Earth Movers in The Big Data Era: A Review of Optimal Transport in Machine Learning

Abdelwahed Khamis, Russell Tsuchida, Mohamed Tarek, Vivien Rolland, and Lars Petersson

Abstract—Optimal Transport (OT) is a mathematical framework that first emerged in the eighteenth century and has led to a plethora of methods for answering many theoretical and applied questions. The last decade is a witness of the remarkable contributions of this classical optimization problem to machine learning. This paper is about where and how optimal transport is used in machine learning with a focus on the question of salable optimal transport. We provide a comprehensive survey of optimal transport while ensuring an accessible presentation as permitted by the nature of the topic and the context. First, we explain optimal transport background and introduce different flavors (i.e. mathematical formulations), properties, and notable applications. We then address the fundamental question of how to scale optimal transport to cope with the current demands of big and high dimensional data. We conduct a systematic analysis of the methods used in the literature for scaling OT and present the findings in a unified taxonomy. We conclude with presenting some open challenges and discussing potential future research directions. A live repository of related OT research papers is maintained in https://github.com/abdelwahed/OT_for_big_data.git

Index Terms—earth mover's distance, optimal transport, Wasserstein distance, Sinkhorn, domain adaptation, Kantorovich, WGAN, Gromov Wasserstein, sliced Wasserstein, generative models, Wasserstein barycenter, optimization

---

1 INTRODUCTION

Selecting the right notion of discrepancy or distance between objects is at the heart of many machine learning problems. This work is about Optimal Transport (OT), a field that defines many such notions between a variety of mathematical objects such as probability histograms, shapes and point clouds. OT views objects as heaps of sand (mass) and quantifies the distance between them as the least costly way to morph the first heap into the second. This seemingly abstract concept has found applications in problems including domain adaptation [1], object detection [2], reinforcement learning [3], graph representation and matching [4], [5], feature matching [6], analyzing deep learning generalization [7], generative modeling [8], knowledge distillation [9], fairness [10], [11] and many others.

Interestingly, the simple idea of moving mass in an optimal way has a rich history. It evolved over the years, branched into many fields, and produced a wealth of theoretical and practical knowledge. A practical wartime optimal allocation problem that was addressed by Leonid Kantorovich in the 1940s shaped the recent treatment of the topic. He realized that the problem is actually of broader interest and devised a "simple general method of solving this group of problems" [12].

By doing so, he unknowingly [13] revisited a similar transport problem studied 150 years earlier by the renowned mathematician Gaspard Monge; the problem of soil transport for construction purposes. Ultimately, Kantrovich's solutions constituted the main tools of linear programming (i.e., linearly-constrained linear optimization) and won him a Nobel Prize in economic sciences in 1975. This historical connection between the theory of OT of its applications persist until today fostering successful developments on both fronts. On the theoretical side, the success can be associated
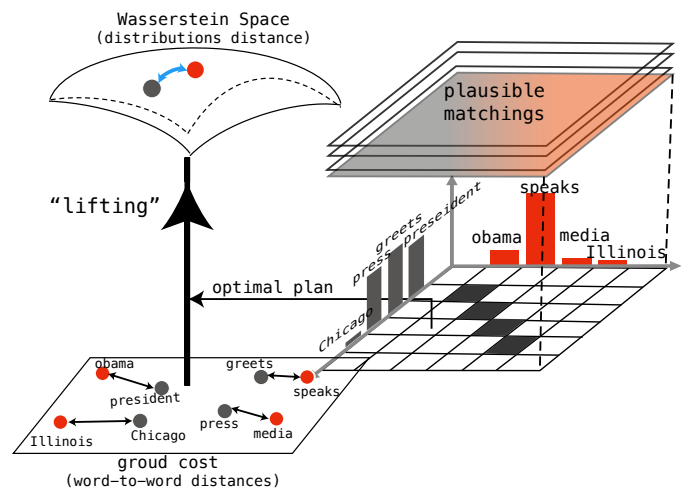

Fig. 1. **Optimal Transport geometry awareness**. Individual word-to-word distances may not capture the semantic similarity between two documents due to the lack of common words. OT leverages the underlying geometry (captured by the *ground cost*) and lifts it to the Wasserstein space where the distance represents the cost of the optimal transportation of a whole document (distribution) to another.

to OT's connection with several branches of mathematics. The parallel practical success is reflected in the growing adoption of OT in diverse applications

So, what does OT bring to the Machine Learning toolbox? Many machine learning problems boil down to comparing probability distributions. Consider the problem of assessing the similarity of two documents containing *"Obama speaks to the media in Illinois"* and *"The president greets the press in Chicago"* [14]. A similarity notion that relies on direct pairwise comparisons of the **individual elements** (i.e. words) would render the documents dissimilar, due to a lack of common words. However, our common sense goes tells us
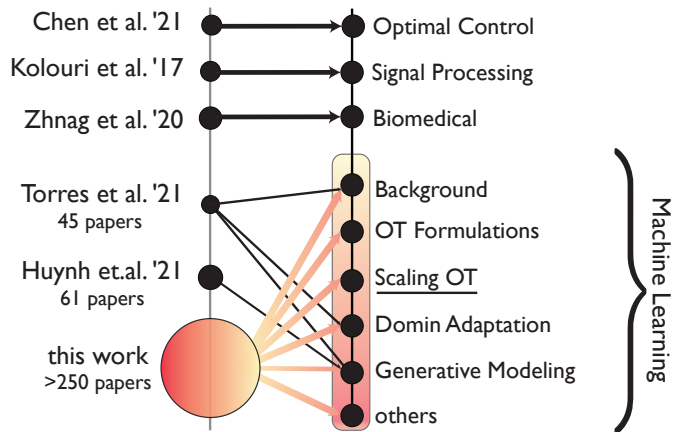
Fig. 2. **Survey Position.** Among OT reviews focusing on machine learning, this work presents a very comprehensive and updated coverage. The circles on the left column represent the existing reviews with their size indicating the number of the reviewed papers. The connections between the left and right columns depict the topics covered in each considered review article. The underlined "Scaling OT" is a key focus of this survey.

that the documents are actually similar. Looking at the two sets of words **as a whole**, the first sentence is semantically very close to the second. The latter insight motivates the OT prospective [15], [16]. OT accounts for the shapes of the distributions by computing a distance that measures how much effort is needed to align a distribution to another (Fig. 1) while allowing the distributions to have different supports. Among all possible ways to do such "global" alignment of distributions, OT picks the one with the least effort (*optimal plan*) guided by a notion of "local" distance/cost between the basic features of the distributions' supports (e.g. words). This local cost captures the geometry of the data and is usually called the *ground cost*[1]. In summary, OT lifts the basic word-to-word ground cost to a global distance in the distributional space. The appealing properties of OT can be summarized in the following two main points.

Underlying geometry awareness. Allowing geometric information of the support to be taken into account (unlike fairly common distances/divergences) can result in meaningful and smooth divergences, even when the two distributions don't overlap. This powerful feature is leveraged in many applications including OT-based embedding [17] and generative modeling [18]. Geometry awareness also means that OT borrows properties of the underlying space [19]. For example, concepts such as interpolation and barycenter from Euclidean ground space carry on to the Wasserstein space. Thus, one can perform smooth interpolation between probability distributions [20] or even perform regression on histograms [21].

Explicit Correspondence: in addition to distances between distributions, OT typically provides correspondences through the optimal plan. Such correspondences can be used for applications such as object matching and registration [22], [23], and interpreting machine learning results (§B.1).

What is the price tag? The above-mentioned desirable characteristics come with longstanding computational and

statistical challenges. First, OT-based distances are hard to compute and explicit closed-form expressions for OT-based distances are rare beyond Gaussian and one-dimensional cases. In its original formulation, OT-based distance computation requires solving a linear program with a super-cubic complexity in the number of data points. Reducing this problematic complexity using various computational techniques and workarounds is an ongoing and interesting trend in the OT literature. We review many of these techniques (§5) and categorize them in a taxonomy that reveals the connections among various techniques. Second, in high dimensional settings, OT suffers from a severe "curse of dimensionality". The sample complexity[2] of estimating the Wasserstein distance using a finite sample from the distributions grows exponentially in the dimension [24], [25]. To understand the gravity of the issue, contrast OT with another divergence, the Maximum Mean Discrepancy [26], whose sample complexity is dimension independent! Finally, there are limitations inherent to the vanilla OT formulation itself such as sensitivity to outliers [27], inability to incorporate context/structure in the optimal plan [28], strong reliance on the chosen ground cost [29], assuming a common ground cost [30] and the inability to handle measures (generalized, un-normalized distributions) of non-equivalent masses (normalization constants) [31]. Unsatisfied with these limitations, researchers have produced a diverse spectrum of OT formulations over the years with potential to meet a broader range of applications and requirements. Indeed, new formulations are still evolving as of this writing. For each formulation, we briefly discuss the characteristics, computational methods for solving it, and ML applications in which the formulation was used.

In the last decade, many significant developments in OT have taken place. New mathematical formulations, highly scalable algorithms (coping with datasets with millions of samples), interactions with deep learning, the rise of learned neural solvers, and empirical advancements in domain adaptation and generative modeling are a few developments to name. These developments call for a refreshed review of the field that sheds light on recent developments and identifies emerging trends and paradigms. The previous excellent reviews (Fig. 2) either have a narrow coverage of OT in machine learning or a specialized coverage irrelevant to machine learning. Our goal is to present a timely and comprehensive survey of the important topic of OT and OT applications in modern machine learning.

We summarize the contributions of this work as follows:

1) We present a comprehensive survey on the topic of OT with a focus on the applied side of the literature in a way that best targets the machine learning research audience.
2) We compose a high-level view that organizes the techniques in literature for addressing the scalability issues of optimal transport (i.e. computational and statistical challenges) in a single framework (§5).
3) We highlight literature gaps and discuss promising future research directions in OT and its applications.

---

1. Note that other competing distances such as total-variation or χ^2-distances are oblivious to any notion of similarity on the ground space.

2. Sample complexity, informally, is the number of samples needed to estimate a function within a certain level of accuracy.
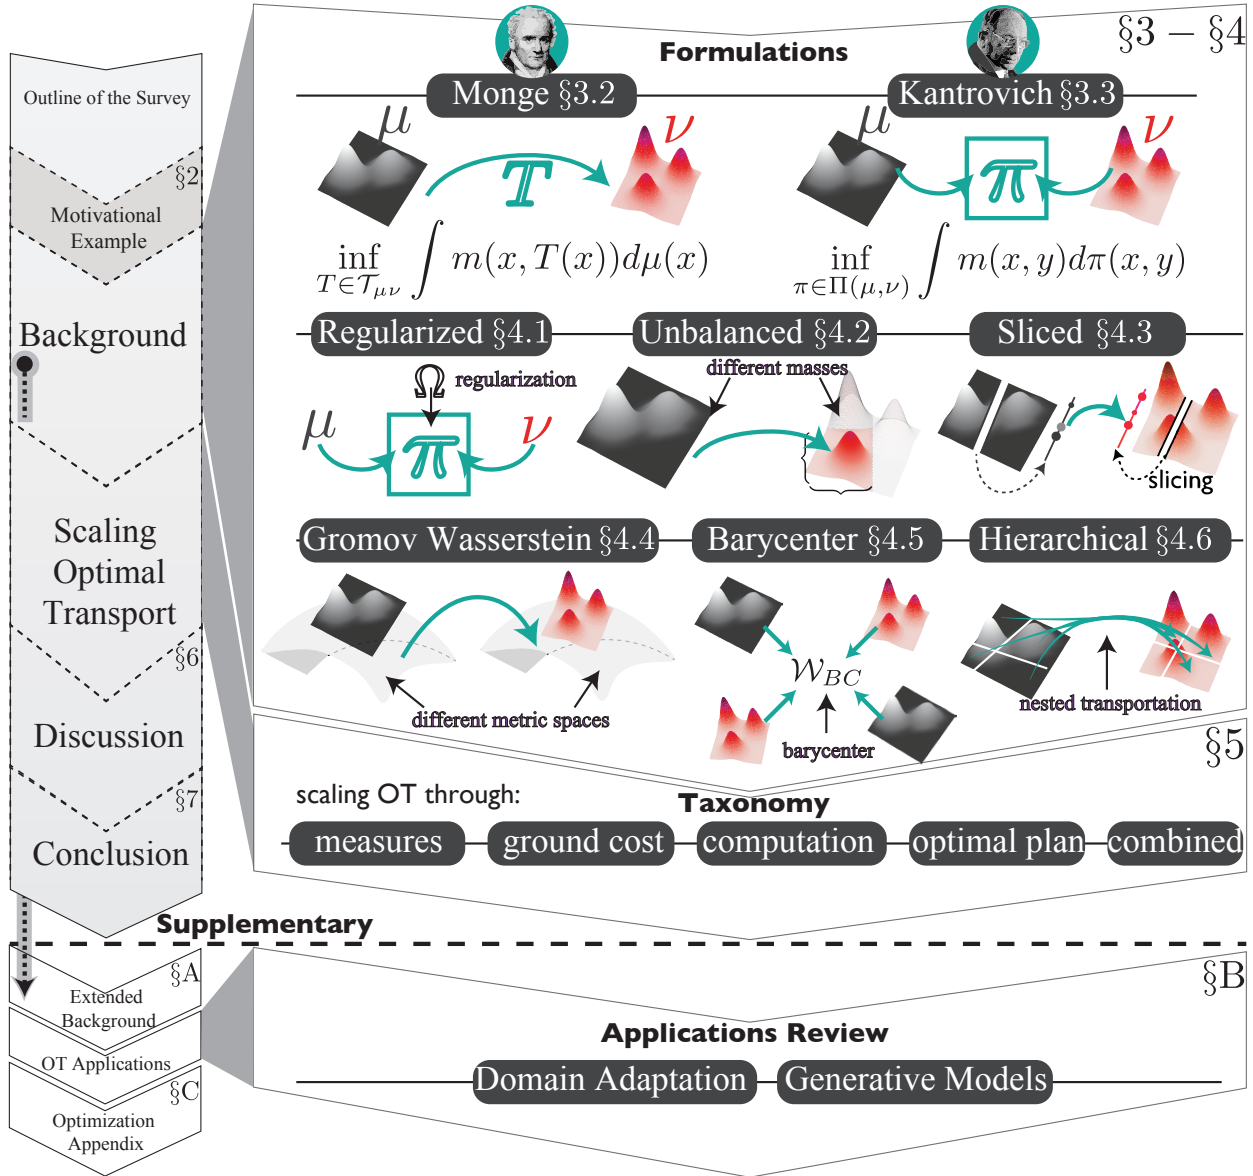
Fig. 3. **Outline of the Survey**. We start by **(§2 Motivational Example)** introducing the reader to OT through a motivational example that paves for **(§3, §4 Background)** the discussion of OT formulations. Next, **(§5 Scaling OT)** we present a taxonomy for scaling OT methods to big data regimes. Then, open issues and future research directions are discussed in **(§6 Discussion)**. In the supplementary material, we extend the background discussion **(§A Extended Background)** and include a summary of OT applications in machine learning**(§B Applications)**. All the figures are best viewed in color.

The remainder of this paper is organized as outlined in Fig. 3. We took some efforts to make the paper self-contained. It is difficult to fit the OT-related optimization background into a few pages; we suggest that §C can help the uninitiated readers before seeking other resources.

2 A MOTIVATIONAL EXAMPLE

In this section, we give a simple motivational example to introduce OT. Consider the problem of comparing point cloud objects. The objects are composed of 3D points that can be acquired by contemporary sensors such as Lidar, Radar, or depth sensors. Figure 4(a) shows examples of point cloud objects. To compare these objects, we seek a matching that redistributes the first object into the second in the least costly way (i.e. optimal transportation). Let's assume that there

are n and m points in the objects μ and ν; respectively. Furthermore, there is some mass (depicted as the volume of the point) that can be exchanged between the two objects during transportation. This mass can be, for example, the *intensity* of the point [3]. The term "mass" may be also used to refer to abstract or conceptual quantities. When unknown, we can assume that the masses are uniformly distributed. The mass of one point in the source can be exchanged with one or multiple points in the target.

Regardless of the notion of mass used, we assume that the points (support) $\{x_i\}_{i=1}^n$ of the object μ have non-negative mass each that sum up to 1 (i.e. normalized). We denote the vector of point masses in μ by $r \in \mathbb{R}^n$. The same

3. For example, in radar's point cloud the intensity is the strength of Radio Frequency (RF) reflection from the sensed point [32].
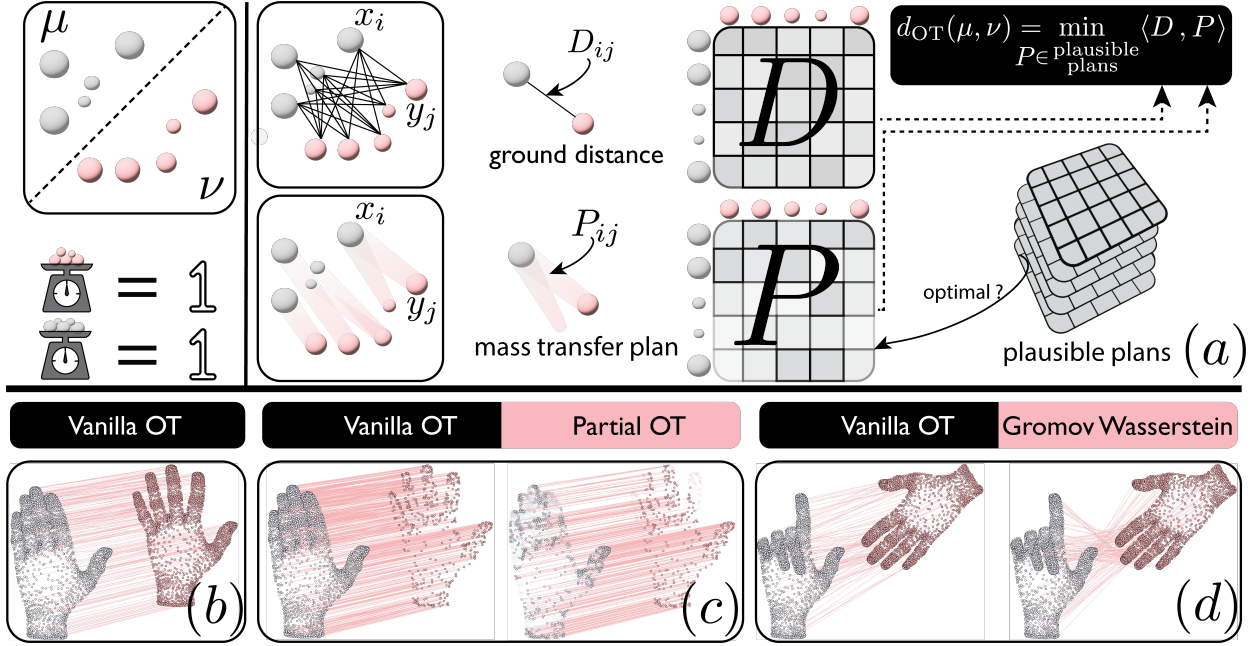
Fig. 4. Conceptual Depiction of Optimal Transport: (a) The optimal transport problem is stated as follows. Given two points cloud objects μ and ν and the knowledge of point-wise distances (i.e. the ground distance), find a legitimate way (i.e. plausible plan) to redistribute the mass of μ into ν in the least costly way P^*. The optimal transport cost is then given by $d_{\text{OT}}(\mu, \nu) = \langle D, P^* \rangle$. **(b)** Matching of two point cloud objects using optimal transport. The point correspondences estimated by the optimal plan are shown as straight pink lines. OT literature is filled with formulations that extend this key concept to new situations and applications. For example, formulations that allow **(c)** partial transportation or **(d)** transportation between incomparable spaces.

assumption is made for the object ν whose points $\{y_i\}_{i=1}^m$ are assumed to have a normalized mass vector $c \in \mathbb{R}^m$.

Suppose we wanted to transport the mass from the particles in μ to the particles in ν in way that minimizes some total measure of cost, e.g. the sum of "mass moved" × "the distance it moved". More rigorously, we would like to decide on how much mass to move from point x_i to point y_j; P_{ij}, considering the distance between x_i and y_j; D_{ij}. The first P_{ij} is sometimes known as a transfer *plan* which is to be identified while the latter D_{ij} is known as *ground distance* (e.g. the squared Euclidean distance $\|x_i - y_j\|^2$) which is given. Let $P \in \mathbb{R}^{n \times m}$ and $D \in \mathbb{R}^{n \times m}$ denote these quantities. The goal of OT is to find an optimal plan P^* that dictates the mass distribution plan from the source points to the destination points in a way that minimizes the total transport cost. The set of all plausible plans (formally defined in **Couplings**) is described as follows:

$$\text{plausible plans} = \{P \in \mathbb{R}_+^{n \times m} | \sum_j P_{ij} = r_i, \sum_i P_{ij} = c_j\}.$$

The constraints ensure that P preserves the mass so no mass is created or destructed during transportation. In light of this, the optimal plan P^* would be one that minimizes

$$d_{\text{OT}}(\mu, \nu) = \min_{P \in \text{plausible plans}} \langle D, P \rangle,$$

where $\langle \cdot, \cdot \rangle$ is the Frobenius inner product. Applying the formulation above on point cloud objects (Fig 4(b) hand gestures from the CAPOD dataset [33]), we can learn about the global proximity of the two shapes through the distance d_{OT}. Moreover, we learn the point correspondences conveyed by the optimal plan P^* (visualized as straight lines).

Posed this way, the OT problem (sometimes known as the Kantrovich formulation) possesses a useful set of characteristics as well as limitations, which motivate other formulations of the problem. We discuss many of them in §4. For example, when the masses of the objects aren't equal, one might wish to relax the mass preservation constraint. As shown in Fig. 4(c), the (sparsely sampled) target object's total mass is much smaller than the source. The vanilla OT forces the transportation of all points from the sources resulting in inaccurate matching and (unrealistically) increased cost. A partial OT formulation (§4.2), however, is a better alternative as it has the flexibility of leaving some points out from the transportation. Another limitation is that transportation that relies only on point correspondences conveyed by the ground distance (without considering the inter-points relations) can be misleading. In 4(d), we see a simple rotation transformation of the target object causes a degradation in the alignment by vanilla OT. This can be clearly seen in the incorrect correspondences of the index and the thumb fingers. This is not the case with a relational formulation (Gromov Wasserstein §4.5) that handles this issue well at the expense of increased computation. This formulation also can compare points defined on distinct (incomparable) spaces which can be useful for comparing graphs or measurements from different modalities.

3 BACKGROUND

We begin with notation and background of the key OT formulations. We state the Monge and Kantrovich formulations below, and refer to an extended discussion of Kantrovich formulation in §A.

3.1 Notation

Probability measures. Let $\mathbb{X}$ and $\mathbb{Y}$ be two non-empty sets. Denote by $\mathcal{P}(\mathbb{X})$ the set of probability measures supported on $\mathbb{X}$. Let $\mu \in \mathcal{P}(\mathbb{X})$ and $\nu \in \mathcal{P}(\mathbb{Y})$. A function $\phi : \mathbb{X} \to \mathbb{R}$ is said to be integrable with respect to μ if it has a finite expectation under μ, i.e. $\mathbb{E}_{X \sim \mu}|\phi(X)| = \int_{\mathbb{X}} |\phi(x)|\, d\mu(x) < \infty$. The space of functions integrable with respect to μ is denoted

$$L^1(\mu) = \{\phi : \mathbb{X} \to \mathbb{R} \mid \mathbb{E}_{X \sim \mu}|\phi(X)| < \infty\}.$$

The mathematical objects that we discuss can take on a discrete, continuous or mixed nature. Averages can be notated using one or more of sum, integral or expectation notation. While in principle these objects can be dealt with using one unified notation as an integral with respect to a measure, for ease of reading we sometimes prefer to explicitly write out the same equation using multiple notations. To simplify notation, we will assume that every probability measure has an associated probability density function or probability mass function.

Nonnegative unnormalised measures. Most of the time, we will be interested in distances defined over the space of probability measures. Unless otherwise stated, all measures are probability measures. In some cases (for example, § 4.2), these distances generalise to spaces containing nonnegative and finite measures that are not necessarily probability measures. Where we deal with such nonnegative unnormalised measures, we will mention this close to where they are introduced. In such cases, the expectation notation $\mathbb{E}_{X \sim \mu}\cdot$ is inappropriate and we resort to integrals $\int_{\mathbb{X}} \cdot d\mu(x)$ or sums. We denote by $\mathcal{M}_+(\mathbb{X})$ the space of finite nonnegative unnormalised measures supported over $\mathbb{X}$.

