transport plan match the target distribution. The plots for COT correspond to the generated distribution having the median value for the metrics among all the (n=50) generated distributions.

Classification We consider the task of multi-class classification and experiment on three benchmark datasets MNIST [LeCun and Cortes, 2010], CIFAR-10 [Krizhevsky et al., 2009] and Animals with Attribute (AWA) [Lampert et al., 2009]. Following the popular approaches of minibatch OT [Fratras et al., 2020, Fratas et al., 2021], we perform a minibatch training. We use the implementation of [Frogner et al., 2015] open-sourced by [Jawanpuria et al., 2021]. We maintain the same experimental setup used in [Jawanpuria et al., 2021]. The classifier is a single-layer neural network with Softmax activation trained for 200 epochs. We use the cost function, c , between labels as the squared l_2 distance between the fastText embeddings [Bojanowski et al., 2017]




Figure 9: Marginals for selected genes ‘ENSG00000165092.12’, ‘ENSG00000175175.5’, ‘ENSG00000173727.12’, where the dosage is 100nM, in the insample setting.

Table 8: In-sample setting: MMD distances (lower is better) between predicted and ground truth distributions where the marker genes are computed at a per-dose level.

Dosage	CellOT	CondOT	Proposed
10nM	0.0089	0.0064	0.00549
100nM	0.0069	0.0054	0.00494
1000nM	0.0117	0.01038	0.00586
10000nM	0.16940	0.01051	0.01011
Average	0.04922	0.00817	0.00660

Table 9: Out-of-sample setting: l_2 (PS) distances (lower is better) between predicted and ground truth distributions where the marker genes are computed at a per-drug level.

Dosage	CellOT	CondOT	Proposed
10nM	2.0889	0.3789	0.3376
100nM	2.0024	0.2169	0.1914
1000nM	1.2596	0.9928	1.002
10000nM	5.9701	34.9016	8.2417

of the labels. The kernel function used in COT is $k(x, y) = 1/(\sigma^2 + c(x, y))^{0.5}$. For MNIST and CIFAR-10, we use the standard splits for training and testing and choose a random subset of size 10,000 from the training set for validation. For AWA, we use the train and test splits provided by [Jawanpuria et al., 2021] and randomly take 30% of the training data for validation.

Following [Jawanpuria et al., 2021], we compare all methods using the Area Under Curve (AUC) score of the classifier on the test data after finding the best hyperparameters on the validation data. Based on the validation phase, the best Sinkhorn regularization hyperparameter in ϵ -OT [Frogner et al., 2015] is 0.2. For COT, we choose the hyperparameters (λ, σ^2) based on the validation set: for MNIST (0.1, 0.1), for CIFAR-10 (0.1, 0.1) and for AWA (10, 0.1).

In Table 13, we also show the per-epoch computation time taken (on an RTX 4090 GPU) by the COT loss as a function of the size of the minibatch, which shows the computational efficiency of the COT loss.

Prompt Learning Let $\mathbf{F} = \{\mathbf{f}_m |_{m=1}^M\}$ denote the set of visual features for a given image and $\mathbf{G}_r = \{\mathbf{g}_n |_{n=1}^N\}$ denote the set of textual prompt features for class r . PLOT [Chen et al., 2023] learns the prompt features by performing an alternate optimization where the inner optimization solves an OT problem between the empirical measure over image features (49) and that over the prompt features (4). We denote the OT distance between the visual features of image \mathbf{x} and the textual prompt features of class r by $d_{OT}(\mathbf{x}, r)$. Then the probability of assigning the image \mathbf{x} to class r is computed as $p(y = r | \mathbf{x}) = \frac{\exp((1 - d_{OT}(\mathbf{x}, r)) / \tau)}{\sum_{r=1}^T \exp((1 - d_{OT}(\mathbf{x}, r)) / \tau)}$, where T denotes the total no. of classes and τ is the temperature of softmax. These prediction probabilities are then used in the cross-entropy loss for the outer optimization.

Following [Chen et al., 2023] and [Zhou et al., 2022b], we choose the last training epoch model. The PLOT baseline empirically found 4 to be the optimal number of prompt features. We follow the same for our experiment.




Figure 10: Marginals for selected genes ‘ENSG00000165092.12’, ‘ENSG00000175175.5’, ‘ENSG00000173727.12’, where the dosage is 100nM, in the outsample setting.

Table 10: Out-of-sample setting: MMD distances (lower is better) between predicted and ground truth distributions where the marker genes are computed at a per-drug level.

Dosage	CellOT	CondOT	Proposed
10nM	0.0369	0.0065	0.0071
100nM	0.0342	0.0061	0.0070
1000nM	0.0215	0.0178	0.0151
10000nM	0.2304	0.3917	0.3591

Table 11: Out-of-sample setting: l_2 (PS) distances (lower is better) between predicted and ground truth distributions where the marker genes are computed at a per-dose level.

Dosage	CellOT	CondOT	Proposed
10nM	1.2130	0.4718	0.3950
100nM	0.8561	0.2846	0.2522
1000nM	0.9707	0.9954	1.0775
10000nM	5.8737	33.5211	7.1487

We also keep the neural network architecture and hyperparameters the same as in PLOT. For our experiment, we choose λ , kernel type and the kernel hyperparameter used in COT. We choose the featurizer in Figure (5) as the same image encoder used for getting the visual features. We use a 3-layer MLP architecture for ψ_r in equation 10. We choose λ from $\{1, 10, 100\}$, kernel type from $k(x, y) = \exp \frac{-\|x-y\|^2}{2\sigma^2}$ (referred as RBF), $k(x, y) = (\sigma^2 + \|x - y\|^2)^{-0.5}$ (referred as IMQ), $k(x, y) = \left(\frac{1+\|x-y\|^2}{\sigma^2}\right)^{-0.5}$ (referred as IMQ2), kernel hyperparameter (σ^2) from $\{\text{median}, 0.01, 0.1, 1, 10, 100\}$. The chosen hyperparameters, (λ , kernel type, kernel hyperparameter), for the increasing number of shots (1 to 8), are (100, RBF, 10), (100, IMQ2, 1), (10, IMQ, 1), (1, IMQ, 0.01).

Figure 11 shows attention maps corresponding to each of the prompts learnt by COT. Table 12 presents an ablation study.

S3 MORE DETAILS

S3.1 Motivation for the Use of MMD

As the abstract motivates, the main challenge in formulating OT over conditionals is the unavailability of the conditional distributions, which is handled by COT using MMD-based kernelized-least-squares terms computed over the joint samples that implicitly match the transport plan’s marginals with the empirical conditionals. This results in the equivalence between Eqn (4) and Eqn (5). Furthermore, the statistical efficiency of MMD (Lemma 1) helps derive the consistency result (Thm. 1). Moreover, as discussed in § 4.2, the MMD metric is meaningful even for distributions with potentially non-overlapping support, enabling us to model the transport plan with implicit models for applications like those in § 5.2. Finally, the closed-form expression for MMD (discussed in § 2) helps in computational efficiency.

Table 12: Out-of-sample setting: MMD distances (lower is better) between predicted and ground truth distributions where the marker genes are computed at a per-dose level.

Dosage	CellOT	CondOT	Proposed
10nM	0.01648	0.00641	0.00638
100nM	0.01133	0.006325	0.00571
1000nM	0.01607	0.01496	0.01462
10000nM	0.24234	0.41845	0.34246

Table 13: Time (in s) for COT loss 7 computation shown for increasing minibatch size. The computation time reported is based on 3 independent runs on the CIFAR-10 dataset.

	16	64	512	1024
Time (s)	0.229±0.0013	0.229±0.0006	0.227±0.0004	0.225±0.0021




Figure 11: The leftmost is an image from the EuroSAT satellite dataset followed by visualization maps corresponding to each of the 4 prompts learnt (using COT loss 10). We can see that the 4 prompts diversely capture different visual features of the image.