Coupling. A useful notion appearing in OT (for example, in the Kantorovich formulation Eq. (2)) is that of coupling. Given two random variables X and Y, we may construct a random vector (X, Y) having marginal distributions that are the same as the distributions of X and Y, and a non-unique joint distribution consistent with those marginals. The space Π of all such joint distributions constitutes a useful space over which to search. Define the set $\Pi(\mu, \nu)$ of joint probability measures with marginals μ and ν. In other words, $\Pi(\mu, \nu)$ is the set of plausible plans

$$\Pi(\mu, \nu) = \Big\{ \pi \in \mathcal{P}(\mathbb{X} \times \mathbb{Y}) \mid \forall A \subseteq \mathbb{X}, B \subseteq \mathbb{Y},$$
$$\int_{A \times \mathbb{Y}} d\pi(x, y) = \int_A d\mu(x), \int_{\mathbb{X} \times B} d\pi(x, y) = \int_B d\nu(y) \Big\}.$$

Matrix representations for discrete and finite support. It is useful to instantiate measures and couplings in terms of vectors and matrices in the case where the random variables in question are discrete and have finite support.

Let $\widetilde{r}$ and $\widetilde{c}$ be the probability mass functions corresponding with μ and ν. Let the support of $\widetilde{r}$ and $\widetilde{c}$ be finite, that is $d_{\mathbb{X}} = |\mathbb{X}| < \infty$ and $d_{\mathbb{Y}} = |\mathbb{Y}| < \infty$. Define the probability simplex $\Sigma_d = \{x \in [0, 1]^d \mid x^\top \mathbf{1} = 1\}$. Let $r \in \Sigma_{d_{\mathbb{X}}}$ and $c \in \Sigma_{d_{\mathbb{Y}}}$ be vector representations of $\widetilde{r}, \widetilde{c}$. That is, for some assignment $\{x_i \mid 1 \le i \le d_{\mathbb{X}}\} = \mathbb{X}$ and $\{y_i \mid 1 \le i \le d_{\mathbb{Y}}\} = \mathbb{Y}$, set $r_i = \widetilde{r}(x_i)$ and $c_i = \widetilde{c}(y_i)$. The set $\Pi(\mu, \nu)$ as corresponds with a set of finite-dimensional matrices $U(r, c)$ representing joint probability mass functions consistent with marginals μ and ν. Define

$$U(r, c) = \{P \in [0, 1]^{d_{\mathbb{X}} \times d_{\mathbb{Y}}} \mid P\mathbf{1} = r, P^\top \mathbf{1} = c\},$$

and note that for any $P \in U(r, c)$, we have that the ijth entry of P satisfies $p_{ij} = \Pr(X = x_i, Y = y_j)$ under some joint probability mass function for random variables (X, Y) with marginal probability mass functions $\widetilde{r}$ and $\widetilde{c}$.

Empirical measures When dealing with real-world data and problems, we usually don't have access to explicit forms of distributions μ and ν. Instead, we have samples $\{x_i\}_{i=1}^m \sim \mu$ and $\{y_j\}_{j=1}^n \sim \nu$. In these settings, it is useful to define and work with the empirical measures $\mu_n = 1/n \sum_i \delta_{x_i}$ and $\nu_n = 1/n \sum_j \delta_{y_j}$. Here δ_a is a point mass centered at a.

3.2 Monge Formulation (Optimal Map)

Given probability measures μ and ν, the OT map $T^* : \mathbb{X} \to \mathbb{Y}$ is defined to be one that maps random variables X following μ to random variables Y following ν with the minimum expected cost between X and $T^*(X) = Y$. The Monge formulation [34, Chapter 1] is :

$$T^* = \inf_{T \in \mathcal{T}_{\mu\nu}} \int_{\mathbb{X}} m(x, T(x)) d\mu(x)$$
$$= \inf_{T \in \mathcal{T}_{\mu\nu}} \mathbb{E}_{X \sim \mu} [m(X, T(X))]. \tag{1}$$

Here $\mathcal{T}_{\mu\nu} = \{T : \mathbb{X} \to \mathbb{Y} \mid T_\#(\mu) = \nu\}$ constitutes the set of all possible mappings. $T_\#(\mu)$ is the pushforward measure of μ under T. In the context of probability measures this means that $T_\#(\mu)$ is the probability measure of the random variable $T(X)$. m is the ground cost responsible for transporting one unit of mass from X to Y. The choice of m is typically application dependent and influenced by domain knowledge and awareness of data.

The Monge formulation Eq. (1) may be ill-posed; there may not exist any mapping T such that $T_\#(\mu) = \nu$. For example, if μ places probability mass over 2 states, it can never be mapped to a measure that places probability mass over 3 states using a deterministic function. The Kantorovich formulation, introduced next, does not suffer from this issue.

3.3 Kantrovich Formulation (Optimal Plan)

Given μ and ν and some cost function m, the Kantorovich OT problem [34, Chapter 1] (the general formulation of the discrete example in §2) is defined as:

$$\mathcal{K}_m(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{X} \times \mathbb{Y}} m(x, y) d\pi(x, y)$$
$$= \inf_{\pi \in \Pi(\mu, \nu)} \mathbb{E}_{(X, Y) \sim \pi} [m(X, Y)]. \tag{2}$$

Any π^* that attains this infimum is called an OT plan. In the case that X and Y are discrete random variables with probability mass vectors r and c, we write

$$\mathcal{K}_m(\mu, \nu) = \min_{P \in U(r, c)} \langle M_{XY}, P \rangle \tag{3}$$

where $M_{XY} = [m(x_i, y_j)]_{ij}$ is the matrix containing the pairwise costs between X and Y. Interestingly, Eq. (2), known as the *primal* formulation, can be written in a number of equivalent formulations (such as the *dual* formulation). **Check §A for intuitive and comprehensive exposition on alternative formulations.**

Wasserstein Distance: It may be shown that if $m(x, y) = d(x, y)^p$ for some distance d on $\mathbb{S} = \mathbb{X} \cup \mathbb{Y}$, i.e. $d : \mathbb{S} \times \mathbb{S} \to \mathbb{R}$ and some $p \geq 1$, then the p-Wasserstein distance [35, Definition 6.1] :

$$W_p(\mu, \nu) = \left(\mathcal{K}_{d^p}(\mu, \nu) \right)^{1/p} \quad (4)$$

is a distance metric over the space of probability measures with finite moments up to order p (Theorem 7.3 (i) [36]). In the case where $0 \leq p < 1$, $W_p^p(\mu, \nu)$ is a distance metric over the space of probability measures with finite moments up to order p (Theorem 7.3 (ii) [36]).

4 COMMON OT FORMULATIONS

Here we discuss OT formulations with a focus on the balance between proper coverage and concise treatment.

4.1 Regularized OT

Formulation and Characteristics: Some of the computational and statistical limitations of the original OT formulation (Eq. 2) can be mitigated by augmenting the original objective with a regularization term. Regularized OT was brought to the forefront of applied machine learning after Cuturi published his seminal work on the topic ten years ago [37]. The work demonstrated the computational superiority of the entropic regularized formulation and its compatibility with modern ML pipelines (i.e. parallelism and differentiability). Yet, the formulation can be traced back to the works of Erwin Schrödinger in 1930s [38][4]. The regularized OT problem is defined as:

$$\mathcal{OT}_{\Omega,\lambda}(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{X} \times \mathbb{Y}} m(x, y) d\pi(x, y) + \lambda\, \Omega(\pi), \quad (5)$$

where $\Omega(P)$ is a regularization operator and λ is a positive regularization coefficient. The regularized discrete OT problem is similarly given by:

$$\mathcal{OT}_{\Omega,\lambda}(\mu, \nu) = \inf_{P \in U(r,c)} \langle M, P \rangle + \lambda\, \Omega(P). \quad (6)$$

One choice for Ω is the KL divergence between the joint distribution P (or measure π) and the product measure $m_1 \times m_2$ for a set of reference measures m_1 and m_2. That is, $\Omega(P) = \text{KL}(P \,||\, m_1 \times m_2)$.

When additionally the reference measures m_1 and m_2 are uniform, the regularisation term (up to a constant) is called entropic regulaisation:

$$\Omega(P) = -H(P) = \sum_{ij} p_{ij}(\log p_{ij} - 1),$$

where H is the Boltzmann-Shannon entropy function. Adding negative of the Boltzmann-Shannon entropy function as a regularization gives us the famous entropic-regularized OT:

$$\mathcal{OT}_\lambda(\mu, \nu) = \inf_{P \in U(r,c)} \langle M, P \rangle - \lambda H(P), \quad (7)$$

which is a version of the original problem Eq. 3 with an additional strictly convex regularization term (negative

4. Recently, OT and Schrödinger bridge have provided insights into diffusion-based generative models [39], [40].

of the Boltzmann-Shannon entropy). Intuitively, this regularization biases the optimal plan towards the uniform distribution. Interestingly, entropic-regularized OT (whose optimal objective value is also known as the Sinkhorn "distance"

can be rewritten as minimisation of a Kullback-Leibler (KL) divergence,

$$\mathcal{OT}_\lambda(\mu, \nu) = \lambda \inf_{P \in U(r,c)} \text{KL}(P \,||\, K), \quad (8)$$

where the matrix K is the element-wise negative exponential of the scaled ground cost $K = \exp(-M/\lambda)$. The KL divergence is known to be a strictly convex function of P for a fixed K which means that there exists a unique P that minimizes the above expression and that is differentiable with respect to K. Note that the Sinkhorn "distance" is symmetric in μ and ν, i.e. $\mathcal{OT}_\lambda(\mu, \nu) = \mathcal{OT}_\lambda(\nu, \mu)$, if the cost function $m(x, y)$ is symmetric in x and y.

Computational Motivation: While the regularization term can be motivated by application requirements through the inclusion of an additional prior [41], one can safely argue that the computational advantage of entropic regularization is a key appealing benefit. Specifically, the formulation in Eq. 7, which was popularized in the GPU era by Cuturi [42], can be solved efficiently using an iterative matrix scaling algorithm known as the Sinkhorn-Knopp [43] algorithm, often shortened to the Sinkhorn algorithm. This algorithm boils down the process of computing the optimal plan to a series of matrix-vector multiplications. The Sinkhorn algorithm has a worst-case quadratic complexity in the number of points per iteration in its base form (with many enhancements possible, e.g. [44]) and linear convergence rate. This is much more efficient and scalable than the linear programming methods employed by solvers for the unregularized formulation.

Sinkhorn Divergence: Despite the computational appeal of the the Sinkhorn "distance", two limitations arise in Eq. 7. First, the regularized OT optimal value is not a distance. Second, the formulation induces a bias in the minimizer known as the entropic bias. This means that μ tends to be transported increasingly to a shrinked version of ν as λ increases. These issues are handled by the Sinkhorn divergence [45], [46]; another variant of regularized OT that builds on Sinkhorn "distance" and is defined as follows:

$$\overline{\mathcal{OT}}_\lambda(\mu, \nu) = \mathcal{OT}_\lambda(\mu, \nu) - \frac{1}{2}\Big(\mathcal{OT}_\lambda(\mu, \mu) + \mathcal{OT}_\lambda(\nu, \nu)\Big). \quad (9)$$

Sinkhorn divergence is smooth and differentiable in the inputs and enjoys a better sample complexity [47] compared to regular OT. Yet, it can be seen that it requires three evaluations of $\mathcal{OT}_\lambda$ rather than one, although the symmetric Sinkhorn algorithm can be used to compute $\mathcal{OT}_\lambda(\mu, \mu)$ and $\mathcal{OT}_\lambda(\nu, \nu)$ more quickly than the regular Sinkhorn algorithm. Additionally, the associated quadratic cost/complexity makes it challenging to compare measures that have tens of thousands of points [48]. Note that we discuss scalable variants of sinkhorn later in §5.3 and §5.5.

Applications: Entropic-regularized OT and the Sinkhorn "distance" have been extensively used in many machine learning applications. Some examples include self labelling [49], adversarial examples [50], permutations learning [51] [52] in deep models, initializing graph correspondence in

graph matching learning [53], differentiable sorting and ranking [54] [55] and enforcing doubly stochastic attention in transformers [56]. Since the kind of regularization used in OT can be adapted depending on the application requirements, some works investigate alternatives to the entropy term in Eq.7. Graph-inspired regularization was used in domain adaptation [57] and color transfer application [41]. This was extended by Laplacian OT [58] and [57]. Other regularization terms that reflect various biases/prior, such as Lasso and Group Lasso [59], also exist in the literature. Temporal regularization [60] was used in human pose alignment [60] and [61] weakly supervised action segmentation.

Limitations: A known limitation of entropic regularized OT is the difficulty of setting the regularization strength λ. Big computational gains are tied to a high level of regularization (i.e. bigger λ) which can cause overspreading (i.e. blurry optimal plan) [59]. For example, in the point matching example (§2), high regularization translates to dense matching where most points are connected to each other. This can hurt the correspondence interpretation. On the other hand, a concentrated sparse plans can theoretically be obtained for very small λ but in practice, this causes numerical issues [57], [62].

4.2 Unbalanced and Partial OT

Motivation: The mass preservation constraint in the Kantorovich formulation Eq. (2) requires that the total mass between the two probability distributions be the same. Unbalanced OT (UOT) refers to formulations that relax this constraint to allow transporting of arbitrary (unnormalized) measures or partial masses. This is useful in applications like multi-object tracking [63], and crowd counting [64], [65].

Formulation and Characteristics: One way of extending OT to arbitrary positive measures is by adopting *marginals relaxation*, in which the hard marginal constraints Eq. (2) are removed and the objective is augmented with soft regularization that penalizes the mass variation. Typically, this is done using Csiszar divergence D_ϕ [66]:

$$\mathcal{UOT}_\phi(\mu,\nu) = \inf_{\pi \in \mathcal{M}_+(\mathbb{X}\times\mathbb{Y})} \int_{\mathbb{X}\times\mathbb{Y}} m(x,y)d\pi(x,y) \\ + \lambda_1 \, D_\phi(\pi_1|\mu) + \lambda_2 \, D_\phi(\pi_2|\nu). \quad (10)$$

where $\mathcal{M}_+(\mathbb{X}\times\mathbb{Y})$ is the space of finite non-negative measures over $\mathbb{X}\times\mathbb{Y}$ and D_ϕ is a divergence induced by ϕ that quantifies mass variation between the marginals of the plan's marginals (π_1,π_2) and the measures (μ,ν). The penalization strength is controlled by (λ_1,λ_2). An obvious consequence of replacing constraints with penalties is that π's marginals no longer need to be equal to (μ,ν). Particular instances of D_ϕ include KL divergence [67] and Total Variation [68]. Recently, [69] proposed a regularization based on Maximum Mean Discrepancy (MMD) and proves that the modified formulation attracts desirable properties such as dimension-free sample complexity.

Another way of extending OT to unbalanced measures, usually leveraged in discrete cases, is *partial assignment*. The formulation is called Partial OT (POT). When transporting from n to m points, one can adapt the formulation Eq. (2) by augmenting the plan $\pi \in \mathbb{R}^{n\times m}$ to be $\widetilde{\pi} \in \mathbb{R}^{(n+1)\times(m+1)}$ and similarly for the ground distance matrix M to be $\widetilde{M} \in$ $\mathbb{R}^{(n+1)\times(m+1)}$. The additional row and column are called *dustbin* [70], [71] or *dummy points* [72] and should absorb the mass from unmatched points. This in effect turns the formulation into a balanced one and thus the traditional computation of Eq. (2) can be recycled.

Computation: Entropic regularization may be combined with UOT [73] or POT [71]. This allows one to invoke the sinkhorn algorithm and speed up the computation.

Applications: The balanced OT formulation in Eq. (2) can be highly non-robust when there are outliers. Since it endeavors to transport all the mass from μ to ν, a single contaminated data point can increase the OT cost arbitrarily. UOT, on the other hand, can be less sensitive to this issue by assigning small masses to outliers. This makes UOT more suitable for machine learning applications that deal with corrupted and noisy data. OT variants that detect and don't transport outliers build on the unbalanced formulation. For example, ROBust OT (ROBOT) [27] performs outlier detection using a formulation that builds on UOT.

4.3 Sliced OT

Motivation: One approach to speed up the OT computation relies on the idea of low-dimensional projections by aggregating OT distances computed on 1D projections of the data. What motivates this approach is the fact that solving OT problems in the univariate/1D case is cheap. Specifically, we obtain a representation for the p-Wasserstein distance between one-dimensional μ and ν with quantile functions of F_μ^{-1} and F_ν^{-1}, respectively, as follows:

$$\mathcal{W}_p(\mu,\nu) = \Big(\int_0^1 m^p\Big(F_\mu^{-1}(t), F_\nu^{-1}(t)\Big)dt\Big)^{\frac{1}{p}} \quad (11)$$

where $m^p(.,.)$ is the ground cost function. When μ and ν are empirical distributions with n and m samples, (11) can be computed very efficiently through simple sorting with a complexity of $\mathcal{O}(n\log(n) + m\log(m))$.

Formulation and Characteristics: Sliced Wasserstein (SW) [74] is motivated by explicit utilization of the above computational efficiency. SW is computed as the average of infinitely many Wasserstein distances between one-dimensional projections of high-dimensional distributions. Formally, for any μ and $\nu \in \mathbb{R}^d$ the Sliced Wasserstein is:

$$\mathcal{SOT}_p^p(\mu,\nu) = \mathbb{E}_{\theta\sim\mathrm{Unif}(\mathbb{S}^{d-1})}\Big[\mathcal{W}_p^p(g_{\theta\#}\mu, g_{\theta\#}\nu)\Big] \quad (12)$$

where $g_\theta(x) = \theta^T x$ is the linear projection map parameterized by θ. Here $g_{\theta\#}\mu$ denotes the pushforward measure of μ, i.e. the probability measure associated with the random variable $g_\theta(X)$, where X is a random variable with probability measure μ. $\mathrm{Unif}(\mathbb{S}^{d-1})$ means uniformly distributed over the d-dimensional unit sphere $\mathbb{S}^{d-1}$. Since the projections of measures on the picked direction ($g_{\theta\#}\mu$ and $g_{\theta\#}\nu$) are one-dimensional, the $\mathcal{W}_p^p$ in (12) can be efficiently calculated using (11). In practice, however, acquiring infinite number of projections is not feasible and thus (12) is usually approximated using a Monte Carlo scheme. The integral is replaced with average calculated over finite number of L random projection directions:

$$\widehat{\mathcal{SOT}}_p^p(\mu,\nu) = \frac{1}{L}\sum_{l=1}^L \mathcal{W}_p^p(g_{\theta_l\#}\mu, g_{\theta_l\#}\nu) \quad (13)$$

where $\theta_1, \cdots, \theta_L \sim \text{Unif}(\mathbb{S}^{d-1})$.

It might be surprising that OT on the real line (i.e. sliced OT) can convey geometric information of high dimensional distributions. However, theoretical studies have shown that SW satisfy the metric axioms [75] and have convergence properties similar to that of the Wasserstein distance [76], [77]. Further theoretical results [78] show that SW has better sample complexity that does not depend on the problem dimension.

Applications: The Sliced Wasserstein is easier to compute and enjoys theoretical properties similar to that of the Wasserstein distance. Thus, it is used in many ML applications as a better alternative. A few examples include defining kernels [79] , generative models [80], [81], pooling [82], neural texture synthesis [81] , model selection [83], learning representations of 3D point clouds [84] and sets [?], and domain adaptation [85]. Additionally, it was employed in alleviating the computational burden of the expensive OT formulations such as Wasserstein barycenter [74] and Gromov Wasserstein [86] (§4.5).

Limitations: Research on extensions of SW is mostly concerned with two questions linked to limitations of (13) : *1) how do we better determine directions for projections θ_l?* and *2) what could be a better mapping than the linear g_θ?*

What motivates the first question is the increased *projection complexity* of randomly chosen θ_l. In other words, to achieve good approximation of 13 we need larger number of projections L [80]. A single experimental example of this is the Sliced Wasserstein Discrepany (SWD) [85] work on unsupervised domain adaptation. SWD shows that the accuracy drops from 98% when using 512 radial projections to below 90% when using only 32 projections in the Synthetic Traffic Signs (SYNSIG) [87] to German Traffic Sign Recognition Benchmark (GTSRB) [88] experiments. To overcome this limitation, one can consider a few informative projections instead of using all random ones. Following this idea, Max Sliced Wasserstein (Max-SW) [80] advocates for the single "best direction". For order $p = 2$, it is defined as follows:

$$max\mathcal{SOT}_2^2(\mu, \nu) = \max_{\theta \in \mathcal{S}^{d-1}} \mathcal{W}_2^2(g_{\theta\#}\mu, g_{\theta\#}\nu) \qquad (14)$$

where the best direction is one that yields the largest Wasserstein distance between the projected measures. Finding this best direction isn't trivial. In practice, it is replaced with direction that results in the largest difference between the means of the projected measures. Distributional Sliced Wasserstein (DSW) [89] claims that focusing only on the most important direction (e.g. Max-SW) ignores other potentially relevant directions. Thus, it takes a middle ground between SW and Max-SW by searching for a "distribution" of the important directions on the unit sphere.

The second question regarding SW limitations concerns the linear g_θ. What motivates this is the hypothesis that nonlinear mappings could be better in high dimensional settings. Generalized Sliced Wasserstein (GSW) [90] formalizes this idea. Also, it is possible to combine nonlinear projection with better projection complexity. Max Generalized Sliced Wasserstein (max-GSW) [90] is an example of this trend.

4.4 Wasserstein Barycenter

Motivation: The barycenter of a collection of measures is an intuitive notion of an "average" of the measures. For a set of

vectors, the usual average of the vectors is the single vector v that minimises the sum of squared Euclidean distances between v and every vector in the set. The barycenter of a set of measures provides a useful first order summary statistic of a set of measures, just as the average of a set of vectors provides a useful summary statistic of those vectors.

Formulation and Characteristics: For a set of measures $\{\mu_i\}_{i=1}^N$, we write the weighted Wasserstein barycenter μ_*

$$\mu_* = \inf_\mu \sum_{i=1}^N w_i W_2^2(\mu, \mu_i) \qquad (15)$$

with nonnegative weights $\{w_i\}_{i=1}^N$ such that $\sum_{i=1}^N w_i = 1$. Note that each Wasserstein distance itself involves an infimum. Existence and uniqueness of solutions to Eq. (15) appear to have first been studied by [91]. They also provide a dual formulation and provide solutions to the problem when each μ_i is zero mean Gaussian.

Computation: Unfortunately, the Wasserstein Barycenter inherits the usual computational cost associated with optimal transport distances. One way to circumvent this issue is to consider entropy regularised divergences (see § 4.1) instead of W_2^2, so that the Sinkhorn algorithm can be exploited. Regularisation in this manner is considered by [92]. They show that in the Gaussian case, entropic regularisation with uniform reference measures Eq. (7) leads to a blurry (or biased) Barycenter, but entropic regularisation of the form Eq. (9) does not suffer from this problem.

The semi-dual formulation of OT (see Eq. (25)) can be used to cast Eq. (15) as a minimum over a minimax problem [93]. Using an input convex neural network (which are universal approximators in the space of convex functions) allows one to search over the set of convex functions. This approach suffers from the usual theoretical intractability of optimisation problems associated with neural networks, but in practice the method appears competetive on toy problems. The computational complexity per training iteration is $\mathcal{O}(NMp)$, where N is the number of marginal distributions, M is the batch size and p is the size of the neural network.

Applications: [93] represent the Barycenter as a generative model that resembles a WGAN. Due to this, they may generate infinitely many samples from the Barycenter. [93] generate samples from toy Barycenters of MNIST images and recoloured high resolution images, however other applications should be possible.

4.5 Gromov-Wasserstein (GW)

Motivation: OT (2) assumes a common meaningful distance m exists for the points in the source and target measures [86]. Thus, implicitly assuming that the two supports lie in the same metric space. This can be limiting in many applications such as cross-domain alignment [94] and unsupervised bi-lingual lexical induction [95] where we seek an alignment between sets of word embeddings from two *different* languages without access to parallel data. This calls for *relational OT formulation* that measures how distances between pairs of words are mapped across languages.

In the lack of a common distance function, d one might seek a solution that relies on the relevant distance functions $d_{\mathbb{X}}, d_{\mathbb{Y}}$ (i.e. metric spaces) corresponding to the measures

considered. Then, define the transportation cost based on the relationships between distances. This turns the problem into a notion of distances between metric spaces [96]. Notably, this offers a way for comparing originally incomparable spaces. Additionally, this notion makes the resulting alignment invariant with respect to arbitrary distance-preserving transformations (e.g. rotations).

Formulation and Characteristics: Gromov-Wasserstein [97] generalizes OT distance to a notion of *distance between distances* [95]. The $\mathcal{GW}$ is defined as

$$\mathcal{GW}(\mu, \nu)$$
$$= \inf_{\pi \in \Pi(\mu, \nu)} \int \int_{(\mathbb{X} \times \mathbb{Y})^2} r(x, x', y, y') d\pi(x, y) d\pi(x', y'), \quad (16)$$

where $r(x, x', y, y') = d(d_{\mathbb{X}}(x, x'), d_{\mathbb{Y}}(y, y'))$ is the relational distance. When $d = L_2$, $\mathcal{GW}^{\frac{1}{2}}$ is a proper distance [97].

Computation: Naively solving the $\mathcal{GW}$ formulation Eq. (16) involves prohibitive computation ($O(N^2 M^2)$) involving fourth-order tensor product [95]. Recent techniques for scaling $\mathcal{GW}$ to big graph and point cloud datasets are discussed in §5.

Applications: On the one hand, $\mathcal{GW}$ in general bypasses a classical OT limitation and allows alignment across "incomparable" spaces [98], [99] (e.g. spaces of different dimensionality or data type). Given this, it has found increased adoption in cross-domain(heterogeneous domains)alignment. In single Cell Omics [100], $\mathcal{GW}$ was used to align heterogeneous measurements from gene expression and Chromatin accessibility. In [72] partial $\mathcal{GW}$ was shown to be efficient for matching point clouds when they exist in different domains and have different features. Very recently, $\mathcal{GW}$ was leveraged in Cross-Domain Imitation Learning [101]to align and compare states between the different spaces of RL agents. A Gromov extension of Dynamic Time Warping [102] was used to align incomparable time series by leveraging the intra-relational geometries. An example is multi-modal data such configuration space of a robotic arm and its representation as pixels of a video frame.

On the other hand, leveraging within-domain similarity ($d_{\mathbb{X}}$ and $d_{\mathbb{Y}}$) allows capturing the topological relations. This can be beneficial for learning relations.

Such a relational perspective is also a reason why $\mathcal{GW}$ is used in graph applications. $\mathcal{GW}$ has thrived applications such as graph classification [103], graph partitioning [104], graph node embedding [105] and supervised graph prediction [106] and others.

Limitations: Relying only on the topological or relational aspects encoded by pairwise similarities during alignment and ignoring the features is one of the limitations of GW (check Fused Gromov Wasserstein below). Computationally, GW turns the OT linear program into a non-convex quadratic program with a prohibitive computational cost. From a computational perspective, as noted earlier the naive implementation of GW requires operating on a fourth-order tensor. This can be of less concern only in a few cases. For example, when working on small-to-moderate data size [95]. Also, for certain classes, [107] the cost can be brought down to $\mathcal{O}(n^3)$ where n is the number of points. Going beyond these, GW can be expensive and the search for ways to speed up the calculation is very active. More discussion about this is provided in Scaling OT (§5).

4.6 Hierarchical OT

Motivation: One might be interested in encouraging OT to consider the inherent structure of the measures, such as cluster structure when doing the transportation. We may aggregate points within a subgroup based on a common feature (e.g. class label) or pre-specified structures and then transport these subgroups. As the number of subgroups to be matched is typically much less than the number of individual points, computational efficiency is gained.

Formulation and Characteristics: There is a diversity of formulations all called Hierarchical OT. In most cases [108], [109] the problem is posed as a nested version of OT by simultaneously solving local assignment problems (e.g. per cluster) in the outer level and a global assignment problems (e.g. across clusters) in the inner level. Decomposing the problem into simpler ones can make optimization easier [108]. In this nested setup, one can use arbitrary formulations at different levels.

Applications: Problems that deal with hierarchical data (e.g. document matching, datasets alignment, and domain adaptation) represent a good playground for hierarchical OT formulations. HOTT [16] leverages the natural hierarchy of documents (words → topics → document) to speed up the OT-based document comparison. The document distance is computed as the 1-Wasserstein distance between the distributions of documents supported on topics. The hierarchical structure comes from the fact the ground cost in the formulation is the Word Movers Distance WMD [14], another OT distance that quantifies the minimal transport cost between multidimensional word embedding vectors. Hierarchical OT for Multimodal Distribution Alignment [108] frames the domain adaptation as a nested two-level transportation. The first level is among the samples in pairs of clusters while the second one is across the clusters themselves. Interestingly, a scalable ADMM algorithm was proposed. It splits across cluster pairs and allows the computation to be done in parallel. The nested transportation strategy was adopted in other domain adaptation settings such as unsupervised domain adaptation [110] and semi-supervised domain adaptation [111].

Another problem that lends itself to hierarchical formulations is quantifying dataset distances [109]. In principle, it is very similar to previous domain adaptation (i.e transportation between classes clusters and another transportation between clusters samples). Unlike domain adaptation, the dataset distance is model agnostic. Notable advantages of using OT here are the possibility of comparing datasets even when the label sets are disjoint and the interpretability conveyed by the optimal matching.

5 SCALING OT COMPUTATION

In this section, we discuss various approaches for scaling OT computation to handle large datasets and high dimensional data. When we compute OT-based discrepancy measures between distributions μ and ν, we usually don't have the explicit forms of these distributions. Instead, we use the empirical approximations $\mu_n = 1/n \sum_i \delta_{x_i}$ and $\nu_n = 1/n \sum_j \delta_{y_j}$ constructed from the random samples $\{x_i\}_{i=1}^m \sim \mu$ and $\{y_j\}_{j=1}^n \sim \nu$. An OT estimator constructed in this straightforward and popular strategy is known as
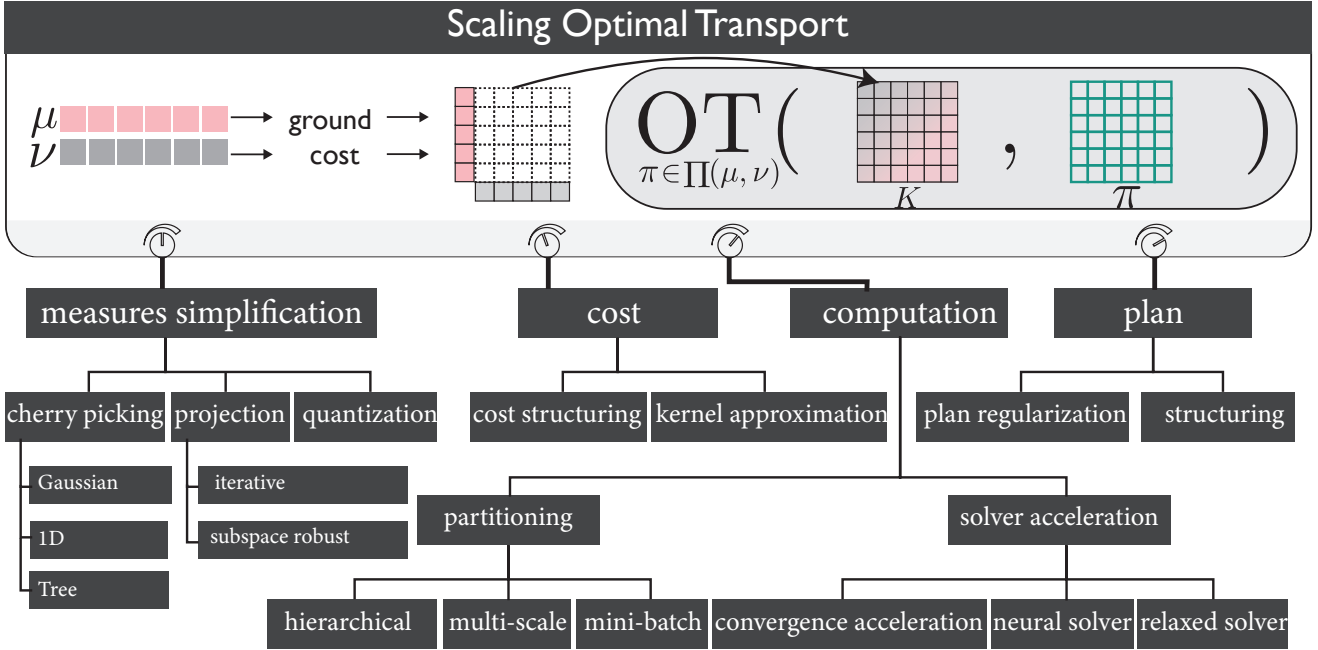
Fig. 5. **OT Scaling.** Exaggerating a bit, we can scale all OT methods by turning one of the knobs (i.e. optimizing components) of the OT computation machinery. This can be done through (§5.1 **measures simplification**) simplifying the input measures, (§5.3 **cost**) structuring the ground cost or approximating the associated kernel, (§5.2 **plan**) enforcing a specific prior on the optimal plan or (§5.4 **computation**) accelerating the optimization. Of course, it is also possible to turn multiple knobs simultaneously and have (§5.6)a **combined** approach.

the *plug-in estimator* [112], [113]. A natural question that arises is how far the plug-in estimate $OT(\mu_n, \nu_n)$ is from the true solution $OT(\mu, \nu)$. The work of Dudley [114] and others [24] show that the plug-in approximation converges with rate $n^{-1/d}$, where d is the points' dimension. Niles-Weed *et al.* [115] notes that this can not be improved without further assumptions. A short proof of this rate is available [116]. Such a convergence rate makes OT-based discrepancy estimation in high dimensions seem hopeless. We need a number of samples n that grows exponentially as d increases before we get a decent approximation accuracy. On the other hand, we know that solving OT problems for a large number of samples presents a computational hurdle. The linear program solver is super-cubic in n and the per-iteration cost of the basic Sinkhorn algorithm is $O(n^2)$. Thus, the statistical and computational challenges of OT are somewhat intertwined, which is why some techniques address the two hurdles simultaneously [44].

So, how do we scale OT? Looking at the big picture of the literature for scaling OT *computation* to large d and/or n, one can argue that all the proposed approaches are either approximating the OT solution or tailoring it to a specific case that is easier to handle. The concrete techniques to achieve this can be categorized into five main themes. In Fig. 5 top, an abstract representation of OT computation is illustrated. It is possible to guess where the optimization can happen. One can work on the individual components corresponding to the *measures*, the *plan*, and the *ground cost* or accelerate the *formulation optimization* by, for example, partitioning. *Combining* some of these approaches can help and some papers fall in more multiple groups.

5.1 Measures Simplification

A logical first stop for scaling OT is to consider simplifying the input to the OT computation pipeline; the measures. Such a simplification can take the following forms.

Cherry Picking (measures with closed form solution): The simplest form of measure simplification is picking from a special set of measures that have closed form solutions. As discussed earlier, the one dimensional case has a closed form (Eq. (11)). The whole line of Sliced OT (§4.3) builds on this fact. The closed form solution for the 2-Wasserstein metric for multivariate Gaussian distributions with means m_1, m_2 and covariance matrices Σ_1, Σ_2 is given by:

$$W_2(\mathcal{N}_1, \mathcal{N}_2)^2 =$$
$$\|m_1 - m_2\|_2^2 + \mathrm{Tr}\left(\Sigma_1 + \Sigma_2 - 2\left(\Sigma_1^{\frac{1}{2}}\Sigma_2\Sigma_1^{\frac{1}{2}}\right)^{1/2}\right) \quad (17)$$

where $\Sigma_1^{1/2}$ is the unique positive definite square root matrix B_1 such that $B_1 B_1 = B_1^T B_1 = \Sigma_1$.

Recent fundamental investigations have revealed the existence of closed-form solutions for regularized OT on Gaussians [117], barycenter of Gaussians [118], [119] and the entropic Gromov-Wasserstein on Gaussians [120]. Another case is the tree Wasserstein for measures that can be supported on a tree, used mostly in document classification tasks [121] as it permits fast comparison of a large number of documents. Formally, if the measures μ and ν supported on a tree $\mathcal{T} = (\boldsymbol{V}, \boldsymbol{E})$ whose root is v_1, then the closed form tree-Wasserstein distance [121] is:

$$W_{d_{\mathcal{T}}}(\mu, \nu) = \sum_{v \in \boldsymbol{V} \setminus \{v_1\}} w_v \left|\mu(\Gamma(v)) - \nu(\Gamma(v))\right|, \quad (18)$$

where w_v is the length of the edge from v to its unique parent node and $\Gamma(v)$ is the set of nodes contained in the

subtree of $\mathcal{T}$ rooted at $v \in \boldsymbol{V}$. Tree-Wasserstein can be computed in linear time with respect to the number of nodes in the tree [31], [122]. Note that this ignores the computational and storage requirements for transforming measures into trees. Noteworthy variants include a sliced version [123] calculated by averaging the Wasserstein distance between the measures using random tree metrics, a fast supervised version that can be trained end-to-end [124], fixed-support Wasserstein Barycenter on a tree [122] and an efficient solver [31] for unbalanced OT with the capability of processing one-million point measures on CPU.

Since the OT solution for these favorable cases is analytically known, they are employed in many applications. Additionally, they are used to alleviate computational burden of other OT formulations [109]. Some other usage examples of simple OT problems with analytic solutions include being used as test cases for establishing the correctness of OT solvers [125], and for providing pre-training data for supervised neural OT solvers [126]. Despite the advantages, cherry-picked measures can be unrepresentative of real-world data. Also, solving certain OT problems in higher dimensions can still be expensive even for closed-form situations. For example, calculating the barycenter of Gaussians [127] incurs a cubic complexity in the dimension [128].

Measures Projection: Low-dimensional projection is another conceptually simple candidate for addressing the curse of dimensionality. The recent findings of spiked transport model [115, Theorem 1] provides a theoretical grounding for this approach. According to the model, if the two measures differ only on a low k-dimensional subspace, then the plug-in estimator convergence rate can be improved from $n^{-1/d}$ to $n^{-1/k}$. Projection robust methods family resides in this research direction. Subspace Robust Wasserstein (SRW) [129] builds on the principles of Sliced Wasserstein (SW §4.3) and extends the projection to a subspace of dimension $k \geq 2$. SRW seeks the subspace that maximizes the OT cost between two measures to capture the major discrepancy (similar in spirit to max-SW §4.3). SRW shares some properties with the 2-Wasserstein distance while being more robust to random perturbations in data. In Wasserstein Barycenter, [130] showed that projecting the distributions into low dimensional spaces can provably preserve the quality of the barycenter. Projection robust Wasserstein barycenter [131] solves the fixed-support discrete Wasserstein barcenter problem.

One issue when working with subspaces is that the subspace-plan is optimal between the projected and not the original high dimensional measures. This can be addressed by considering only the plan on the original space that is constrained to be optimal after the projection. This is the key idea behind the concept of *subspace detours* that was introduced by [132] and later extended to other formulations such as Gromov-Wasserstein in [133].

In cases where the input measures can be grouped and and the grouping is known a-priori, one can further speed up SRW and gain additional advantages. Feature Robust OT (FROT) [134] speeds up OT computation and makes it more robust to noise. While similar in spirit to SRW, FROT is convex and more scalable than SRW.

Another projection-based family of methods uses iterative random projection where informative projections are determined sequentially. In this category, there is the Projection Pursuit Monge Map (PPMM) method [135] that picks the best projection in the kth iteration guided by information (e.g. residuals) from $k-1$ projections. *Check [136] for a focused tutorial on iterative projection approaches.*

Measures Quantization: Another simplification approach is to quantize the measures before computing OT. Examples include [137], [138] over-sampling the input measures and applying the solver only on their summary (i.e. a k representative samples acquired by k-means like quantization). Usually, it is easy to obtain more samples when the data come from large datasets or generative models. The approach is solver agnostic in the sense it can be applied to different formulation (OT vs regOT) with no change [137]. Along the same line, a simplification that considers transporting the measures centroids (instead of the actual measures) can be more efficient. For example, Word Centroid Distance [14] speeds up the OT computation by representing each document as the weighted average of its word embedding.

In addition to the approaches discussed for measure simplification, *smoothing* is being explored. Smoothing simplifies measures by adding Gaussian noise. Gaussian Smoothed OT (GOT) [139] simply convolves the input measures with isotropic Gaussian kernel before performing the computation. Research investigation into the theoretical properties of the formulation [140] and potential applications [141], [142] is ongoing. A notable trait is alleviating the curse of dimensionality by enjoying a convergence rate $n^{-1/2}$ in all dimensions [143].

5.2 Structuring the Optimal Plan

Going beyond input pre-processing and looking into the OT optimization itself, Eq. (4), OT computation involves an optimization over the infinite-dimensional set of couplings $\pi \in \Pi(\mu, \nu)$ between the measures μ and ν. Enforcing a structure or regularization on the optimal plan can often lead to structural changes in the optimization problem making it simpler to solve. For example, searching for the transport plan only in a low-rank sub-space [144], [145] can be used to reduce the dimension of the problem, speeding up the optimization process.

Plan Regularization: The regularized OT formulation (§4.1) is the most famous group of plan structuring methods. The entropy regularizer (Eq. (7)) biases the optimal coupling towards low entropy couplings. The algebraic properties of the added KL term enable arranging the unique solution P_λ in a simple form. This form can be solved using successive matrix scaling iterations. Matrix scaling iterations are significantly faster than using linear programming for solving the original OT problem. In other words, regularization changes the structure of the optimization problem, enabling the use of a faster algorithm. However, this speedup comes with a limitation. In regularized OT, there is a trade-off between convergence speed and precision, controlled by λ. As reported by many studies [146], larger λ results in a faster convergence but poor approximation. On the other hand, very small λ leads to numerical instability. To obtain better approximation, recent works investigate gradual reduction in λ [147], [148]. This simulates gradually approaching the exact (unregularized) OT which is recovered with $\lambda \to 0$.

Low-Rank Structuring: Another family of algorithms uses the idea of low-rank approximate factorization of the plan. In this family, the plan is decomposed into a product of low-rank matrices, improving the statistical stability of OT and reducing the sample complexity. Factored Coupling (FC) [149] was the first to impose a low-rank structure on the transport plan. This is a popular strategy in high-dimensional statistics [24]. Their *low-rank transport plan* is an approximation of the optimal coupling that builds on the assumption that μ and ν may only vary across a few dimensions. FC transports factored partitions or "soft clusters" (in which the clusters can include fractional points) of the measures and the problem is solved as regularized 2-Wasserstein barycenters. Building on this, the authors of the Latent OT (LOT) method [150] observed that the low-rank transport plan is also more resilient to outliers and noise. A low-rank plan transports measures through common "anchors" and is observed to be better at preserving clusters and tolerating outliers. In FC, the anchors are shared between the measures, and their number is dictated by the rank r. To improve the interpretability of the plan, LOT allows each measure to have a different number of anchors. However, the anchors are additional hyperparameters that need to be chosen carefully. A generalization of FC is the Low-Rank Sinkhorn Factorization method [144]; a formulation compatible with all ground costs and not only the squared Euclidean distance. In the Low-Rank Sinkhorn Factorization formulation, the coupling P is explicitly factorized into a product of two sub-couplings Q and R linked by a common marginal g, $P = Q \cdot \mathrm{diag}(g)^{-1} \cdot R^T$. The solution is obtained by jointly optimizing Q, R, and g. The demonstrated computational improvement of the factorization of the transport plan in various applications motivated a further investigation [151] into its theoretical properties. A debaised version of the Low-Rank Sinkhorn Factorization method [144] was introduced in [151] and was shown to interpolate between Maximum Mean Discrepancy [26] and OT.

An obvious issue with these approaches is the poor approximation when the low-rank assumption is violated (for example, when μ can be obtained by permuting the support of ν). To attempt to circumvent this limitation, Liu *et al.* [145] proposed OT approximation in which the transport plan can be decomposed into the sum of a low-rank and a sparse matrix as a possible solution.

5.3 Ground Cost Approximation

Algorithm 1: Sinkhorn(K, a, b, δ)

---

1 **Inputs:** K, a, b, δ, u
2 **repeat**
3     $v \leftarrow b/K^T u, \ u \leftarrow a/Kv$
4     $r \leftarrow \|u \odot Kv - a\|_1 + \|v \odot K^T u - b\|_1$
5 **until** $r < \delta$;
   Result: u, v

Application-wise, the fact that the OT distance admits a customizable ground cost has been leveraged in many applications. Since we know that the OT cost is shaped by the structure of the ground cost m, we can control

the transport plan by introducing an application-specific structure in m. For example in domain adaptation, we can bias similar (e.g. same-class) source domain points to be transported together to the target domain without splitting, e.g. using a submodular cost [156]. Computationally, having a structured cost (perhaps one that can be decomposed) can be leveraged to improve the execution time of OT.

Structuring Ground Cost: Enforcing a structure on the ground cost can be seen as an implicit way of reflecting that structure on the optimal plan. For example, using a cost of infinity (or a large default value) for the ground cost between 2 points guarantees that the final plan will have no transport between these 2 points. So a structural infinity in the cost matrix maps to a structural zero in the optimal plan which can be used to sparsify the optimal plan. This relationship between the cost and the plan can be exploited to speedup the OT computation reducing the memory and time requirement of the various algorithms. Relatedly, in applications that deal with a large number of samples (e.g. document and image retrieval) such that storing the pairwise distances $M_{i,j}$ for all i and j samples can be prohibitive, a useful structure that can be used is the *ground cost saturation*. Specifically, we threshold the ground cost $M_{i,j} = M_{max}$ if $m \notin r$ where r denotes the relative (nearest) samples. From a bi-partite graph matching (between the measures) perspective, this is equivalent to constraining edges connectivity and it can result in significant savings. The approach was used by [15] and [157] for speeding up document and image retrieval; respectively. In Multi-marginal OT (MOT), the structure of the cost tensors is exploited to speed up the computation. In sensor fusion and tracking, [158] shows that the cost functions enjoy a structure that allows sequential decoupling, central decoupling or a combination of both. Computing the Sinkhorn projections in these cases can be done efficiently [159].