S3.2 The Choice of Baselines

Other baselines for § 5.1.1 and 5.1.2: CKB and CondOT are inapplicable to Fig 2. CondOT requires multiple samples for each conditioned variable (§ 3 and Table 1). Using CKB, Wasserstein distance conditioned at an x can't be computed, which is needed for § 5.1.1. Also, it does not provide an OT plan/map needed for § 5.1.2. Hence, these are inapplicable. We will add this clarification in § 5. For the downstream applications in § 5.2 and § 5.3, we compare with the state-of-the-art baselines. However, for completeness's sake, we extended other baselines to these applications. The results obtained by [Tabak et al., 2021] for Table 2 are (7.1758, 56.682, 559.42, 5588.14), for Table 3 are (0.2438, 0.587, 0.582, 0.600) and for Table 4 are 0.49(MNIST), 0.52(CIFAR10), 0.52(AWA). As Tables 2 and 3 need an OT map, CKB can't be applied. CondOT doesn't apply to Table 4 as they need multiple samples for each conditioned variable. Table 5 results with ([Tabak et al., 2021], CKB, CondOT) are: (29.13±0.90, 29.7±2.41, 23.97±0.98) for $K = 1$, (38.87±2.00, 26.1±5.31, 21.8±6.39) for $K = 2$, (33.07±1.94, 28.87±1.58, 22.3±6.98) for $K = 4$ and (32.10±0.49, 27.3±4.61, 21.67±3.09) for $K = 8$. The results in the manuscript (§ 5) can be seen better than the above newly added.

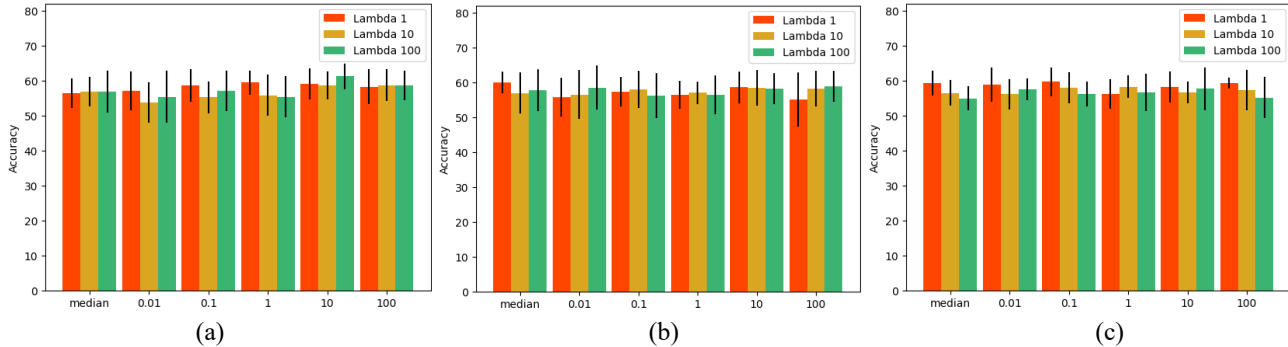


Figure 12: Ablation study for the prompt learning experiment $K = 1$. For different kernel types: (a) RBF kernel $k(x, y) = \exp\left(-\frac{\|x-y\|^2}{2\sigma^2}\right)$ (b) IMQ kernel $k(x, y) = (\sigma^2 + \|x - y\|^2)^{-0.5}$ (c) IMQ kernel $k(x, y) = \left(\frac{1+\|x-y\|^2}{\sigma^2}\right)^{-0.5}$, we show the average accuracy for different kernel hyperparameters (shown on the x-axis) and different lambda values.

S4 NEGATIVE SOCIETAL IMPACT

We present a formulation for solving optimal transport between conditional distributions. This problem has many socially beneficial applications, like predicting cell responses to cancer treatment, as shown in our paper. However, if a malicious task is selected, the proposed COT formulation may have a negative societal impact, similar to most other methods in machine learning.