Kernel Matrix Approximation: another popular strategy is structuring the ground cost M such that the dependent kernel matrix $K(M)$ can be approximated efficiently. The Sinkhorn algorithm uses the kernel matrix $K(M)$ by multiplying it (or its transpose) by 2 vectors in each iteration. In practice, K can be very large (two 256×256 gray images require 32GB memory for storing kernel with double percision [160]). Understanding kernel approximation techniques discussed next requires a brief interlude into the Sinkhorn algorithm. In Figure. 6 we have the measures μ and ν supported on X and Y respectively with corresponding weights of a and b. The kernel matrix capturing the measures dissimilarity is denoted as $K_{ij} = e^{-\frac{M_{ij}}{\lambda}}$ where M is the ground cost matrix and λ is the regularization coefficient in (Eq. (6)). The algorithm seeks the unique optimal solution that must satisfy the following:

$$P_\lambda = \mathrm{diag}(u) K \mathrm{diag}(v) \qquad \text{(Sinkhorn theorem)}$$

$$P_\lambda \mathbb{1}_m = a \ , P_\lambda^T \mathbb{1}_n = b \qquad \text{(optimality condition)}$$

Finding the vectors u and v that satisfy the requirements above is achieved simply by performing repeated alternate scaling (line 3 in Alg. 1) until convergence. The number of iterations needed to converge to a specific tolerance δ is controlled by the scale of elements in M relative to λ [161]. The iterations are merely applying kernel matrix K (or its
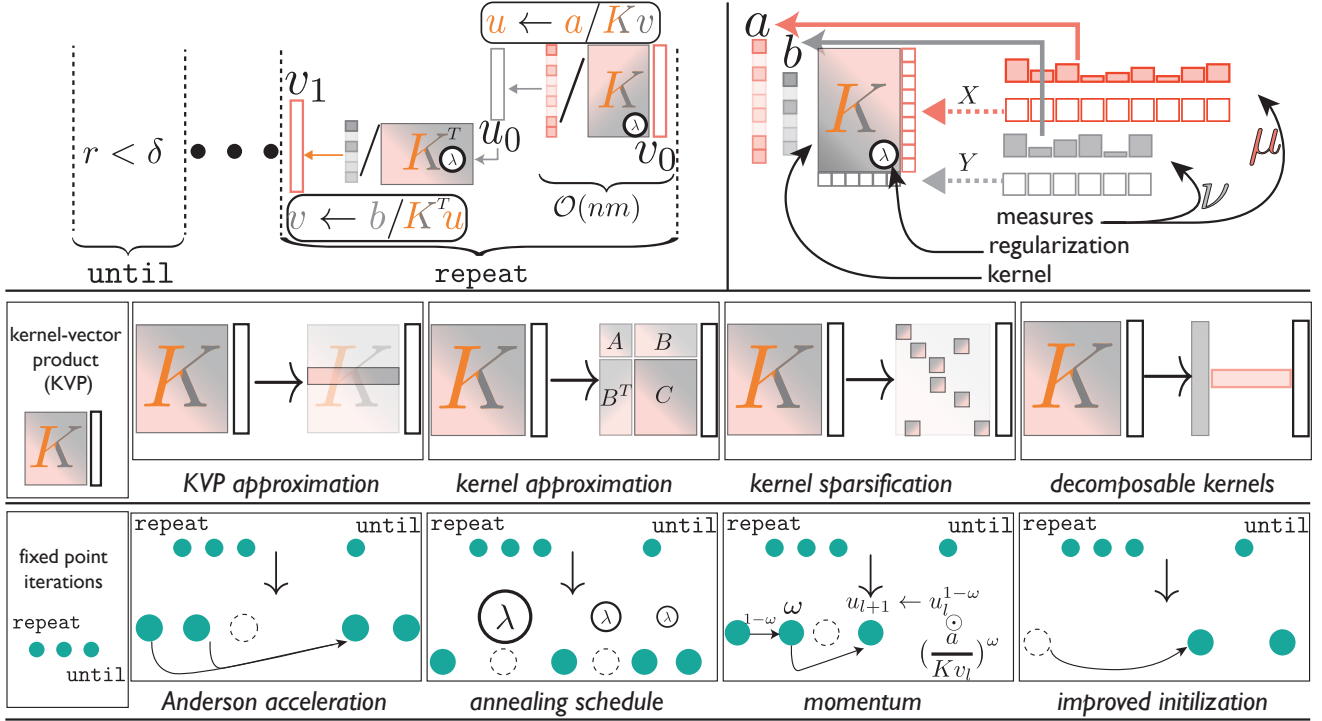
Fig. 6. **Speeding up Sinkhorn.** Sinkhorn procedure (top row) basically applies kernel (K) vector (u, v) product operations till convergence. Given this, Sinkhorn can be accelerated through a faster kernel-vector product (KVP) or a reduced number of iterations (bottom row). Example techniques for improved kernel vector product (middle row) include KVP approximation [44], kernel approximation [44], kernel sparsification [152], or decomposable kernels [48]. Fixed point iterations can be accelerated (discussed in §5.5) using methods such as Anderson acceleration [153], annealing schedule [147], momentum [154] and improved initialization [155].

transpose) to vectors u and v. The computational cost of the whole procedure is driven by this and a naive matrix-vector multiplication will cause each iteration to have a quadratic cost. It is natural to think of reducing the computational burden of the matrix-vector product.

Reducing the computational cost of kernel matrix-vector product is a common problem with a number of known workarounds, often used in kernel-based machine learning algorithms such as Gaussian processes. In OT, the same computational tricks were adopted by some works. Nys-Sink [44] combines Nyström approximation [162] with Sinkhorn scaling. LCN-Sinkhorn [152] follows in the same direction and notes that the exponential used in K makes its entries negligibly small everywhere except at each (support) point's closest neighbors. This motivates sparse approximation of K where only nearby neighbors are accounted for. This is equivalent to approximating the cost matrix M using a "generalized sparse" cost matrix M^{sparse} where $M_{ij}^{\text{sparse}} = \infty$ for "far" points, i.e. the structural zeros of the sparse matrix have a default value of ∞. Filtering "near" points is done via Locality Sensitive Hashing (LSH) [163]; points within a certain distance r_1 are much more likely to be assigned to the same hash bucket than points with a distance $r_2 > r_1$. Ultimately, this yields a version of Sinkhorn that can scale log-linearly with the number of points, (i.e. $\mathcal{O}(n \log n)$). Scetbon *et al.* [48] considers the ground cost function of the form $m(x, y) = -\log\langle\phi(x), \phi(y)\rangle$ where ϕ is a map from the ground space onto the positive orthant $\mathbb{R}_+^r$, with $r \ll n$. This results in a decomposable kernel (K in algorithm 1) making sinkhorn computation more efficient

with iterations complexity $\mathcal{O}(nr)$.

5.4 Computation Partitioning

Computation Partitioning refers to breaking OT computation into smaller OT problems where the outcomes of the sub-problems is aggregated into a solution of the original OT. Check the top panel of Fig. 7 for visual representation. Computation Partitioning is done in the following ways:

Multi-scale, recursive or hierarchical: a number of similar methods were proposed in the literature relying on arranging the points in a tree structure of some sort. "Multi-scale" OT was proposed in [164], [165]. In [164], the transport problem is decomposed into a sequence of problems with varying scales and solved in a top-down fashion (i.e. coarse-to-fine scale). These methods proved to be useful when the ambient dimension d is small. Yet they fail to scale to high dimensional settings as they utilize space discretization.

In "recursive" OT, a similar divide-and-conquer strategy is used to perform recursive partitioning of the measures/data. Examples in this category are the recent techniques for scaling GW to big datasets [104]. The general theme is to partition the measures into smaller blocks, match the blocks based on their representatives (e.g. centroids), then match the paired blocks by proceeding in a recursive manner as needed (Fig. 7 top). MERC [166] adopts this approach using black box clustering for blocking and cluster centers as block representatives. Quantized Gromov-Wasserstein (qGW) [167] supports this simple approach with theoretical guarantees based on [168] and proposes an algorithm that scales to datasets containing up to 1
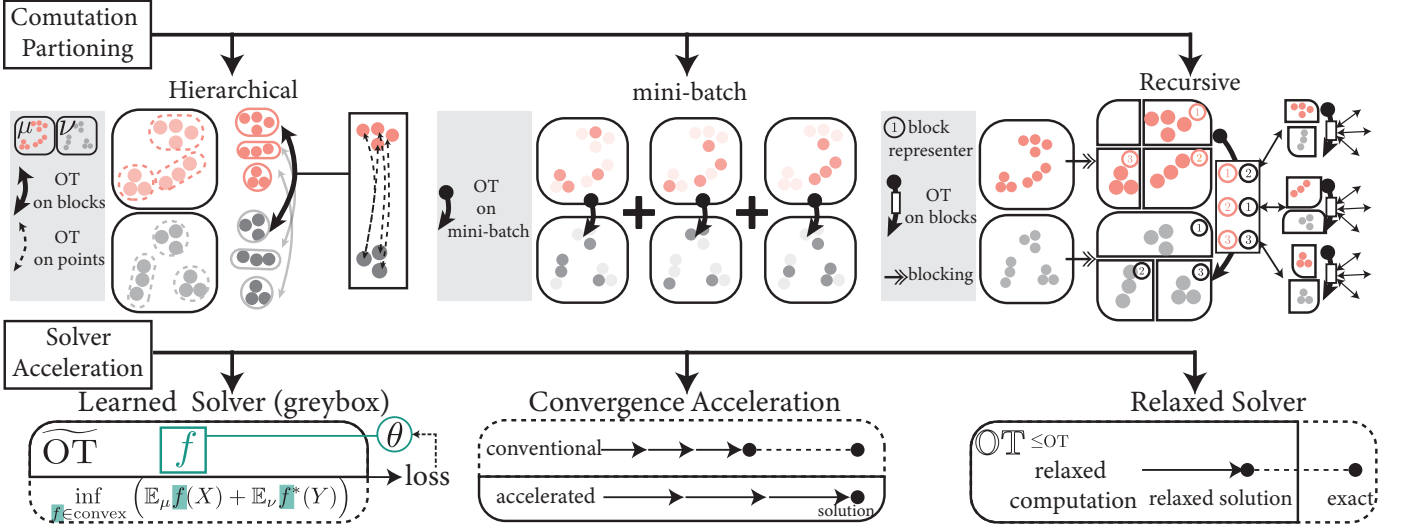
Fig. 7. **Computation Partitioning and Solver Acceleration in OT**.

million points. It should be noted that the blocking the representative's choices (e.g., graph data community detection algorithms [169] and maximal PageRank [170]) are crucial and careful selection of the methods might be needed depending on the data.

In "hierarchical" OT (§4.6), a similar tree structure is adopted but the structure is more intrinsic to the data itself. For example, measuring documents similarity using OT [16] can leverage the natural hierarchical structure where documents are modeled as distributions over topics which, in turn, are modeled as distribution over words. Thus, nested computation of documents distances can be performed where the outer level of the computation is over the topics. Noting that the number of topics (or clusters) is much less than that of the words, the computation will be greatly reduced thanks to the hierarchy. Similarly, aligning datasets (potentially of different modalities) using OT [108] can be enhanced as the datasets usually admit clustered structure.

Mini-batch: Sommerfeld *et al.* [171] shows that the average OT distances (computed by any black-box solver) on random sub-sampling of the input measures is a fast approximation of the exact OT. Theoretical non-asymptotic guarantees reveal that the approximation error can be independent of the size of the problem in some cases. This surprisingly simple approach can scale very well in practice using one percent of the exact computation in images [171]. Recent OT acceleration strategies along the same line take inspiration from deep learning. Mini-batch OT performs the computation on mini-batches in a way that mimics the typical training of deep neural networks. The empirical success of mini-batch OT motivated recent theoretical investigation [172], [173], [174]. One concern here is the mini-batch strategy can lead to an undesirable smoothing effect for which Fatras *et al.* [175] suggests relying on unbalanced OT formulation to counter the issue.

5.5 Solver Acceleration

Solver acceleration targets optimizing the computation process either by learning to predict the solution quickly (usually for specialized problems) or admitting a relaxed version of the original problem which is easier to compute. In this category, we identify *Convergence Accelerations* , *learned solvers* and *relaxed solvers*. A depiction is shown in Fig. 7 bottom panel.

Convergence Acceleration. Convergence Acceleration aims at bringing the solver to convergence in fewer iterations. In regularized OT, for example, we can use a number of strategies to accelerate the Sinkhorn algorithm (see Fig. 6 bottom row). The Sinkhorn algorithm (Alg. 1) is fundamentally a fixed point iteration (FPI) algorithm for finding fixed points (u^*, v^*). Therefore, all FPI algorithm acceleration techniques can be investigated. In [153], the classical Anderson acceleration for FPI algorithms was used. Anderson acceleration can speed up FPI algorithms by making them more robust to residual oscillation. Besides Andreson acceleration, the over-relaxed Sinkhorn method [154] adopts a relaxed FPI algorithm where the auxiliary vector update $u_{l+1} \leftarrow \frac{a}{Kv_l}$ is replaced with the relaxation $u_{l+1} \leftarrow u_l^{1-\omega} \odot \left(\frac{a}{Kv_l}\right)^{\omega}$, where l denotes the iteration and $\omega > 0$ is the chosen relaxation parameter. [154], [176] show that for $\omega > 1$ the algorithm can converge faster than traditional Sinkhorn. Recently, [177] identified an apriori range for ω for which global convergence is guaranteed. Other fixed point acceleration techniques were investigated in [178]. In the Sinkhorn literature, better initialization has been mostly neglected as the potential speedup seems insignificant given the convex nature of regularized OT. However, a recent work [155] showed that careful initialization of the Sinkhorn subroutine (in a neural network) initialized by pre-training on problems with closed-form solutions (Gaussians and 1D) can be accelerated. The technique can also be complementary to previous techniques. Also, for smooth approximation of OT problems (e.g. regularized OT) it is possible to use advanced optimization techniques such as accelerated gradient descent [179], [180], L-BFGS [59], [181] and FISTA [59]. The previous approaches have little in common apart from seeking faster convergence. Yet, they all are holistic approaches that consider the inner working of the solver and a potential revisit of the original formulation [179] before formulating an acceleration strategy.

Learned Solvers: A promising direction that is accumulating interest recently is based on the idea of "neural solvers". Neural solvers are neural networks that attempt to predict the OT output from the input measures. Data-driven neural solvers reduce the expensive traditional computation by leveraging knowledge from past OT solutions. By constructing a model and training it on ground truth OT solutions and parameters, we can predict OT solutions during test time very quickly (single forward pass) without working it out from the measures from scratch (iterative) every time. The approach becomes more attractive in deployment scenarios that requires solving OT repeatedly (e.g. coupling stream of image pairs). In addition to the computational improvements, a notable advantage of neural solvers is the ability to cope with parametric continuous μ and ν (densities) [182], [183]. Empirically, neural solvers show great success in generative modeling (§B.2) and domain adaptation (§B.1).

Grey-box Solvers: These approaches replace a component of the OT formulation by a neural network reparameterizing the optimization to be in terms of the weights and biases of the neural network. Some of these semi-transparent models strike a good balance between the regular OT computation, which can be expensive and hard to scale, and the data-rich (yet opaque) black-box learned solvers. So far in the literature the Kantrovich-Rubinstein duality Eq. (26) and the semi-dual formulaion Eq. (20) [18], [184] have received the most attention. This is because the (semi-)dual solution is comprised of 2 pointwise solutions: a dual value for each x in the support of the measure μ, and another dual value for each y in the support of the measure ν. This pointwise structure motivates the use of a neural network to replace the mapping from x and/or y to their respective dual solutions. However, there are formulation tricks that need to be employed here to get rid of any constraints or non-smoothness in the (semi-)dual formulation. Current works alleviate this issue either by relying on soft constraints [183] or structuring the neural solver to ensure feasibility [18], [185].

Grey-Box Solvers with Soft Constraints: To understand the strategy of these solvers, we revisit the Kantrovich dual:

$$\sup_{\phi \in L^1(\mu), \psi \in L^1(\nu)} \mathbb{E}_{(X,Y) \sim \mu \times \nu} \big[\phi(X) + \psi(Y) \big]$$
$$\text{subject to } \phi(x) + \psi(y) \leq c(x,y) \text{ for all } (x,y). \quad (19)$$

The formulation shows the OT computation as maximizing an expectation term thus suggesting that the maximization can be achieved using stochastic gradient methods on sampled batches from the coupling $\mu \times \nu$. However, it is not clear how the constraints on ϕ and ψ can be satisfied in this scheme. It turns out that the dual of the entropy regularized OT problem is unconstrained and doesn't exhibit this problem, Eq. (6), which can be obtained by the Fenchel-Rockafellar theorem:

$$\sup_{\phi, \psi} \mathbb{E}_{(X,Y) \sim \mu \times \nu} \big[\phi(X) + \psi(Y) + R_\lambda \big] \quad (20)$$

where $R_\lambda = -\lambda e^{\frac{1}{\lambda}(\phi(x) + \psi(y) - c(x,y))}$ is a penalty term. Intuitively, R_λ represents a smooth version of the constraints in Eq 19. Thus, we bypass the need to enforce the hard constraints. Given that the problem is now unconstrained,

one can simply parameterize the dual variables ϕ and ψ using neural networks and use stochastic graident methods to maximize the objective in (20) [184].

Alternatively, the semi-dual formulation Eq. (20) of the unregularized OT problem is also unconstrained but its objective is non-smooth, so a smooth approximation of the semi-dual formulation of the unregularized OT problem can also be used. The semi-dual's objective is not smooth because it uses the max operator. One way to alleviate this is to construct a smooth approximation of the max operator using the conjugate function of some strictly convex regularization/penalty function Ω over the unit Simplex Δ. The entropy or L2 norm are possible choices for Ω. The following function is known as the conjugate function of $\Omega(x)$:

$$\max_\Omega(x) = \sup_{y \in \Delta} x' \cdot y - \Omega(y). \quad (21)$$

The above conjugate function is smooth and differentiable in x and is equal to $\max(x)$ when $\Omega(y) = 0$ for all y (i.e. no regularization). Replacing max with $\max_\Omega$ for some choice of Ω in the semi-dual results in a smooth unconstrained optimization problem. The gradient of the above function can also be analytically derived for common choices of Ω.

Grey-Box Structured Solvers: These solvers rely on structural constraints in the neural model itself (e.g. specific layers, training strategies, etc). Thus, we bypass the hard constraints by simply relying on the constrained structure of the neural network. A prominent example here is W-GAN solver [18]. The potential f_w (where $w \in W$) in Wasserstein-1 formulation Eq. (29) has to satisfy K-Lipschitz constraint. This is achieved by forcing compactness on the parameters space W. In practice, this is achieved by weight clipping after each gradient update [5]. Another way to import structure that ensures constraints feasibility is to use special layers. What gives rise to this approach is the formulation in Eq. (19) which allows using Input Convex Neural Networks (ICNN) [188]. Specifically, one can rephrase Eq. (19) as an optimisation problem over the space of convex functions while eliminating the distance constraints. This trick appears in [185], [189], although elements of the convexification trick first appear in Villani's work [36]. Reparameterising the problem in terms of convex functions allows one to leverage a deep learning approach to solving OT problems. The formulation is given by

$$C_{\mu,\nu} - \inf_{\phi \in \mathsf{CVX}(\mu)} \Big[\mathbb{E}_\mu \phi(X) + \mathbb{E}_\nu \phi^*(Y) \Big], \quad (22)$$

where $C_{\mu,\nu} = \frac{1}{2} \Big(\mathbb{E}_\mu \|X\|_2^2 + \mathbb{E}_\nu \|Y\|_2^2 \Big)$ and $\mathsf{CVX}(\mu)$ is the space of all convex functions integrable with respect to μ. The space $\mathsf{CVX}(\mu)$ can be parameterised by input convex neural networks (ICNNs). Generally, one can compose ICNN from two principal blocks: linear blocks consisting of linear layers, and convexity preserving blocks consisting of linear layers with non-negative weights. The key idea here to use special parametric models based on deep neural networks to approximate the set of convex functions

---

5. Beyond GANs, 1-Lipschitz networks were used for adversarial robustness and better generalization [186]. More strategies exist for enforcing Lipschitzness property [187] including the penalty adopted by W-GAN followup papers (§B.2), making them actually members of *Grey-box w/ Soft Constraints* category.

in Eq. (22). Note also that Eq. (22) involves only a single convex function ϕ. ICNNs were used for computing OT maps [185] and Wasserstein Barycenter [93]. A number of variant architectures conforming to the same guidelines were also proposed, e.g. ConvICNN [190] and DenseICNN [182], [190] were used in generative networks (§B.2).

A question pertinent to grey-box solvers is *how to differentiate through the OT computation?* The OT plan $\pi^* : \mathbb{X} \times \mathbb{Y}$ that attains the infimum (Eq. (2)) (or any discrete and/or regularised variant) is a function defined implicitly as a solution to an optimisation problem. As such, given sufficient smoothness, functional operators such as differentiators for such an implicitly defined mapping can be examined. Given an iterative algorithm that estimates the optimal plan, one may estimate such derivatives using automatic differentiation by *unrolling* the steps of the algorithm [191] and applying the chain rule to compute derivatives through the algorithm. A popular example of this approach is unrolling the step of the Sinkhorn algorithm [55]. Clearly, this approach becomes more expensive as more iterations are required for the algorithm.

Another approach is use the *implicit function theorem* [192] to compute derivatives without any knowledge of the steps involved in the algorithm. This idea has been explored more generally for solutions to optimisation problems [193] and fixed point equations [194]. When used inside neural networks as layers, these are called implicit layers. Layers that solve optimisation problems such as OT problems are sometimes called Deep Declarative Networks [193].

A unified framework for implicit differentiation of Sinkhorn layers is presented in [195]. They provide expressions for the implicit derivatives of loss functions that accept the result of the pre or post composition of a neural network with a Sinkhorn layer as an input. This allows for gradients to be backpropagated through the neural network pipeline so that variants of stochastic gradient descent can be used to update the neural network parameters. An important consideration is that approximate solutions found using Sinkhorn iterations will propagate errors into the gradient when used with expressions obtained via implicit differentiation. [195] provide a bound on the error of the gradients in terms of the error in the forward Sinkhorn solver in a Frobenius norm sense. Such implicitly differentiated Sinkhorn layers have been used to solve the blind Perspective-n-Point problem [196]. Also, implicit layers were used in few-shot image classification where EMD (1-Wasserstein) was used to compute structural similarities between images [197]. Exploiting the problem structure with implicit differentiation can further increase computational gains significantly as shown by [198]. However, it requires case-by-case handling, lengthy mathematical derivations and custom implementation. Modular implicit differentiation [199] presents an potential technology for handling such difficulty.

But, how good are the neural solvers ? While a neural solver can result in computational savings for a pre-specified class of OT problems on which it was trained, neural solvers can fail to approximate OT solution [200] or exhibit poor generalization beyond the training dataset (see future directions §6). The accuracy of neural solvers is a topic of recent investigation. Recent studies [190] show that neural solvers might not faithfully recover the optimal map yet continue to perform well in downstream tasks. Architectural constraints (e.g. ICNN) also can be a concern. Despite architectural restrictions, [201] demonstrated the richness of this class by showing its capability in capturing the temporal dynamics of complex physical systems accurately. Yet, [202] noted that ICNNs are difficult to train, which led [126] to investigate a better initialization. Beyond the approximation capability, the neural solvers can have characteristic variations. For some solvers, it is possible to obtain the OT plan/map. For example, from the optimal dual potential ϕ and ψ in Eq. (20), one can recover the optimal plan using the primal-dual relationships [203]. For other solvers obtaining the correspondence information (through a plan or a map) isn't feasible.

Relaxed Solver: Sometimes the context in which OT computation is used provides clues for optimizing the computation. If the OT distance is not the ultimate target, then one can consider relaxing OT computations according to subsequent application requirements. Relaxed Word Movers Distance (RWMD) [14] computes a lower bound of WMD by relaxing one of the marginal constraints on the coupling. This reduces the running time complexity from cubic to quadratic in the number of datapoints. Linear-Complexity RWMD [204] further pushes it down to a linear complexity.

5.6 Combined Approaches

While each of the previous strategies can bring benefits on the computational/statistical side, it can be noticed that there is always an associated side effect. Consider, for example, the entropy-based regularization (§5.2). It produces an easy-to-parallelize algorithm and a well-behaved optimization problem (e.g. unique optimal coupling and a differentiable objective). Yet, it can produce a blurry optimal plan ill-suited for some applications [59] or numerical instability issues depending on the magnitude of the regularization λ. Given this, a valid line of thought is combining multiple of the previously discussed (compatible) strategies above in a single framework and aggregating the benefits. Schmitzer *et al.* [148] combines ideas of λ-scaling heuristic (§5.3), an adaptive truncation of the kernel, and the multi-scale coarse-to-fine scheme (§5.4) in a single algorithm to obtain the combined benefits. [144] proposes structuring both the coupling (§5.2) and the kernel matrix (§5.3) to be low rank. Similarly, [205] used the same strategy to obtain a linear time complexity of Gromov Wasserstein.

6 FUTURE DIRECTIONS

Circumventing OT computational challenges almost a decade ago [37] through optimization advancements has led to its emergence in machine learning. Subsequent advances [203] in the same direction enabled embracing OT in many applications. Recently, with more adoption, computational limitations have become more pronounced in the big data era. OT is being used increasingly as a component of various contemporary learning pipelines [206], [207]. Algorithms with high complexity (e.g. super cubic) can not match scalability ambitions and would be rendered as computational bottlenecks. Notably, theoretical advances are no less important in this aspect. They pave the way for principled

scaling rather than 'hacking the way' to a better complexity. Below, we highlight, in broad strokes, some challenges we see as worthy of future research investment.

Computational complexity: The emphasis of this survey is on the longest-standing research challenge in OT, the computational budget. Despite the algorithmic advancements, more research is needed to develop computational approaches better suited for the massive amounts of data available. A recurring pattern observed in computational tricks in §5 is the presence of some sort of compromise. In general, the computational tricks involve a transformation of the original transportation problem into an easier *approximate* version that potentially works better in a specific niche of problems. Unfortunately, this process typically incurs losing some connections and properties of the original formulation. Research in this direction will continue to flourish.

Research at the intersection of optimization and deep learning [208] presents one promising direction in this regard. The recent trend of *learned optimization* [209] and *amortized optimization* [210] builds on the intuition that we can leverage the structure of multiple related optimization problems and learn a model to predict the solutions faster (without running the naive optimization from scratch). Evidently, this was shown to produce reliable results with orders of magnitude faster than classical optimization. Recently, some works have started exploring this direction [211] [212] in the OT literature. Interesting research questions in this direction are generalizing predicted solutions to out-of-(training data) distribution and, more importantly, addressing the lack of theoretical convergence guarantees in this setup. Formulations that can benefit the most from these efforts are the complex (higher order) formulations such as Multi-marginal OT, Gromov Wasserstein, and hierarchical formulations. To date, leveraging these formulations in learning tasks is subject to careful judgment that weighs the potential benefits against the computational price.

Benchmarking. Given the wealth of OT formulations and solvers, standardized evaluation of solvers would greatly benefit future research. Yet, very few works [182], [213], [214] have addressed OT solver benchmarking in the literature. These works are mostly [182], [213] focused on discrete OT solvers with low-dimensional inputs that don't reflect the current nature of big datasets. Benchmarking OT solvers is challenged by the diversity of the OT landscape in terms of formulations and Machine Learning applications considered. One potential direction is to initiate specialized benchmarks that focus on specific tasks in which OT is used extensively (e.g. domain adaptation §B.1) or specific group of formulations. Another benchmarking challenge is designing metrics that assess the OT computation accuracy. This is difficult given that in some cases (such as OT between continuous measures [215]) it isn't straightforward to obtain the ground truth OT solution including the cost and plan. Recent efforts [215] address this issue for special cases. It would be interesting to see more research in this direction. Another issue with the metric design is that the performance on the downstream task isn't an indicator of a good approximation of the Wasserstein distance [200].

New data types. Extending the scalability of OT beyond data quantity to new data types is a promising future direction. Complex data such as structured objects (e.g.

graphs and trees) and sequences and temporal streams have received little attention in the literature. OT on structured objects requires generalizing the notion of transport to accommodate both nodes and relations while ensuring other properties. Current approaches imitate OT on these objects either by having multiple levels of computation (e.g. [106] and nested formulations §4.6) or embedding the data in a new space that preserves the structure (e.g. hyperbolic embedding [216]). Both approaches come with increased computation challenges. A native formulation that embraces the structure is yet to be realized. Potential advantages of native formulations include having the structure reflected in marginal constraints and the capacity to perform interpolation and Barycentric mapping on structured data.

Novel applications. By treating the weights of deep models as probability measures, one can cast the framework of OT on problems involving transportation among (trained) model objects. This can be useful, for example, in model fusion and interpolation without requiring retraining. For example, OT can be used [217] for fusing pre-trained neural networks as an easy 'one-shot' transfer mechanism. Another use of this can be in *federated learning* [218], in which multiple clients collaborate to train a global model in a coordinated server without exchanging their local data. Traditional clients model aggregation (i.e. averaging) on the coordinated server can be replaced with OT-based advanced averaging (e.g. Wasserstein barycenter) that can address clients heterogeneity [219]. Along the same line, OT can be used as a geometry-aware analytics tool for studying the weights of deep models and addressing issues of interpretability.

Scalable OT Tools. Python OT (POT) [125] is a popular long-standing open-source OT package. Alignments in Fig. 4 were computed using POT. Given the growing need for faster and more scalable computation of OT, new tools that can transport millions of samples are needed. In addition to the ongoing acceleration of POT, a few alternative packages are emerging. OTT [220] is a JAX toolbox for OT. OTT leverages accelerators, such as GPU or TPU, for speeding up computations and it includes scalable implementation of popular OT algorithms. The current implementation includes various variants of Sinkhorn, complex formulations such as Gromov Wasserstein, and neural solvers based on ICNN. KeOps package [221] handles the issue of OT computations not fitting in GPU memory when the number of samples is large. It stores computation as formulas that can be streamed on the fly. It is difficult to extend the package and use it and deep learning [175]. Looking forward, overcoming the limitations of current tools will drive wider applicability of optimal transport.

7 CONCLUSION

Optimal Transport provides a coherent mathematical toolbox for formulating and solving problems whose core is about similarity quantification and correspondence estimation between objects. Problems like these, which arise naturally in machine learning, are addressed by transporting collections of samples or features in a geometric sensible way. We discussed various transportation mechanisms (i.e. formulations), their characteristics, their application value,

and their limitations. Then, a formulation-agnostic taxonomy for scaling OT computation to large datasets was presented. We reviewed a broad spectrum of techniques ranging from those native to the optimization framework (e.g. additional constraints or structures) all the way to those inspired by recent advances in machine learning (e.g. neural learned solvers).

ACKNOWLEDGMENTS

APPENDIX A
EXTENDED BACKGROUND (KANTOROVICH)

$$\mathcal{K}_m(\mu,\nu) = \inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathbb{X} \times \mathbb{Y}} m(x,y)d\pi(x,y)$$
$$= \inf_{\pi \in \Pi(\mu,\nu)} \mathbb{E}_{(X,Y) \sim \pi}\left[m(X,Y)\right]. \quad (23)$$

The benefit of reformulating the Kantrovich formulation in Eq. (23) is that multiple computational approaches and interpretations of the same problem can be obtained. For example, we can invoke the basic concept of "duality" [6] from convex optimization to obtain a sup (dual) formulation equivalent to the inf (primal) formulation in Eq. (23). The Kantorovich dual of Eq. (23) [35, Theorem 5.10] is :

$$\mathcal{K}_m(\mu,\nu) = \sup_{(\varphi,\psi) \in \Phi_m} \left[\int_{\mathbb{X}} \varphi(x)d\mu(x) + \int_{\mathbb{Y}} \psi(y)d\nu(y) \right]$$
$$= \sup_{(\varphi,\psi) \in \Phi_m} \mathbb{E}_{X \sim \mu}\left[\varphi(X)\right] + \mathbb{E}_{Y \sim \nu}\left[\psi(Y)\right], \quad (24)$$

where

$$\Phi_m = \{(\varphi,\psi) \in L^1(\mu) \times L^1(\nu) \mid \varphi(x) + \psi(y) \leq m(x,y)\}.$$

The dual is obtained by noting that Eq. (23) is a linear programming problem. For more on linear programs, dual problems, and the Kantorovich problem, refer to the appendix section on optimization (§C).

Given that Eq. (24) bears little resemblance to Eq. (23), one might wonder how intuitively the original minimum-effort transportation can be interpreted as a maximization problem. This can be understood by imagining that the transportation is carried out by a third-party beneficiary. Their goal then would be to maximize the total sending $\int_{\mathbb{X}} \varphi(x)d\mu(x)$ and receiving $\int_{\mathbb{Y}} \psi(y)d\nu(y)$ costs as long their rate stays below a baseline cost $m(x_i,y_j)$ for all the points x_i and y_j. Otherwise, the interested parties would opt out (assuming a competitive market climate).

Note also that in Eq. (24), we no longer have the optimal plan π but rather the variables $\varphi(x)$ and $\psi(y)$, known as the *dual potentials*. They still encode the optimal plan and

---

6. Any convex minimization problem admits a *dual* maximization problem whose optimal value lower-bounds that of the minimization problem. The optimal values coincide if the problem is a linear program.

---

the optimal plan can be recovered from them. However, the dual problem Eq. (24) can be more computationally tractable than the primal problem Eq. (23). In the discrete problem in §2, the $n \times m$ memory footprint of the optimal plan is reduced to a linear (w.r.t the input) footprint required by the dual potential. This can be very useful when scaling the computation to thousands or even millions of points. On the other hand, we still have $n \times m$ constraints to check against when solving Eq. (24).

A simplification that makes the objective dependent only on one potential φ (or ψ) instead of two potentials φ and ψ is called the *semi-dual*:

$$\mathcal{K}_m(\mu,\nu) = \sup_{\varphi \in L^1(\mu), \varphi \in L^1(\nu)} \left[\int_{\mathbb{X}} \varphi(x)d\mu(x) + \int_{\mathbb{Y}} \varphi^c(y)d\nu(y) \right]$$
$$= \sup_{\varphi \in L^1(\mu), \varphi \in L^1(\nu)} \mathbb{E}_{X \sim \mu}\left[\varphi(X)\right] + \mathbb{E}_{Y \sim \nu}\left[\varphi^c(Y)\right], \quad (25)$$

where $\psi(y)$ was substituted with $\varphi^c(y) = \inf_x[m(x,y) - \varphi(x)]$ which is the c-transform of φ. The term *semi-dual* was named as such by [222] since this formulation underlies semi-discrete OT methods. Notably, the flexibility by which this formulation allows us to pick one potential over the other can be useful in applications. In some problems (e.g. crowd counting [65]) that require transporting a sparse measure (e.g. a few annotation points) to a dense one (e.g. many image pixels), one can eliminate the dense potential and optimize over the other.

If additionally the cost function $m(x,y)$ is a distance function $d(x,y)$ satisfying the triangle inequality, e.g. $m(x,y) = d(x,y) = ||x - y||_2$, the c-transform becomes $\varphi^c(y) = -\varphi(y)$ where $\varphi(y)$ is a 1-Lipschitz function w.r.t the distance d:

$$|\varphi(X) - \varphi(Y)| \leq d(X,Y) \quad \forall(X,Y)$$

We denote the set of 1-Lipschitz functions w.r.t the distance d by $L_d(1)$. So the semi-dual formulation [35, Theorem 5.10] becomes:

$$\mathcal{K}_m(\mu,\nu) = \sup_{\varphi \in L_d(1)} \left[\int_{\mathbb{X}} \varphi(x)d\mu(x) - \int_{\mathbb{Y}} \varphi(y)d\nu(y) \right]$$
$$= \sup_{\varphi \in L_d(1)} \mathbb{E}_{X \sim \mu}\left[\varphi(X)\right] - \mathbb{E}_{Y \sim \nu}\left[\varphi(Y)\right], \quad (26)$$

This is known as the *Kantrovich-Rubinstein formulation* or the *Kantrovich-Rubinstein norm* [223]. Note that this only applies to $m(X,Y) = d(X,Y)$, not to the more general (and popular) case $m(X,Y) = d(X,Y)^p$ for some power p.

The exact formulations discussed here are typically solved using classic linear programming solvers. Unfortunately, these methods can't exploit the parallel computing power of modern hardware. As we will see later, approximate formulations that embrace parallel computation and leverage GPUs (e.g. Regularized Optimal Transport §4.1) had gained more popularity in many recent applications.

APPENDIX B
APPLICATIONS OF OT IN MACHINE LEARNING

In this section, we review the applications of OT. We focus primarily on domain adaptation and generative models.

B.1 Domains Alignment (DA)

Domain adaptation is a very vast and vibrant sub-field of machine learning. Its motivation is the frequent violation of a classical learning assumption in the real-world application, namely, the assumption that the data used for model training and testing data come from the same probability distributions. A common example of this assumption failure is an image classification task in the acquisition conditions causing a non-negligible (distributional) shift between training and test datasets [224]. In this vast topic, one can find these excellent recent surveys focusing on specific themes of domain alignment whether it is generalization [225], open set adaptation [226] or theoretical advances [227]. Here, we rather collect the many domain alignment paradigms to which optimal transport contributes. It should be noted that the subtle, but crucial, differences between these domain alignment paradigms can be confusing. Figure 8 provides a quick visual connection between various domain alignments[7] paradigms and highlights the differences.

Formally, domain adaptation considers two domains *source domain* and *target domain* denoted as s and t; respectively. The goal is to learn a robust model given that the source and target domains have different data distributions (i.e $\mathcal{D}_s \neq \mathcal{D}_t$). In the most common setup, the labels for the two domains are assumed to be available in abundance. In this case, it is called *supervised alignment*. Yet, the problem also comes in many other variations. For example, the target data can be unlabeled (*unsupervised alignment*) or partly labeled (*semi-supervised alignment*). The classes in the source and target domain can be disjoint (*open set alignment*). The learning goal, in this case, may go beyond merely rejecting new classes as 'unknown' to include learning the new classes from one (*one-shot alignment*) or multiple (*few-shot alignment*) samples. When the target domain is very dynamic, domain alignment can be done in a continuous manner (*continuous domain alignment*). All the previous configurations assume that the training samples from the source domain are available during adaptation. This is unrealistic in cases where it's impractical to take training data to deployment due to privacy or storage considerations. This motivates *source-free domain alignment*, which makes different assumptions as shown in Figure 8.

Since the first use of OT in domain alignment problems by [57], OT has been considered a strong baseline for domain adaptation. The mathematical flexibility of optimal transport formulations allows for catering to the various domain alignment configurations. Before proceeding with the review of these works, it is instructive to see why one might consider OT in domain alignment problems. The following are potential motivations:

- Distributions matching is at the core of domain adaptation techniques. Given that, the problem naturally lends itself to optimal transport. Moreover, the previously discussed OT properties (e.g. geometry awareness and diversity of formulations) can offer a new perspective on the problem.
- As a metric on the space of probability measures, the Wasserstein distance was used to derive theoretical bounds [8] on the domain adaptation error [228] and deep learning generalization [7]. Thus, many techniques building on OT for alignment [229] provide alignment theoretical guarantees.
- It is possible to encode a prior structure useful for domain alignments (e.g. preserving labels or graphical relations between samples [57]) in the OT mapping between the source and target distributions.
- A degree of interpretability can be gained by having explicit correspondences, represented in a transport plan, between domains/datasets. This might be even more true in works that consider the matching in the raw input space [109], [230] rather than (a less interpretable) feature space.

The factors discussed above are possible reasons for OT to be a versatile building block in domain alignment areas. Below is a selective review of these areas.

Transferability Assessment using OT. Transfer learning across domains or tasks is one the most promising and widely adopted machine learning methods and it remains to be a topic of extensive investigation [231]. Yet, judging the effectiveness of a specific transfer learning setting prior to the actual transfer in a quantifiable way has received much less attention. In multi-task learning, one may want to save the joint training burden of weakly related tasks as the final performance will be actually worse compared to a single-task training. Also, this can be used in source model selection [232] (i.e. select from a pool of source pre-trained models the best one for a target task). Rather than relying on human experience for this kind of relatedness assessment, estimating a transferability metric [233], [234] could be much more ubiquitous and practical. General metrics for transferability can be categorized into the *analytical transferability metrics* [234], [235], [236] and *empirical transferability* [237]. The first has stricter assumptions but is computationally efficient and enjoys theoretical generalization bounds. The latter is scalable but lacks theoretical guarantees.

Optimal Transport-based Conditional Entropy OTCE [233] was the first analytical transferability metric in a simultaneous cross-domain cross-task. It adopts the popular "retrain head" transfer model in which a source model with frozen feature extractor is transferred to target data by fine-tuning only the classifier head following the extractor. The proposed metric predicts the model's performance after the tuning without actual re-training.

Unsupervised DA (UDA) using OT. Several works applied OT in UDA in which labels are available only for the source while samples are available from both the source and target domain. Courty *et al.* [57] established a framework, that was extended by others later, for using optimal transport as a principled method for UDA. The key idea is to unify the training and test samples in one domain in which the classifier can be learned. First, OT is used to align the empirical measures of source and target domains. A process that can be done without using any labels. Then, using the estimated transport plan, the training samples are transported to the target distribution. Finally, a classifier can

---

7. we use 'domain adaptation' and 'domain alignment' interchangeably in this section.

8. For the interested reader, Redko *et al.* [227] nice exposition of theoretical domain adaptation covers learning bounds based on various statistical frameworks and divergences (Wasserstein included).
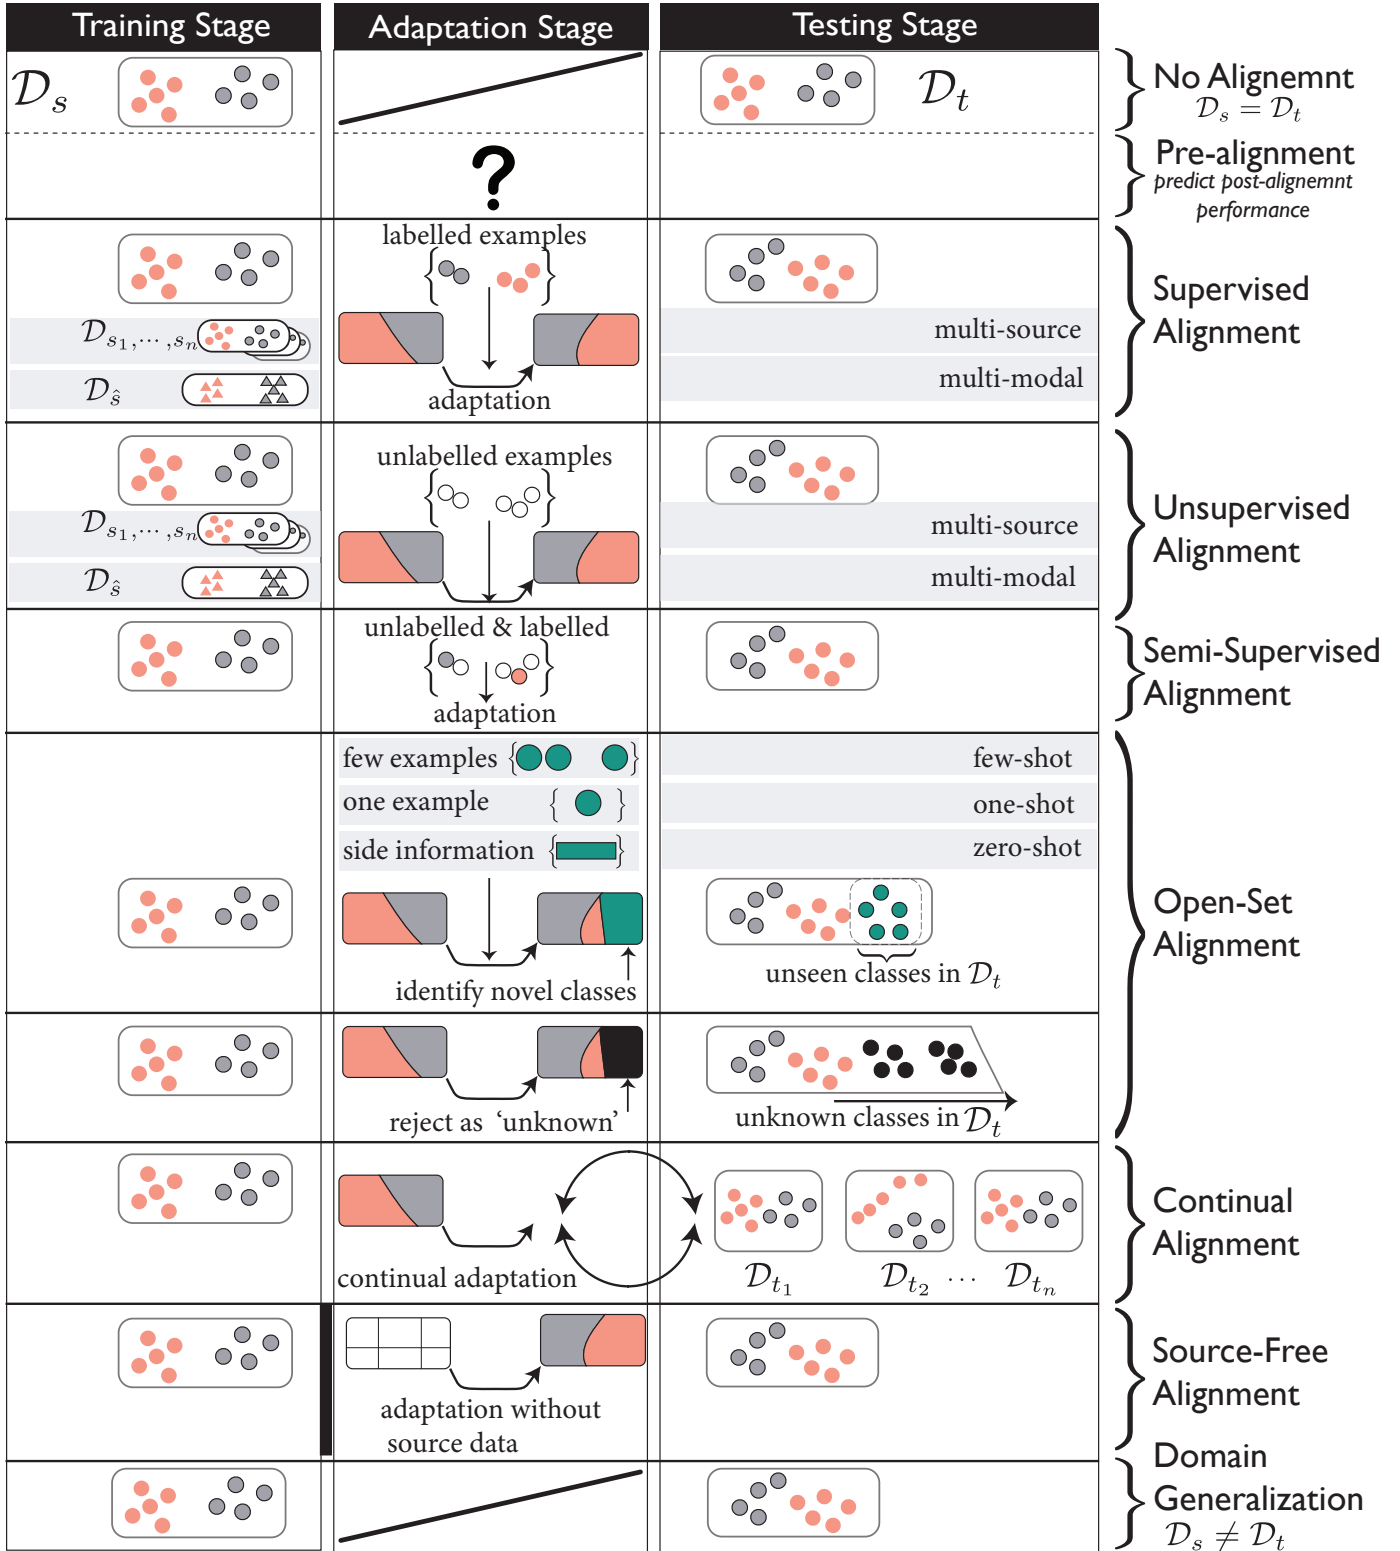
Fig. 8. Various domain alignment configurations addressed by OT works in §B.1

be learned only on the target domain. A graph-inspired regularization is used to preserve the after-transportation proximity of the source samples sharing the same label.

Instead of sample alignment, follow-up works seek to align the model's internal features corresponding to $\mathcal{D}_s$ and $\mathcal{D}_t$. Sliced Wasserstein Discrepancy (SWD) [85] builds on

Maximum Classifier Discrepancy (MCD) [238] that maximizes the discrepancy between task-specific classifiers as a part of an adversarial alignment strategy. As the discrepancy is a key component of MCD, SWD proposes upgrading it with the Sliced Wasserstein distance. Interestingly, by merely replacing the L_1 discrepancy in MCD with SWD, the
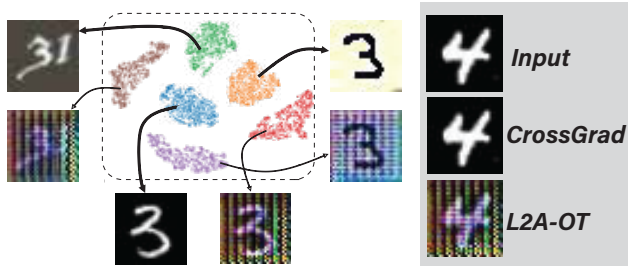
Fig. 9. (Left) L2A-OT augments source domains (thick arrows) with synthesized novel domains (thin arrows) by learning to maximize the Wasserstein distance between the two sets. (Right) synthesized novel samples are qualitatively better compared to adversarial perturbation approach CrossGrad [243] (figures adapted from [244])

system consistently delivers better performance in image classifications, segmentation, and object detection tasks.

Enhanced Transport Distance (ETD) [233] uses a neural solver (§**5.5**) for OT-based features alignment and notes the mini-batch instability issue. Specifically, optimizing OT on mini-batch can lead to inconsistent transport plans across iterations. To mitigate the bias caused by mini-batch, the system weighs (calibrates) the ground cost using a network-integrated attention.

Wasserstein Guided Representation Learning (WDGRL) [239] proposes augmenting Adversarial Domain Adaptation (ADA) architectures, such as DANN [240], with additional Wasserstein loss to guide the domain alignment process. ADA extends the concept of Generative Adversarial Networks (GAN) to domain adaptation. While GANs leverage the adversary to generate samples indistinguishable from the real data, ADA seeks a mapping that makes source features indistinguishable from target domain features. WD-GRL adds a WGAN-like domain critic (WGAN discussed in §B.2) for calculating the 1-Wasserstein distance between the source and target features. The motivation behind this is the possibility of having vanishing gradients when the supports of the mapped features are disjoint, a setting in which OT is assumed to work better compared to other divergences. In fact, WDGRL significantly outperforms DAN in this setting, albeit the results for this are based on synthetic data.

Unsupervised DA using OT (multi-source). JCPOT [241] considers the mutli-source DA problem with a focus on cases where the label proportions between source and target domains are different. Target shift (or class imbalance) has practical values in applications like anomaly detection. Theoretically [242] the error on the target domain can be minimized by re-weighting the distribution of the source classes to match the target one. Motivated by the impracticality of acquiring the target class distribution, an optimal transport matching that automatically acquires the weights is proposed. Entropy regularised OT (§**4.1**) is used for mapping between every single source and the target and the problem is formulated as a Bregman projection that has a simple solution (proposition 1 [241]). The approach shows improvement over vanilla OT on the satellite imagery dataset.

Open Set DA using OT. In open set domain alignment, the target domain can contain classes that weren't present in the source domain. While more practical, it is a challenge where domain alignment methods fail due to interference with extra unknown classes. The key goal in open set domain alignment [245] (OSDA) is to correctly classify known classes while rejecting the extra classes as "unknown". Xu *et al.* [246] noticed that aligning the whole source and target domains followed by previous OT methods is actually inferior to the adaptation in an open set setting. Causing a problem known as *negative transfer* [247]. Thus, Joint Partial Optimal Transport (JPOT) [246] employs partial optimal transport to align only "well-matched"/most-correlated samples and avoid far-fetched pairings that can cause the negative transfer. To obtain the partial matching, the mean cost of the complete optimal transport matrix [248] is used as a threshold. Coupled pairs whose distance is greater than this threshold are marked as non-transferable.

Open Set DA using OT (zero-shot). For the more challenging cases of totally unseen classes (i.e. zero shot), GZSL [249] proposes synthesizing features for the unseen classes from the corresponding auxiliary textual descriptions. The feature generator is encouraged to match real feature distributions by minimizing regularized OT distance. OT-based matching was motivated by the concern of sampling outliers (brought by domain shift) that can make point-wise matching approaches less robust. IPOT [147] is used instead of Sinkhorn as it maintains near-linear time complexity while being less sensitive to regularization weight.

Multi-modal and heterogeneous DA using OT. Applications that require models to operate across different modalities motivate the need for heterogenous domain alignment. Clearly, aligning features living in heterogeneous (multimodal) spaces is more complicated than the alignment of homogeneous (unimodal) features. In a neural population decoding application, HiWA [108] used data from populations of (human brain) neurons to predict the arm movement direction in reach-out tasks. This involves aligning two domains of different modalities; the neural data distribution and the 3D movement data distribution. Given the difficulty of the alignment, they suggest leveraging the additional clustering structure of the data to regularize OT and constrain the solution space. Thus, a nested formulation of OT is proposed for aligning clustered and multi sub-spaces data and the framework hierarchical optimal transport (§**4.6**) was invoked.

Knowledge Distillation (KD) is a common option for cross-modal domain alignment. That is; aligning latent features of two models each of them trained on a different modality yet solving the same or related tasks. Wasserstein Contrastive Knowledge Distillation (WCKD) [9] encourages a student (compact) model features h^s to be distributionally similar to those of an advanced teacher h^t. Specifically, a GAN-like discriminator ensures that source samples are not distinguishable from target samples. The alignment is done using a neural version of OT (WGAN [18]).

Continual DA using OT. Very few works developed techniques that address gradual distributional shifts in the target domain. Continual Domain Alignment [229] is a problem that resembles Continual Learning (CL) [250] albeit without assuming the availability of labels in the target domain. LDAuCID [229] learns a source model capable of adapting to several sequential target domains without
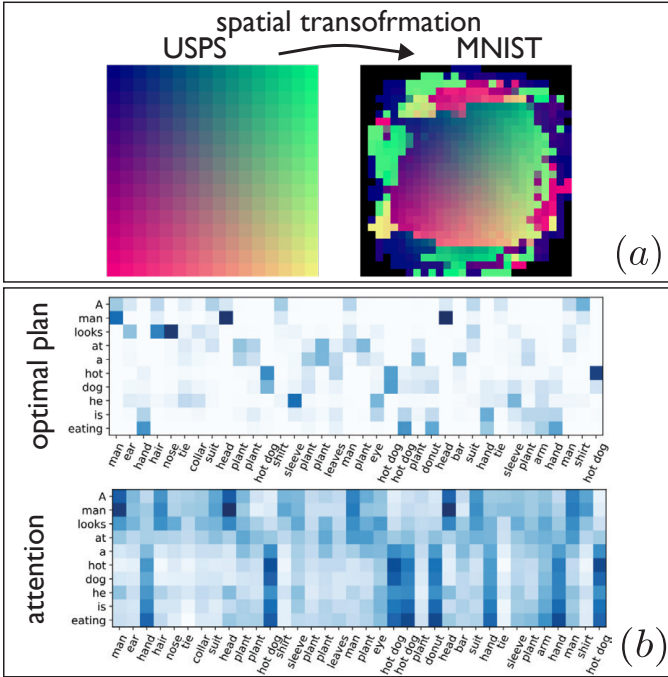
Fig. 10. **Interpretability of optimal transport plan**. Optimal plan is used to reveal in (a) features (represented by raw pixels) correspondences between USPS and MNIST datasets using an unsupervised approach [230] and (b) images object-to-word correspondences [251]. The latter was shown to be favorable from an interpretability perspective compared to Transformer's attention due to sparsity. The figures were adapted from original references

labels. A strong overlap between the source and target classes is assumed. LDAuCID aligns the *consolidated* internal learned distribution with the target domain it using Sliced Wasserstein. To mitigate the catastrophic forgetting issue, representative samples from all tasks are stored and replayed in the adaptation stage. This technique is known as *experience reply* in CL literature.

Source-free DA using OT. Traditional UDA performs the alignment between domains using the source and target data. Requiring the presence of source data during alignment is impractical in cases when the data is inaccessible due to privacy (e.g. Federated Learning clients [252]), storage limitations (e.g. on edge device), or other reasons. Augmented Self-Labelling (ASL) [253] uses the pre-trained source model (instead of the source data) and the unlabeled target data for alignment. ASL generates pseudo-labels for target data using the pre-trained model to be used during the model adaptation. It builds on a technique [49] that frames *self-labeling* as an optimal assignment between the model's predictions (soft-labels) and the set of possible labels under an equipartition constraint (that avoids degenerate solution where a single label is assigned to all samples).

Domain generalization using OT. UDA assumes that data can be collected in advance (prior to training) from the target domain since this is required for training. This assumption doesn't hold in practice. Domain Generalization (DG) thus is envisioned as a more practical alternative that learns a model capable of operating in unseen target domains without requiring model update (adapta-

tion). Learning to Augment by Optimal Transport (L2A-OT) [244] proposes learning to synthesize novel domains from multiple source domains. By augmenting source domains with synthesized ones, DG can be realized through increased training data diversity. The problem is framed as image-to-image translation and despite being similar in spirit to gradient-based perturbation [243](i.e. perturbing input using adversarial gradients from a domain classifier). Interestingly, the conditional generator in their pipeline is trained to maximize (not minimize) the optimal transport distance between the source domain and the synthesized novel domains. Such divergence loss is balanced with cycle consistency loss [254] and task loss to ensure preserving semantic content. L2A-OT advocates the use of Wasserstein distance compared to f-divergences for divergence maximization as the near-zero denominators for f-divergences tend to generate large but numerically unstable divergence values.

DA interpretability using OT. One of the virtues of OT-based domain alignment is the possibility of inspecting the optimal plan (when available) to gain insights into the domains' correspondence. This offers an opportunity to go beyond merely detecting and handling the distributional shift between the source and target domains. Specifically, an explanation of the domain shift between the source and target datasets through the samples or features correspondences can be attained.

We note that many OT works [59], [150], [230], [251], [255], [256] refer to *interpretability* of the optimal plan without formal specification of what are the interpretability requirements. Yet, there is consensus among them that some characteristics are desirable and favorable from an interpretability point of view. For example, sparse optimal plans [59], [255] are favorable. Sparsity is motivated by the principle of parsimony in which simple solutions should be preferred. For example, a sparse optimal plan in a weakly-supervised alignment setup [251] between images-sentences pairs can be easier to interpret compared to the dense attention from Transformers. As shown in Fig. 10(b) the two constructs (the optimal plan and the attention) visualize the strength of association between words (e.g. 'man') on the vertical axis and visual objects, denoted by words on the horizontal axis. Clearly, having the word strongly associated with a few visual objects (e.g. the 'man' in the optimal plan) is easier to interpret. Note, however, that solving OT can yield an arbitrarily complex plan that's not necessarily interpretable as per the characteristics mentioned above. Thus, a number of works explicitly place additional constraints on the optimal plan reflecting the interpretability requirements. For example, [255], [256] suggest enforcing user-defined interpretable mappings (e.g. k-sparsity) on the optimal plan as additional constraints.

Also, it is easier to interpret the correspondences when they can be traced back to the raw samples [109], [230] . For example, CO-Optimal Transport [230] simultaneously estimates the correspondence between samples (images) and features (independent pixels values across the dataset images) from two handwritten digit datasets. The features transport plan (Fig. 10(a) shows the spatial transformation of USPS into MNIST dataset. Although it is oblivious to the geometric structure of the images, it reveals how pixels

would spatially rearrange in transportation.

DA theoretical guarantees using OT. Another virtue of OT is the possibility of leveraging theoretical error bounds that build on the Wasserstein distance. For example, Courtey *et al.* [257] proved that minimizing the joint distribution optimal transport (JDOT) quantity is equivalent to minimizing the learning bound on the domain adaptation problem. Wasserstein Guided Representation Learning (WDGRL) [239] builds on [228] and proves that the target error is bounded by the Wasserstein distance for empirical measures under the assumption that hypothesis class is K-Lipschtiz continuous for some K. LDAuCID [229] analysis within the PAC-learning framework [258]

B.2 Generative Modeling

Learning models that generate images, audio, video, text or other data is a major sub-filed of machine learning. Generative models have a vast range of possible applications and enormous potential can result from training on endlessly available unlabeled data. Let $X \sim P_r$, where P_r be a fixed distribution from which we have access to samples. The fundamental problem of generative modelling is to make distributions P_z and P_r as close as possible, where P_z is a distribution we may sample from. In high-dimensional and large data regimes, this is often done by minimizing a divergence between the two distributions.

Generative Adversarial Networks (GANs). The GAN consists of two networks — a generator network $g_\theta \in \mathcal{G}$ and a $(0,1)$-valued discriminator network $d \in \mathcal{D}$ — that are jointly trained according to a procedure that may be understood as a minimax game [259]. Given some easy-to-sample distribution P_0, the generator implicitly defines a sample from a distribution $g_\theta(Z_0)$, where $Z_0 \sim P_0$. The discriminator network d is then used to evaluate whether the sample $g_\theta(Z_0)$ comes from P_z or P_r. Whether or not the sample comes from P_z or P_r is measured through a quantity

$$\mathcal{L}_{\text{GAN}}(g_\theta) = \sup_{d \in \mathcal{D}} \mathbb{E}_X[\log(d(X))] + \mathbb{E}_{Z_0}[\log(1 - d(g_\theta(Z)))].$$

$$(27)$$

The GAN's minimiax objective between the generator g_θ and the discriminator d is then

$$\inf_\theta \mathcal{L}_{\text{GAN}}(g_\theta).$$

$$(28)$$

The quantity $\mathcal{L}_{\text{GAN}}$ provides a lower bound on the Jensen-Shannon divergence (JSD), up to a multiplicative and additive constant [259]. Furthermore, if $\mathcal{D}$ were replaced by the set of all $(0,1)$-valued functions, the lower bound would be an equality, and Eq. (28) amounts to minimising the JSD [260].

WGAN and followups. The notorious training instability of the vanilla GAN has been argued to be an issue associated with the JSD metric [18]. Specifically, the JSD remains constant when the two distributions are disjoint (not overlapping). This can happen when P_g and P_r are low dimensional manifold in high dimensional space [260]. Additionally, even if they overlap, the limited sampling (during training) might render them disjoint. The WGAN [18] minimisation objective with respect to a generative network g_θ is

$$\mathcal{L}_{\text{WGAN}}(g_\theta) = \sup_{f_w \in \mathcal{F}} \mathbb{E}_{x \sim P_r}[f_w(x)] - \mathbb{E}_{z \sim P_g}[f_w(g_\theta(z))],$$

$$(29)$$

$$\inf_{g_\theta} \quad \sup_{f_w} \quad \mathbb{E}_{x \sim P_r}[f_w(x)] - \mathbb{E}_{z \sim P_g}[f_w(g_\theta(z))] + \lambda R$$

WGAN	$f_w \in \mathcal{L}ip(f)$	$R = 0$
WGAN-GP		$R = \mathbb{E}_{\hat{x} \sim \hat{X}}[\|\nabla f_w(\hat{x})\|_2 - 1]$ $\hat{x} = tx + (1-t)g_\theta(z), \quad t \sim U[0,1]$
WGAN-LP		$R = \mathbb{E}_{\hat{x} \sim \hat{X}}[(\max\{0, \|\nabla f_w(\hat{x})\| - 1\})^2]$
SWGAN	$f_w \in \mathcal{S}ob(f)$	$R = 0$

Fig. 11. Connection between WGAN [18] and the followup works WGAN-GP [262], WGAN-LP [261] and SWGAN [263]. $\mathcal{L}ip$ and $\mathcal{S}ob$ denote 1-Lipschitz constraint and Sobolev integral constraint; respectively.

where $\mathcal{F}$ is a space of 1-Lispchitz neural networks. f_w here replaces GAN's $(0,1)$-valued discriminator and is called the critic. Recall that the equivalent form of Kantrovich-Rubinstein duality [35] is

$$W_1(\mu, \nu) = \sup_{f \in \overline{\mathcal{F}}} \mathbb{E}_{X \sim \mu}[f(X)] - \mathbb{E}_{Z \sim \nu}[f(Z)],$$

$$(30)$$

where $\overline{\mathcal{F}}$ is the space of all 1-Lispchitz functions. Since $\mathcal{F} \subset \overline{\mathcal{F}}$, it follows that the GAN objective Eq. (29) is a lower bound of the 1-Wasserstein distance Eq. (30). The 1-Lipschitz constraint ($\mathcal{L}ip(f)$) in Eq. (29) also ensures a smooth critic and, thus, a more meaningful gradient for updating the generator. In practice, WGAN enforced this constraint using weight clipping heuristic to enforcing the parameters to be in a compact space ($[-0.01, 0.01]$). WGAN-LP [261] augments the loss with a regularization term that penalizes the deviation of the critic's gradient norm from one. Followup works argued that explicit enforcement of $\mathcal{L}ip(f)$ might be unnecessary. WGAN-GP [262] observed that the optimal critic has gradient norm equal to 1 almost everywhere. They replace the weight clipping with a gradient penalty term that encourages this characteristic. Sobolev WGAN [263] argues that the strong Lipschitz constraint might be unnecessary for optimization. Another direction for tackling the instability that can be caused imperfect optimization of the critic is to use a fully tractable divergence. Genavy et al. [264] used regularised OT evaluated on mini-batches.

Diffusion and score based methods. Score based [265] and diffusion generative models [266] both rely on a similar mechanism. Input data (such as images, or text) is perturbed by incrementally adding small amounts of noise by a procedure called the forward process. The forward process is a so-called diffusion process. The input data is recovered by running another diffusion process called the reverse process. The reverse process involves the score function — the derivative of the log probability density function of the perturbed data distribution — which is unknown. Instead of using the exact score function, score based methods use a neural network approximation. The parameters of the neural network are learned using the so-called (conditional) score-matching loss, which is coarsely an expected weighted mean-squared error between the neural network and the true score function [267]. Samples from the approximate

data distribution can be generated by running the reverse process.

The forward and reverse mechanisms do not explicitly minimise a divergence between the input data and the output of the reverse process. Surprisingly however, score based methods are equivalent to minimising an upper bound on the KL divergence between the reverse process and the input data distribution [267]. More recently, it has been shown that the 2-Wasserstein distance is upper bounded by the (conditional) score-matching loss, up to multiplicative and additive constants [268].

APPENDIX C
OPTIMIZATION BACKGROUND

C.1 Minimum cost flow problems

An exact optimal transport problem on discrete measures can be formulated as the classic minimum cost flow (MCF) network problem which is a graph-theoretic problem that can be formulated using the following linear program:

$$\underset{\boldsymbol{p}}{\text{minimize}} \quad \mathcal{L}_c(\boldsymbol{\mu}, \boldsymbol{\nu}) = \sum_{i=1}^{m} \sum_{j=1}^{n} c(x_i, y_j) p_{ij}$$

subject to

$$\sum_{j=1}^{n} p_{ij} = \mu_i \quad \forall i \in [1, m], \tag{31}$$

$$\sum_{i=1}^{m} p_{ij} = \nu_j \quad \forall j \in [1, n],$$

$$p_{ij} \geq 0 \quad \forall i \in [1, m], j \in [1, n]$$

where c is a cost function, usually a distance measure between x_i and y_j. General linear programs can be solved using a number of algorithms with the Simplex method [269], [270], [271] and the interior point method [272] being 2 of the most popular choices. The computational complexity of the Simplex method is exponential in the number of decision variables in the worst case while the interior point method has polynomial time complexity in the number of decision variables. However, in practice the Simplex method tends to be significantly faster for most practical problems. Besides computational complexities and running time, the main difference between the Simplex method and the interior point method from a user's perspective is that the Simplex method always finds a so-called "extreme" or corner point of the feasible domain as the optimal solution, while the interior point method finds a point in the relative interior of the feasible domain as the optimal solution. An extreme point is a point at the intersection of N hyperplanes where N is the number of decision variables ($m \times n$ above). Since the optimal transport problem has $m + n$ equality constraints, assuming $m \times n > m + n$, the remaining $m \times n - m - n$ constraints to be satisfied at equality will be bound constraints of the form $p_{i,j} \geq 0$. The Simplex method will therefore naturally tend to find sparse optimal plans. The relative interior of a feasible domain is the set of points that satisfy all the linear equality constraints (e.g. $\sum_{j=1}^{n} p_{ij} = \mu_i$) at equality, but with all the inequality constraints ($p_{ij} \geq 0$) satisfied at a strict inequality.

That said, using the standard Simplex method or interior point method for generic linear programs is not the most efficient way to solve minimum cost flow (MCF) problems. MCF problems have a wide variety of more efficient algorithms with polynomial time complexities [273], one of the most famous of which is the so-called network Simplex method [274]. The network Simplex method is a variation of the general Simplex method specialized for MCF problems to have a cheaper per iteration cost and a polynomial bound on the number of iterations needed to converge. The general Simplex algorithm is both more expensive per iteration and has an exponential worst case complexity for the number of iterations needed to converge. More recently, faster algorithms for the MCF problem have also been proposed [275].

C.2 Convex optimization

C.2.1 Set convexity

A set of points $\mathcal{S} = \{\boldsymbol{x}\}$ is said to be convex if every weighted average of some points inside the set is also inside the set. Equivalently for a set $\mathcal{S}$ to be convex, the following must be satisfied for all $(\boldsymbol{x}_2, \boldsymbol{x}_2) \in \mathcal{S}$ and for all $0 < \alpha < 1$.

$$\alpha \times \boldsymbol{x}_1 + (1 - \alpha) \times \boldsymbol{x}_2 \in \mathcal{S} \tag{32}$$

The above weighted average is also known as the convex combination of $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$.

The intersection of any number of convex sets is either the empty set or another convex set. Therefore in an optimization problem with multiple constraints if the set of feasible points to each constraint individually is a convex set, the set of feasible points to all the constraints is either the empty set (problem is infeasible) or another convex set.

C.2.2 Function convexity

A function $f(\boldsymbol{x})$ with domain $\mathcal{X}$ is called convex if the set of all points in $\mathcal{X}$ satisfying the constraint $f(\boldsymbol{x}) \leq c$ is either an empty or a convex set for all $c \in \mathbb{R}$. Note that the set of such points doesn't have to be a compact set, that is it can extend to $\pm\infty$ along one or more directions. Equivalently, one can say that a function is convex if the set of feasible points to its epigraph is a convex set. The eipgraph of the function $f(\boldsymbol{x})$ is the set of points $\{(\boldsymbol{x}, c) \in \mathcal{X} \times \mathbb{R}\}$ such that $c - f(\boldsymbol{x}) > 0$. Figure 12 shows the epigraph of $f(x) = x^2$.
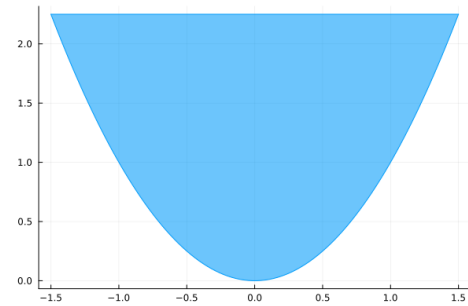

Fig. 12. Epigraph of $f(x) = x^2$.

Showing that a function is convex can be done by proving any of the above properties. There are however other ways to check if a function is convex. Another common check for twice differentiable functions is to check if the

function's Hessian is positive semidefinite everywhere in $\mathcal{X}$. Identifying and building convex functions in a disciplined way is a well-studied field, often termed "disciplined convex programming" (DCP). For more on DCP, the readers are advised to check the excellent online resources available in https://dcp.stanford.edu.

C.2.3 Constraint convexity

In optimization, a constraint is said to be a convex constraint if the set of points satisfying this constraint is a convex set. One common way to construct a convex constraint is using a convex function $f(\boldsymbol{x})$. If f is a convex function, the following constraint is convex:

$$f(\boldsymbol{x}) \leq c \tag{33}$$

for any constant c, assuming there exists at least 1 point satisfying this constraint. Other common convex constraints are affine (aka linear) constraints of the form:

$$f(\boldsymbol{x}) = 0 \tag{34}$$

where f is an affine function of $\boldsymbol{x}$, e.g. $f(\boldsymbol{x}) = \boldsymbol{b}^T \boldsymbol{x} + a$ for some constants $\boldsymbol{b}$ and a. Note that it's customary to use the term linear to refer to affine functions and/or constraints in optimization literature.

A constrained optimization problem with a convex *minimization* objective function and convex constraints is called a "convex program". Formulating decision problems as convex programs and solving them with convex optimization algorithms is often called "convex programming". Convex programs are theoretically appealing because under mild assumptions, they have efficient algorithms that can find a so-called *global optimal solution*. A global optimal solution is a (not necessarily unique) feasible point $\boldsymbol{x}$ that absolutely minimizes the objective function, i.e. no other solution can do strictly better.

Note that it has to be a minimization problem. Maximizing a general convex function is NP-hard. If your objective $f(\boldsymbol{x})$ is a quantity you want to maximize, that's equivalent to minimizing $-f(\boldsymbol{x})$ instead. If $-f(\boldsymbol{x})$ is a convex function, f is known as a concave function. Therefore for a maximization optimization problem to be convex, the objective function must be concave and the constraints must be convex.

C.2.4 Structured constraints

Among convex constraints, there are some special classes of constraints that are generally considered nicer than "generic" convex constraints. These constraints have a specific structure which can be exploited in more efficient algorithms. These are broadly known as "*structured constraints*". A simple class of structured constraints are the affine/linear constraints, e.g:

$$\boldsymbol{b}^T \boldsymbol{x} + a \quad (\leq / =) \quad 0 \tag{35}$$

There are some algorithms that can handle linear inequality constraints more efficiently than generic convex constraints. Additionally, linear equality constraints are the only equality constraints that are convex. Any other equality constraint where the function in the constraint is not affine/linear can

never be convex. Linear constraints are therefore among those special classes of so-called *structured constraints*.

Another broad class of structured constraints is conic constraints. A conic constraint is a convex constraint whose feasible set is a special convex set known as a cone. A convex set S is considered a cone if for every $\boldsymbol{x}$ inside the set S, $c\boldsymbol{x}$ was also inside the set S for all $c \geq 0$. Note how a cone must be a non-compact set since c is allowed to go to ∞.

A structured conic constraint can be written as:

$$\boldsymbol{x} \in \mathcal{C} \tag{36}$$

where $\mathcal{C}$ is a cone. There are various algorithms that can handle conic constraints more efficiently than regular convex constraints. It is therefore often appealing to convert regular mathematical constraints that are known to be convex to their conic equivalents to make use of these efficient algorithms.

C.2.5 Algorithms

There are various classes of optimization algorithms for solving each class of optimization problems. For linear programs, the most popular algorithms are:

- Primal/dual Simplex algorithm
- Interior point algorithm
- Primal-dual hybrid gradient algorithm

For unstructured convex and general nonlinear programs, the following are popular algorithms:

- Interior point algorithm
- Sequential linear programming (with trust region)
- Sequential quadratic programming (with trust region)
- Method of moving asymptotes algorithm
- Augmented Lagrangian algorithm

For convex and general nonlinear programs with an unstructured objective but structured convex (conic) and linear constraints, the following are popular algorithms:

- Frank-Wolfe algorithm
- Alternating direction method of multipliers (ADMM) algorithm
- Projected gradient descent algorithm
- Mirror descent algorithm

More generally, "proximal algorithms" are a family of algorithms for handling this class of problems which encompasses multiple of the above algorithms, such as ADMM, projected gradient descent, mirror descent and many more algorithms.

For structured convex (conic) programs, the following are popular algorithms:

- Interior point algorithm
- ADMM with conic operator splitting

C.3 Differentiability

A function is any one-to-one or many-to-one (but not one-to-many) mapping from a set called the domain to another set called the range of the function. The uniqueness of the value in the range for each value in the domain is a characteristic property of what a function is. For some

parametric optimization problems with parameters θ, there exists a unique optimal solution $x^*(\theta)$. In such cases, x^* can be considered a proper function of θ. Functions can be continuous, dis-continuous or semi-continuous. Continuity of a function means that for every infinitesimal change in the inputs, an infinitesimal change takes place in the output. In the context of optimization problems, this means that for every small change in θ, x^* only changes by a small amount. Such continuity assumption is generally violated if the solution x^* is discrete since discrete solutions can only move in non-infinitesimal steps, except in the trivial case where the solution x^* does not change at all with θ. Semi-continuous functions are continuous almost everywhere but have a zero-measure subset of the domain where the function has a dis-continuity. Optimal solutions to non-convex optimization problems with multiple modes can exhibit such semi-continuity.

C.4   Duality

C.4.1   What is duality?

Duality is one of the most important concepts in optimization. Sometimes, instead of solving an optimization problem in its "natural formulation", also known as the primal formulation, one can solve a related optimization problem, known as the dual problem, or one can solve the primal and dual problems together. Let the primal optimization problem be:

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & f(x) \\ \text{subject to} \quad & \\ & x \in P \end{aligned} \quad (37)$$

where x is the vector of primal decision variables, f is the primal objective function and P is the feasible domain of the primal problem, i.e. the set of all points that satisfy any constraints on x. Additionally, let x^* be the unknown optimal solution of the primal optimization problem. The dual problem would be another optimization problem of some so-called dual variables y:

$$\begin{aligned} \underset{y}{\text{maximize}} \quad & g(y) \\ \text{subject to} \quad & \\ & y \in D \end{aligned} \quad (38)$$

where g is the dual problem's objective and D is the feasible domain of the dual problem such that for any feasible $y \in D$, $g(y) \le f(x^*)$. The fact that every dual feasible solution y gives a lower bound $g(y)$ on the optimal primal objective value $f(x^*)$ is what makes the primal and dual problems entangled with each other. The optimal dual solution $y^* \in D$ that maximizes g gives the best/tightest lower bound on the optimal primal objective. In problems where the so-called strong duality exists, the best bound $g(y^*)$ is exactly equal to the best primal objective value $f(x^*)$. For linear programs and convex programs (subject to a mild constraint qualification condition), strong duality holds.

C.4.2   Importance of the dual problem

There are various reasons to formulate and solve the dual problem. One reason is to identify if a primal solution is near optimal. Assuming the primal problem is a minimization problem, a dual solution that's feasible to the dual problem provides a lower bound on the optimal objective value of the primal problem. Having a good (tight) lower bound on the optimal value of our primal problem helps us to know how far we are from optimality in the worst case. For instance, let x be the best feasible primal solution found and let y be a feasible dual solution. In this case, we know that $g(y) \le f(x^*) \le f(x)$. A feasible dual solution with a higher $g(y)$ or a feasible primal solution with a lower $f(x)$ provide tighter bounds around the optimal value $f(x^*)$. Therefore knowing y, we know that $f(x)$ can never be more than $f(x) - g(y)$ away from $f(x^*)$. When strong duality applies, the best bound $g(y^*)$ is exactly equal to the best primal objective value $f(x^*)$. In this case, the dual solution can be used to provide a proof/certificate of convergence to the optimal solution if we evaluate $f(x)$ and $g(y)$ for the best x and y found and they turn out to be equal. However, strong duality doesn't always exist for all classes of optimization problems.

Another use of the dual problem is that sometimes it's beneficial to solve the primal and dual problems simultaneously like in the primal-dual hybrid gradient algorithm for large scale linear programs or the primal-dual interior point method for unstructured convex and nonlinear programs. In these cases, the dual problem can guide the optimization algorithm improving the convergence rate of the primal solution. In some other cases, the dual can also have a special structure enabling more efficient algorithms or faster convergence. In such cases, it's also common that there exist an easy way to get the optimal primal solution x^* corresponding to the optimal dual solution y^*. In this case, it might make sense to solve the dual problem alone instead of the primal one. We will see examples of such duality being exploited for algorithmic reasons in optimal transport.

C.4.3   Deriving the dual problem (general)

To derive the dual of an optimization problem, the following steps should be followed:

1) Derive the Lagrangian function
2) Derive the Lagrangian relaxation problem
3) Derive the Lagrangian dual function
4) Derive the Lagrangian dual problem
5) Optional: simplify the dual problem

The **Lagragnian function** is a function of the objective and some (or all) of the constraints of the primal optimization problem such that the minimum of the Lagrangian function, subject to the non-relaxed constraints, is a lower bound on the optimal value of the primal problem. The standard way to construct the Lagrangian function is by relaxing some or all of the constraints in the primal problem and adding a function of them to the objective that ensure the above property. For instance, let the primal problem be an inequality constrained optimization problem:

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & f(x) \\ \text{subject to} \quad & \\ & g(x) \le 0, \\ & x \in X \end{aligned} \quad (39)$$

where f is the objective function, g is the constraint function, and X encodes all the "other constraints" if any exist or just the Euclidean space if no other constraints exist. The Lagrangian function resulting from relaxing the inequality constraint would be:

$$L(x, \lambda) = f(x) + \langle \lambda, g(x) \rangle \qquad (40)$$

where $\lambda \in \mathcal{R}^+$ is a fixed vector of non-negative coefficients known as the **Lagrangian multipliers** or the **dual decision variables**. The so-called **Lagrangian relaxation** of the primal problem would then be a parametric optimization problem parameterized by λ as follows:

$$\begin{aligned}
&\underset{x}{\text{minimize}} && L(x, \lambda) = f(x) + \langle \lambda, g(x) \rangle \\
&\text{subject to} && \\
&&& x \in X
\end{aligned} \qquad (41)$$

Since any solution $x \in X$ that is *feasible to the primal problem* will satisfy $g(x) \leq 0$ and since λ is non-negative, $L(x, \lambda) \leq f(x)$ is true for all primal feasible (i.e. satisfying $g(x) \leq 0$) $x \in X$. This is true for the optimal solution to the primal problem x^* as well, i.e. $L(x^*, \lambda) \leq f(x^*)$.

Let the minimum objective value of the Lagrangian relaxation problem for some choice of λ be:

$$L^*(\lambda) = \min_{x \in X} L(x, \lambda) \qquad (42)$$

$L^*(\lambda)$ is known as the **Lagrangian dual function**. It is important to note that the $x \in X$ that minimizes $L(x, \lambda)$ in the Lagrangian relaxation problem may or may not satisfy the inequality constraint $g(x) \leq 0$.

The condition of the Lagrangian dual function is that $L^*(\lambda) \leq f(x^*)$ is satisfied for any $\lambda \in \mathcal{R}^+$. A sufficient condition to ensure this property is satisfied is that $L(x, \lambda) \leq f(x)$ for any x that is feasible to the primal problem since $L^*(\lambda) \leq L(x^*, \lambda) \leq f(x^*)$ would then be true. For different constraint types: 1) appropriate Lagrangian multipliers λ, 2) multiplier domains D_L (instead of $\mathcal{R}^+$), and 3) Lagrangian functions are defined to ensure this property. This property guarantees that minimizing the Lagrangian relaxation problem gives us a lower bound $L^*(\lambda)$ on the optimal value $f(x^*)$. This should be true for every choice of $\lambda \in D_L$ where D_L encodes the constraints on the Lagrangian dual variables in the dual problem.

Since for every $\lambda \in D_L$, $L^*(\lambda) \leq f(x^*)$, the choice of $\lambda \in D_L$ that gives the tightest bound is the one that maximizes $L^*(\lambda)$. This optimization problem in λ is known as the **dual problem**.

$$\begin{aligned}
&\underset{\lambda}{\text{maximize}} && L^*(\lambda) = \min_{x \in X} L(x, \lambda) \\
&\text{subject to} && \\
&&& \lambda \in D_L
\end{aligned} \qquad (43)$$

This is a max-min optimization problem which from a first look may seem more complicated than the primal problem. However, for classes of optimization problems where the inner minimization of the Lagrangian function wrt x has a closed form solution, the dual problem can be significantly simplified. We will see an example for linear programs in the next section.

When no closed form solution exists for the Lagrangian relaxation, instead of solving the dual problem as a max-min problem, primal-dual algorithms solve for x and λ simultaneously. For instance in the primal-dual interior point method, the Lagrangian relaxation problem's and dual problem's optimality conditions are written out as a system of (potentially nonlinear) equations in terms of x and λ which are then solved using a variation of Newton's method.

The Lagrangian dual feasible domain D_L can more generally be defined as the Cartesian product of the domains of the dual variables associated with all the relaxed constraints. Each relaxed set of constraints has a dual variable domain associated with it to ensure the lower bound property of the Lagrangian. The following are some of the most common dual variable domains:

Constraint type	Dual variable domain
$g(x) \leq 0$	$\mathcal{R}^+$: non-negative numbers
$g(x) \geq 0$	$\mathcal{R}^-$: non-positive numbers
$g(x) = 0$	$\mathcal{R}$: all real numbers
$x \in \mathcal{C}$ where $\mathcal{C}$ is a cone	The polar cone $\mathcal{C}^o$ of $\mathcal{C}$

The polar cone of a cone $\mathcal{C}$ is the set of points $\{\lambda\}$ in the Euclidean space such that the inner product $\langle \lambda, x \rangle \leq 0 \, \forall x \in \mathcal{C}$. This definition ensures that the Lagrangian is a lower bound on the objective function for any feasible x.

C.5 Dual of a general linear program

In this section, we will apply the general steps from the previous section to a linear program in its canonical form. Consider the following primal linear program:

$$\begin{aligned}
&\underset{x}{\text{minimize}} && \langle c, x \rangle \\
&\text{subject to} && \\
&&& Ax \geq b
\end{aligned} \qquad (44)$$

for some constants c, A and b of the appropriate dimensions where $\langle .. \rangle$ is the inner product operator. There is no loss in generality in assuming the above form. Any $\leq$ constraints can be trivially transformed to $\geq$ constraints by multiplying the constraints by -1. Additionally, equality constraints of the form $Cx = d$ can be transformed to the following 2 sets of constraints:

1) $Cx \geq d$
2) $Cx \leq d$ which is equivalent to $-Cx \geq -d$

Therefore setting $A = \begin{bmatrix} C \\ -C \end{bmatrix}$ and $b = \begin{bmatrix} d \\ -d \end{bmatrix}$, we recover the canonical form.

The Lagrangian relaxation of the above linear program is:

$$\begin{aligned}
&\underset{x}{\text{minimize}} && L(x, \lambda) = \langle c, x \rangle + \langle \lambda, b - Ax \rangle \\
&&& = \langle x, c - A^T \lambda \rangle + \langle b, \lambda \rangle \\
&\text{subject to} && \\
&&& \lambda \in \mathcal{R}^+
\end{aligned} \qquad (45)$$

where λ is the vector of Lagrangian multipliers. This is just minimizing an affine function of x. The optimal value

$L^*(\lambda)$ of the above relaxation has the following closed form solution:

$$L^*(\lambda) = \begin{cases} \langle b, \lambda \rangle & \text{if } c - A^T\lambda = 0 \\ -\infty & \text{otherwise} \end{cases} \quad (46)$$

The general form of the Lagrangian dual problem can be written as:

$$\begin{aligned} \underset{\lambda}{\text{maximize}} \quad & \min_x L(x, \lambda) \\ \text{subject to} \quad & \\ & \lambda \geq 0 \end{aligned} \quad (47)$$

but it can then be simplified given the closed form solution as follows:

$$\begin{aligned} \underset{\lambda}{\text{maximize}} \quad & \langle b, \lambda \rangle \\ \text{subject to} \quad & \\ & A^T\lambda = c, \\ & \lambda \geq 0 \end{aligned} \quad (48)$$

Note that the $-\infty$ branch when $c - A^T\lambda \neq 0$ can be ignored in the simplified form since $-\infty$ is clearly less than $\langle b, \lambda \rangle$ and we are only interested in the λ that maximizes the dual function $L^*(\lambda)$. Also note that the simplification led to some additional constraints on λ compared to the general Lagrangian dual formulation. The above problem is known as the dual linear program. Similar derivation steps can be followed for conic and other special classes of optimization problems to derive their dual formulation.

In some cases, it can be more efficient to solve the dual linear program instead of the primal one. For linear programs, strong duality exists and there exists an efficient way to recover the optimal primal solution x^* from the optimal dual solution λ^* so one has the choice to solve either the primal or dual problems. The dual can also be exploited in algorithms directly without changing the problem's formulation. For instance, the dual Simplex method is an efficient way to perform the regular Simplex method on the dual formulation but without having to explicitly derive the dual formulation.

C.6 Dual of exact optimal transport problem

In this section, we shall derive the dual problem of the exact discrete optimal transport problem:

$$\begin{aligned} \underset{P}{\text{minimize}} \quad & \langle P, M \rangle \\ \text{subject to} \quad & \\ & P\mathbf{1} = r, \\ & P^T\mathbf{1} = c, \\ & P_{ij} \geq 0 \quad \forall i, j \end{aligned} \quad (49)$$

The Lagrangian relaxation of the above linear program is:

$$\begin{aligned} \underset{P}{\text{minimize}} \quad & \langle P, M \rangle + \langle \lambda_r, r - P\mathbf{1} \rangle + \langle \lambda_c, c - P^T\mathbf{1} \rangle \\ & + \langle \lambda_-, P \rangle \\ \text{subject to} \quad & \\ & \lambda_r \in \mathcal{R}, \\ & \lambda_c \in \mathcal{R}, \\ & \lambda_- \in \mathcal{R}^- \end{aligned} \quad (50)$$

Note that $\langle a, Pb \rangle$ can be re-written as $\langle P, ab^T \rangle$ for any vector a and b:

$$\begin{aligned} \langle a, Pb \rangle &= a^T Pb \\ &= \text{tr}(a^T Pb) \\ &= \text{tr}(Pba^T) \\ &= \langle P^T, ba^T \rangle \\ &= \langle P, ab^T \rangle \end{aligned}$$

Similarly, $\langle a, P^T b \rangle = \langle P, ba^T \rangle$.

Re-writing the Lagrangian relaxation using these identities and removing the $\in \mathcal{R}$ constraints:

$$\begin{aligned} \underset{P}{\text{minimize}} \quad & L(P, \lambda_r, \lambda_c, \lambda_-) \\ &= \langle P, M - \lambda_r \mathbf{1}^T - \mathbf{1}\lambda_c^T + \lambda_- \rangle \\ & \quad + \langle \lambda_r, r \rangle + \langle \lambda_c, c \rangle \\ \text{subject to} \quad & \\ & \lambda_- \in \mathcal{R}^- \end{aligned} \quad (51)$$

This is just minimizing an affine function of P. The optimal value $L^*(\lambda_r, \lambda_c, \lambda_-)$ of the above relaxation has the following closed form solution:

$$L^*(\lambda_r, \lambda_c, \lambda_-) = \\ \begin{cases} \langle \lambda_r, r \rangle + \langle \lambda_c, c \rangle & \text{if } M - \lambda_r \mathbf{1}^T - \mathbf{1}\lambda_c^T + \lambda_- = \mathbf{0} \\ -\infty & \text{otherwise} \end{cases} \quad (52)$$

The general form of the Lagrangian dual problem can be written as:

$$\begin{aligned} \underset{\lambda_r, \lambda_c, \lambda_-}{\text{maximize}} \quad & \min_P L(P, \lambda_r, \lambda_c, \lambda_-) \\ \text{subject to} \quad & \\ & \lambda_- \in \mathcal{R}^- \end{aligned} \quad (53)$$

but it can then be simplified given the closed form solution as follows:

$$\begin{aligned} \underset{\lambda_r, \lambda_c, \lambda_-}{\text{maximize}} \quad & \langle \lambda_r, r \rangle + \langle \lambda_c, c \rangle \\ \text{subject to} \quad & \\ & M - \lambda_r \mathbf{1}^T - \mathbf{1}\lambda_c^T + \lambda_- = \mathbf{0}, \\ & \lambda_- \in \mathcal{R}^- \end{aligned} \quad (54)$$

Simplifying this further by eliminating λ_- from the constraints, we can write the same problem using inequality constraints:

$$\begin{aligned} \underset{\lambda_r, \lambda_c}{\text{maximize}} \quad & \langle \lambda_r, r \rangle + \langle \lambda_c, c \rangle \\ \text{subject to} \quad & \\ & M - \lambda_r \mathbf{1}^T - \mathbf{1}\lambda_c^T \geq 0 \end{aligned} \quad (55)$$

The main advantage of the dual optimal transport problem is that the decision variables now scale linearly with the number of points in the supports of the measures $O(m+n)$. Additionally, the local constraints in the dual formulation over each (x_i, y_j) pair enable stochastic optimization methods.

C.7 Dual of entropy regularized optimal transport problem

The final example of duality and its application in optimal transport which we shall present is that of the entropy regularized optimal transport problem:

$$\underset{\boldsymbol{P}}{\text{minimize}} \quad \langle \boldsymbol{P}, \boldsymbol{M} \rangle - \lambda H(\boldsymbol{P})$$
$$\text{subject to}$$
$$\boldsymbol{P}\mathbf{1} = \boldsymbol{r}, \tag{56}$$
$$\boldsymbol{P}^T\mathbf{1} = \boldsymbol{c}$$

where $H(\boldsymbol{P}) = -\sum_{ij} P_{ij}(\log P_{ij}-1)$. The partial derivative of H wrt P_{ij} is:

$$\frac{\partial H}{\partial P_{ij}} = -\log P_{ij}$$

The Lagrangian relaxation of the above convex program is:

$$\underset{\boldsymbol{P}}{\text{minimize}} \quad \langle \boldsymbol{P}, \boldsymbol{M} \rangle - \lambda H(\boldsymbol{P})$$
$$+ \langle \boldsymbol{\lambda_r}, \boldsymbol{r} - \boldsymbol{P}\mathbf{1} \rangle + \langle \boldsymbol{\lambda_c}, \boldsymbol{c} - \boldsymbol{P}^T\mathbf{1} \rangle$$
$$\text{subject to} \tag{57}$$
$$\boldsymbol{\lambda_r} \in \mathcal{R},$$
$$\boldsymbol{\lambda_c} \in \mathcal{R}$$

which can be re-written as:

$$\underset{\boldsymbol{P}}{\text{minimize}} \quad L(\boldsymbol{P}, \boldsymbol{\lambda_r}, \boldsymbol{\lambda_c})$$
$$= \langle \boldsymbol{P}, \boldsymbol{M} - \boldsymbol{\lambda_r}\mathbf{1}^T - \mathbf{1}\boldsymbol{\lambda_c}^T \rangle$$
$$+ \langle \boldsymbol{\lambda_r}, \boldsymbol{r} \rangle + \langle \boldsymbol{\lambda_c}, \boldsymbol{c} \rangle \tag{58}$$
$$- \lambda H(\boldsymbol{P})$$

This is just minimizing an unconstrained strongly convex function of $\boldsymbol{P}$. The unique optimal solution $\boldsymbol{P}^*(\boldsymbol{\lambda_r}, \boldsymbol{\lambda_c})$ minimizing L can be obtained using the stationarity conditions:

$$\lambda \log P_{ij}^* + M_{ij} - \lambda_{ri} - \lambda_{cj} = 0$$
$$P_{ij}^* = e^{\frac{1}{\lambda}(\lambda_{ri} + \lambda_{cj} - M_{ij})}$$

The corresponding optimal value $L^*(\boldsymbol{\lambda_r}, \boldsymbol{\lambda_c})$ is therefore:

$$L^*(\boldsymbol{\lambda_r}, \boldsymbol{\lambda_c}) = \langle \boldsymbol{\lambda_r}, \boldsymbol{r} \rangle + \langle \boldsymbol{\lambda_c}, \boldsymbol{c} \rangle$$
$$+ \sum_{ij} P_{ij}^* \times (M_{ij} - \lambda_{ri} - \lambda_{cj} + \lambda \log P_{ij}^* - \lambda) \tag{59}$$

Since $-\lambda \log P_{ij}^* = M_{ij} - \lambda_{ri} - \lambda_{cj}$, the above expression simplifies to:

$$L^*(\boldsymbol{\lambda_r}, \boldsymbol{\lambda_c}) = \langle \boldsymbol{\lambda_r}, \boldsymbol{r} \rangle + \langle \boldsymbol{\lambda_c}, \boldsymbol{c} \rangle - \lambda \sum_{ij} P_{ij}^*$$
$$= \langle \boldsymbol{\lambda_r}, \boldsymbol{r} \rangle + \langle \boldsymbol{\lambda_c}, \boldsymbol{c} \rangle - \lambda \sum_{ij} e^{\frac{1}{\lambda}(\lambda_{ri} + \lambda_{cj} - M_{ij})}$$

Let $\boldsymbol{R}_\lambda(\boldsymbol{\lambda_r}, \boldsymbol{\lambda_c})$ be the matrix whose $(i, j)^{th}$ element is $-\lambda e^{\frac{1}{\lambda}(\lambda_{ri} + \lambda_{cj} - M_{ij})}$. The Lagrangian dual problem can therefore be written as:

$$\underset{\boldsymbol{\lambda_r}, \boldsymbol{\lambda_c}}{\text{maximize}} \quad \langle \boldsymbol{\lambda_r}, \boldsymbol{r} \rangle + \langle \boldsymbol{\lambda_c}, \boldsymbol{c} \rangle + \langle \boldsymbol{R}_\lambda(\boldsymbol{\lambda_r}, \boldsymbol{\lambda_c}), \mathbf{1} \rangle \tag{60}$$

This is an unconstrained maximization problem.

REFERENCES

[1] S. Dhouib, I. Redko, and C. Lartizien, "Margin-aware adversarial domain adaptation with optimal transport," in *International conference on machine learning*, pp. 2514–2524, PMLR, 2020.

[2] Z. Ge, S. Liu, Z. Li, O. Yoshie, and J. Sun, "Ota: Optimal transport assignment for object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 303–312, 2021.

[3] S. Haldar, V. Mathur, D. Yarats, and L. Pinto, "Watch and match: Supercharging imitation with regularized optimal transport," in *Conference on Robot Learning*, pp. 32–43, PMLR, 2023.

[4] Y. Dong and W. Sawin, "Copt: Coordinated optimal transport on graphs," *Advances in Neural Information Processing Systems*, vol. 33, pp. 19327–19338, 2020.

[5] A. Wijesinghe, Q. Wang, and S. Gould, "A regularized wasserstein framework for graph kernels," in *2021 IEEE International Conference on Data Mining (ICDM)*, pp. 739–748, IEEE, 2021.

[6] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, "Loftr: Detector-free local feature matching with transformers," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8922–8931, 2021.

[7] C.-Y. Chuang, Y. Mroueh, K. Greenewald, A. Torralba, and S. Jegelka, "Measuring generalization with optimal transport," *Advances in Neural Information Processing Systems*, vol. 34, pp. 8294–8306, 2021.

[8] C.-W. Huang, R. T. Chen, C. Tsirigotis, and A. Courville, "Convex potential flows: Universal probability distributions with optimal transport and convex optimization," in *International Conference on Learning Representations*, 2021.

[9] L. Chen, D. Wang, Z. Gan, J. Liu, R. Henao, and L. Carin, "Wasserstein contrastive representation distillation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16296–16305, 2021.

[10] P. Gordaliza, E. Del Barrio, G. Fabrice, and J.-M. Loubes, "Obtaining fairness using optimal transport theory," in *International Conference on Machine Learning*, pp. 2357–2365, PMLR, 2019.

[11] M. Buyl and T. De Bie, "Optimal transport of classifiers to fairness," in *Advances in Neural Information Processing Systems*, 2022.

[12] L. V. Kantorovich, "Mathematical methods of organizing and planning production," *Management science*, vol. 6, no. 4, pp. 366–422, 1960.

[13] L. V. Kantorovich, "On a problem of monge," *Journal of Mathematical Sciences*, vol. 4, no. 133, pp. 1383–1383, 2006.

[14] M. Kusner, Y. Sun, N. Kolkin, and K. Weinberger, "From word embeddings to document distances," in *International conference on machine learning*, pp. 957–966, PMLR, 2015.

[15] M. Werner and E. Laber, "Speeding up word mover's distance and its variants via properties of distances between embeddings," in *ECAI 2020*, pp. 2204–2211, IOS Press, 2020.

[16] M. Yurochkin, S. Claici, E. Chien, F. Mirzazadeh, and J. M. Solomon, "Hierarchical optimal transport for document representation," *Advances in neural information processing systems*, vol. 32, 2019.

[17] C. Frogner, J. Solomon, and F. Mirzazadeh, "Learning embeddings into entropic wasserstein spaces," in *International Conference on Learning Representations*, International Conference on Learning Representations, ICLR, 2019.

[18] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International conference on machine learning*, pp. 214–223, PMLR, 2017.

[19] G. Peyré, M. Cuturi, *et al.*, "Computational optimal transport: With applications to data science," *Foundations and Trends® in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019.

[20] C. Bunne, L. Papaxanthos, A. Krause, and M. Cuturi, "Proximal optimal transport modeling of population dynamics," in *International Conference on Artificial Intelligence and Statistics*, pp. 6511–6528, PMLR, 2022.

[21] N. Bonneel, G. Peyré, and M. Cuturi, "Wasserstein barycentric coordinates: Histogram regression using optimal transport," *ACM Transactions on Graphics*, vol. 35, no. 4, pp. 71–1, 2016.

[22] Z. Shen, J. Feydy, P. Liu, A. H. Curiale, R. San Jose Estepar, R. San Jose Estepar, and M. Niethammer, "Accurate point cloud registration with robust optimal transport," *Advances in Neural Information Processing Systems*, vol. 34, pp. 5373–5389, 2021.

[23] G. Puy, A. Boulch, and R. Marlet, "Flot: Scene flow on point clouds guided by optimal transport," in *16th European Conference on Computer Vision*, pp. 527–544, Springer, 2020.

[24] J. WEED and F. BACH, "Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance," *Bernoulli*, vol. 25, no. 4A, pp. 2620–2648, 2019.

[25] J. LEI, "Convergence and concentration of empirical measures under wasserstein distance in unbounded functional spaces," *Bernoulli*, vol. 26, no. 1, pp. 767–798, 2020.

[26] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 723–773, 2012.

[27] D. Mukherjee, A. Guha, J. M. Solomon, Y. Sun, and M. Yurochkin, "Outlier-robust optimal transport," in *International Conference on Machine Learning*, pp. 7850–7860, PMLR, 2021.

[28] Y. C. F. Lim, L. Wynter, and S. H. Lim, "Order constraints in optimal transport," in *International Conference on Machine Learning*, pp. 13313–13333, PMLR, 2022.

[29] S. Dhouib, I. Redko, T. Kerdoncuff, R. Emonet, and M. Sebban, "A swiss army knife for minimax optimal transport," in *International Conference on Machine Learning*, pp. 2504–2513, PMLR, 2020.

[30] L. Zheng, Y. Xiao, and L. Niu, "A brief survey on computational gromov-wasserstein distance," *Procedia Computer Science*, vol. 199, pp. 697–702, 2022.

[31] R. Sato, M. Yamada, and H. Kashima, "Fast unbalanced optimal transport on a tree," *Advances in neural information processing systems*, vol. 33, pp. 19039–19051, 2020.

[32] A. Khamis, B. Kusy, C. T. Chou, M.-L. McLaws, and W. Hu, "Rfwash: a weakly supervised tracking of hand hygiene technique," in *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pp. 572–584, 2020.

[33] P. Papadakis, "The canonically posed 3d objects dataset," in *Eurographics Workshop on 3D Object Retrieval*, pp. 33–36, 2014.

[34] F. Santambrogio *et al.*, "Optimal transport for applied mathematicians: calculus of variations, pdes, and modeling," *Springer Mathematics and Statistics eBooks 2015 English/International*, 2015.

[35] C. Villani *et al.*, *Optimal Transport: Old and New*, vol. 338. Springer, 2008.

[36] C. Villani, *Topics in optimal transportation*, vol. 58. American Mathematical Soc., 2021.

[37] M. Cuturi, "Sinkhorn Distances: Lightspeed Computation of Optimal Transport," *Advances in Neural Information Processing Systems*, 2013.

[38] Y. Chen, T. T. Georgiou, and M. Pavon, "Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrodinger bridge," *Siam Review*, vol. 63, no. 2, pp. 249–313, 2021.

[39] V. De Bortoli, J. Thornton, J. Heng, and A. Doucet, "Diffusion schrödinger bridge with applications to score-based generative modeling," *Advances in Neural Information Processing Systems*, vol. 34, pp. 17695–17709, 2021.

[40] E. Tamir, M. Trapp, and A. Solin, "Transport with support: Data-conditional diffusion bridges," *arXiv preprint arXiv:2301.13636*, 2023.

[41] S. Ferradans, N. Papadakis, G. Peyré, and J.-F. Aujol, "Regularized discrete optimal transport," *SIAM Journal on Imaging Sciences*, vol. 7, no. 3, pp. 1853–1882, 2014.

[42] F.-P. Paty and M. Cuturi, "Regularized optimal transport is ground cost adversarial," in *International Conference on Machine Learning*, pp. 7532–7542, PMLR, 2020.

[43] P. A. Knight, "The sinkhorn–knopp algorithm: convergence and applications," *SIAM Journal on Matrix Analysis and Applications*, vol. 30, no. 1, pp. 261–275, 2008.

[44] J. Altschuler, F. Bach, A. Rudi, and J. Niles-Weed, "Massively scalable sinkhorn distances via the nyström method," *Advances in neural information processing systems*, vol. 32, 2019.

[45] J. Feydy, T. Séjourné, F.-X. Vialard, S.-i. Amari, A. Trouvé, and G. Peyré, "Interpolating between optimal transport and mmd using sinkhorn divergences," in *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2681–2690, PMLR, 2019.

[46] T. Séjourné, J. Feydy, F.-X. Vialard, A. Trouvé, and G. Peyré, "Sinkhorn divergences for unbalanced optimal transport," *arXiv preprint arXiv:1910.12958*, 2019.

[47] A. Genevay, L. Chizat, F. Bach, M. Cuturi, and G. Peyré, "Sample complexity of sinkhorn divergences," in *The 22nd international conference on artificial intelligence and statistics*, pp. 1574–1583, PMLR, 2019.

[48] M. Scetbon and M. Cuturi, "Linear time sinkhorn divergences using positive features," *Advances in Neural Information Processing Systems*, vol. 33, pp. 13468–13480, 2020.

[49] Y. Asano, C. Rupprecht, and A. Vedaldi, "Self-labelling via simultaneous clustering and representation learning," in *International Conference on Learning Representations*, 2019.

[50] E. Wong, F. Schmidt, and Z. Kolter, "Wasserstein adversarial examples via projected sinkhorn iterations," in *International Conference on Machine Learning*, pp. 6808–6817, PMLR, 2019.

[51] G. Mena, J. Snoek, S. Linderman, and D. Belanger, "Learning latent permutations with gumbel-sinkhorn networks," in *ICLR 2018 Conference Track*, vol. 2018, OpenReview, 2018.

[52] R. Santa Cruz, B. Fernando, A. Cherian, and S. Gould, "Visual permutation learning," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 12, pp. 3100–3114, 2018.

[53] M. Fey, J. E. Lenssen, C. Morris, J. Masci, and N. M. Kriege, "Deep graph matching consensus," in *International Conference on Learning Representations*, 2020.

[54] M. Cuturi, O. Teboul, and J.-P. Vert, "Differentiable Ranking and Sorting using Optimal Transport," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[55] R. P. Adams and R. S. Zemel, "Ranking via sinkhorn propagation," *stat*, vol. 1050, p. 14, 2011.

[56] M. E. Sander, P. Ablin, M. Blondel, and G. Peyré, "Sinkformers: Transformers with doubly stochastic attention," in *International Conference on Artificial Intelligence and Statistics*, pp. 3515–3530, PMLR, 2022.

[57] R. Flamary, N. Courty, D. Tuia, and A. Rakotomamonjy, "Optimal transport for domain adaptation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 1, 2016.

[58] R. Flamary, N. Courty, A. Rakotomamonjy, and D. Tuia, "Optimal transport with laplacian regularization," in *NIPS 2014, Workshop on Optimal Transport and Machine Learning*, 2014.

[59] M. Blondel, V. Seguy, and A. Rolet, "Smooth and sparse optimal transport," in *International conference on artificial intelligence and statistics*, pp. 880–889, PMLR, 2018.

[60] B. Su and G. Hua, "Order-preserving wasserstein distance for sequence matching," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1049–1057, 2017.

[61] S. Kumar, S. Haresh, A. Ahmed, A. Konin, M. Z. Zia, and Q.-H. Tran, "Unsupervised action segmentation by joint representation learning and online clustering," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20174–20185, 2022.

[62] B. Schmitzer, "Stabilized sparse scaling algorithms for entropy regularized transport problems," *SIAM Journal on Scientific Computing*, vol. 41, no. 3, pp. A1443–A1481, 2019.

[63] Q. C. Le, D. Conte, and M. Hidane, "Unbalanced optimal transport in multi-camera tracking applications," in *Pattern Recognition. ICPR International Workshops and Challenges: Virtual Event, January 10–15, 2021, Proceedings, Part V*, pp. 327–343, Springer, 2021.

[64] J. Wan, Z. Liu, and A. B. Chan, "A generalized loss function for crowd counting and localization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1974–1983, 2021.

[65] Z. Ma, X. Wei, X. Hong, H. Lin, Y. Qiu, and Y. Gong, "Learning to count via unbalanced optimal transport," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 2319–2327, 2021.

[66] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard, "Unbalanced optimal transport: Dynamic and kantorovich formulations," *Journal of Functional Analysis*, vol. 274, no. 11, pp. 3090–3123, 2018.

[67] M. Liero, A. Mielke, and G. Savaré, "Optimal entropy-transport problems and a new hellinger–kantorovich distance between positive measures," *Inventiones mathematicae*, vol. 211, no. 3, pp. 969–1117, 2018.

[68] T. Séjourné, G. Peyré, and F.-X. Vialard, "Unbalanced optimal transport, from theory to numerics," *arXiv preprint arXiv:2211.08775*, 2022.

[69] P. Manupriya, J. Saketha Nath, and P. Jawanpuria, "MMD-regularized Unbalanced Optimal Transport," *arXiv e-prints*, p. arXiv:2011.05001, Nov. 2020.

[70] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superglue: Learning feature matching with graph neural networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4938–4947, 2020.

[71] Z. Dang, F. Wang, and M. Salzmann, "Learning 3d-3d correspondences for one-shot partial-to-partial registration," *arXiv preprint arXiv:2006.04523*, 2020.

[72] L. Chapel, M. Z. Alaya, and G. Gasso, "Partial optimal tranport with applications on positive-unlabeled learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 2903–2913, 2020.

[73] K. Pham, K. Le, N. Ho, T. Pham, and H. Bui, "On unbalanced optimal transport: An analysis of sinkhorn algorithm," in *International Conference on Machine Learning*, pp. 7673–7682, PMLR, 2020.

[74] N. Bonneel, J. Rabin, G. Peyré, and H. Pfister, "Sliced and radon wasserstein barycenters of measures," *Journal of Mathematical Imaging and Vision*, vol. 51, pp. 22–45, 2015.

[75] N. Bonnotte, *Unidimensional and evolution methods for optimal transportation*. PhD thesis, Université Paris Sud-Paris XI; Scuola normale superiore (Pise, Italie), 2013.

[76] E. Bayraktar and G. Guo, "Strong equivalence between metrics of wasserstein type," *arXiv e-prints*, pp. arXiv–1912, 2019.

[77] K. Nadjahi, A. Durmus, U. Simsekli, and R. Badeau, "Asymptotic guarantees for learning generative models with the sliced-wasserstein distance," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[78] K. Nadjahi, A. Durmus, L. Chizat, S. Kolouri, S. Shahrampour, and U. Simsekli, "Statistical and topological properties of sliced probability divergences," *Advances in Neural Information Processing Systems*, vol. 33, pp. 20802–20812, 2020.

[79] S. Kolouri, Y. Zou, and G. K. Rohde, "Sliced wasserstein kernels for probability distributions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5258–5267, 2016.

[80] I. Deshpande, Y.-T. Hu, R. Sun, A. Pyrros, N. Siddiqui, S. Koyejo, Z. Zhao, D. Forsyth, and A. G. Schwing, "Max-sliced wasserstein distance and its use for gans," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10648–10656, 2019.

[81] E. Heitz, K. Vanhoey, T. Chambon, and L. Belcour, "A sliced wasserstein loss for neural texture synthesis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9412–9420, 2021.

[82] N. Naderializadeh, J. F. Comer, R. Andrews, H. Hoffmann, and S. Kolouri, "Pooling by sliced-wasserstein embedding," *Advances in Neural Information Processing Systems*, vol. 34, pp. 3389–3400, 2021.

[83] B.-H. Tran, S. Rossi, D. Milios, P. Michiardi, E. V. Bonilla, and M. Filippone, "Model selection for bayesian autoencoders," *Advances in Neural Information Processing Systems*, vol. 34, pp. 19730–19742, 2021.

[84] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10478–10487, 2021.

[85] C.-Y. Lee, T. Batra, M. H. Baig, and D. Ulbricht, "Sliced wasserstein discrepancy for unsupervised domain adaptation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10285–10295, 2019.

[86] V. Titouan, R. Flamary, N. Courty, R. Tavenard, and L. Chapel, "Sliced gromov-wasserstein," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[87] B. Moiseev, A. Konev, A. Chigorin, and A. Konushin, "Evaluation of traffic sign recognition methods trained on synthetically generated data," in *Advanced Concepts for Intelligent Vision Systems: 15th International Conference, ACIVS 2013, Poznań, Poland, October 28-31, 2013. Proceedings 15*, pp. 576–583, Springer, 2013.

[88] J. Stallkamp, M. Schlipsing, J. Salmen, and C. Igel, "The german traffic sign recognition benchmark: a multi-class classification competition," in *The 2011 international joint conference on neural networks*, pp. 1453–1460, IEEE, 2011.

[89] K. Nguyen, N. Ho, T. Pham, and H. Bui, "Distributional sliced-wasserstein and applications to generative modeling," in *International Conference on Learning Representations*, 2021.

[90] S. Kolouri, K. Nadjahi, U. Simsekli, R. Badeau, and G. Rohde, "Generalized sliced wasserstein distances," *Advances in neural information processing systems*, vol. 32, 2019.

[91] M. Agueh and G. Carlier, "Barycenters in the wasserstein space," *SIAM Journal on Mathematical Analysis*, vol. 43, no. 2, pp. 904–924, 2011.

[92] H. Janati, M. Cuturi, and A. Gramfort, "Debiased sinkhorn barycenters," in *International Conference on Machine Learning*, pp. 4692–4701, PMLR, 2020.

[93] J. Fan, A. Taghvaei, and Y. Chen, "Scalable computations of wasserstein barycenter via input convex neural networks," in *International Conference on Machine Learning*, pp. 1571–1581, PMLR, 2021.

[94] T.-D. Truong, C. N. Duong, H. A. Pham, B. Raj, N. Le, K. Luu, *et al.*, "The right to talk: An audio-visual transformer approach," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1105–1114, 2021.

[95] D. Alvarez-Melis and T. Jaakkola, "Gromov-wasserstein alignment of word embedding spaces," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1881–1890, 2018.

[96] M. Gromov, "Groups of polynomial growth and expanding maps," *Publications Mathématiques de l'Institut des Hautes Scientifiques*, vol. 53, no. 1, pp. 53–78, 1981.

[97] F. Mémoli, "Gromov–wasserstein distances and the metric approach to object matching," *Foundations of computational mathematics*, vol. 11, pp. 417–487, 2011.

[98] T. Vayer, *A contribution to Optimal Transport on incomparable spaces*. PhD thesis, Lorient, 2020.

[99] C. Bunne, D. Alvarez-Melis, A. Krause, and S. Jegelka, "Learning generative models across incomparable spaces," in *International conference on machine learning*, pp. 851–861, PMLR, 2019.

[100] P. Demetci, R. Santorella, B. Sandstede, W. Stafford Noble, and R. Singh, "Gromov-Wasserstein Optimal Transport to Align Single-Cell Multi-Omics Data," in *ICML 2020 Workshop on Computational Biology (WCB)*, 2020.

[101] A. Fickinger, S. Cohen, S. Russell, and B. Amos, "Cross-domain imitation learning via optimal transport," in *International Conference on Learning Representations*, 2022.

[102] S. Cohen, G. Luise, A. Terenin, B. Amos, and M. Deisenroth, "Aligning time series on incomparable spaces," in *International Conference on Artificial Intelligence and Statistics*, pp. 1036–1044, PMLR, 2021.

[103] V. Titouan, N. Courty, R. Tavenard, and R. Flamary, "Optimal transport for structured data with application on graphs," in *International Conference on Machine Learning*, pp. 6275–6284, PMLR, 2019.

[104] H. Xu, D. Luo, and L. Carin, "Scalable gromov-wasserstein learning for graph partitioning and matching," *Advances in neural information processing systems*, vol. 32, 2019.

[105] H. Xu, D. Luo, H. Zha, and L. C. Duke, "Gromov-wasserstein learning for graph matching and node embedding," in *International conference on machine learning*, pp. 6932–6941, PMLR, 2019.

[106] L. Brogat-Motte, R. Flamary, C. Brouard, J. Rousu, and F. d'Alché Buc, "Learning to predict graphs with fused gromov-wasserstein barycenters," in *International Conference on Machine Learning*, pp. 2321–2335, PMLR, 2022.

[107] G. Peyré, M. Cuturi, and J. Solomon, "Gromov-wasserstein averaging of kernel and distance matrices," in *International conference on machine learning*, pp. 2664–2672, PMLR, 2016.

[108] J. Lee, M. Dabagia, E. Dyer, and C. Rozell, "Hierarchical optimal transport for multimodal distribution alignment," *Advances in neural information processing systems*, vol. 32, 2019.

[109] D. Alvarez-Melis and N. Fusi, "Geometric dataset distances via optimal transport," *Advances in Neural Information Processing Systems*, vol. 33, pp. 21428–21439, 2020.

[110] M. El Hamri, Y. Bennani, and I. Falih, "Hierarchical optimal transport for unsupervised domain adaptation," *Machine Learning*, vol. 111, no. 11, pp. 4159–4182, 2022.

[111] F. Taherkhani, A. Dabouei, S. Soleymani, J. Dawson, and N. M. Nasrabadi, "Transporting labels via hierarchical optimal transport for semi-supervised learning," in *ECCV 2020: 16th European Conference on Computer Vision*, pp. 509–526, 2020.

[112] N. Deb, P. Ghosal, and B. Sen, "Rates of estimation of optimal transport maps using plug-in estimators via barycentric projections," *Advances in Neural Information Processing Systems*, vol. 34, pp. 29736–29753, 2021.

[113] T. Manole, S. Balakrishnan, J. Niles-Weed, and L. Wasserman, "Plugin estimation of smooth optimal transport maps," *arXiv preprint arXiv:2107.12364*, 2021.

[114] R. M. Dudley, "The speed of mean glivenko-cantelli convergence," *The Annals of Mathematical Statistics*, vol. 40, no. 1, pp. 40–50, 1969.

[115] J. Niles-Weed and P. Rigollet, "Estimation of wasserstein distances in the spiked transport model," *Bernoulli*, vol. 28, no. 4, pp. 2663–2688, 2022.

[116] V. Divol, "A short proof on the rate of convergence of the empirical measure for the wasserstein distance," *arXiv e-prints*, pp. arXiv–2101, 2021.

[117] A. Mallasto, A. Gerolin, and H. Q. Minh, "Entropy-regularized 2-wasserstein distance between gaussian measures," *Information Geometry*, vol. 5, no. 1, pp. 289–323, 2022.

[118] E. del Barrio and J.-M. Loubes, "The statistical effect of entropic regularization in optimal transportation," *arXiv preprint arXiv:2006.05199*, 2020.

[119] H. Janati, B. Muzellec, G. Peyré, and M. Cuturi, "Entropic optimal transport between unbalanced gaussian measures has a closed form," *Advances in neural information processing systems*, vol. 33, pp. 10468–10479, 2020.

[120] K. Le, D. Q. Le, H. Nguyen, D. Do, T. Pham, and N. Ho, "Entropic gromov-wasserstein between gaussian distributions," in *International Conference on Machine Learning*, pp. 12164–12203, PMLR, 2022.

[121] M. Yamada, Y. Takezawa, R. Sato, H. Bao, Z. Kozareva, and S. Ravi, "Approximating 1-wasserstein distance with trees," *Transactions of Machine Learning Research*, 2023.

[122] Y. Takezawa, R. Sato, Z. Kozareva, S. Ravi, and M. Yamada, "Fixed support tree-sliced wasserstein barycenter," in *International Conference on Artificial Intelligence and Statistics*, pp. 1120–1137, PMLR, 2022.

[123] T. Le, M. Yamada, K. Fukumizu, and M. Cuturi, "Tree-sliced variants of wasserstein distances," *Advances in neural information processing systems*, vol. 32, 2019.

[124] Y. Takezawa, R. Sato, and M. Yamada, "Supervised tree-wasserstein distance," in *International Conference on Machine Learning*, pp. 10086–10095, PMLR, 2021.

[125] R. Flamary, N. Courty, A. Gramfort, M. Z. Alaya, A. Boisbunon, S. Chambon, L. Chapel, A. Corenflos, K. Fatras, N. Fournier, *et al.*, "Pot: Python optimal transport," *The Journal of Machine Learning Research*, vol. 22, no. 1, pp. 3571–3578, 2021.

[126] C. Bunne, A. Krause, *et al.*, "Supervised training of conditional monge maps," in *Advances in Neural Information Processing Systems*, 2022.

[127] J. Altschuler, S. Chewi, P. R. Gerber, and A. Stromme, "Averaging on the bures-wasserstein manifold: dimension-free convergence of gradient descent," *Advances in Neural Information Processing Systems*, vol. 34, pp. 22132–22145, 2021.

[128] A. Korotin, V. Egiazarian, L. Li, and E. Burnaev, "Wasserstein iterative networks for barycenter estimation," in *Advances in Neural Information Processing Systems*, 2022.

[129] F.-P. Paty and M. Cuturi, "Subspace robust wasserstein distances," in *International conference on machine learning*, pp. 5072–5081, PMLR, 2019.

[130] Z. Izzo, S. Silwal, and S. Zhou, "Dimensionality reduction for wasserstein barycenter," *Advances in neural information processing systems*, vol. 34, pp. 15582–15594, 2021.

[131] M. Huang, S. Ma, and L. Lai, "Projection robust wasserstein barycenters," in *International Conference on Machine Learning*, pp. 4456–4465, PMLR, 2021.

[132] B. Muzellec and M. Cuturi, "Subspace detours: Building transport plans that are optimal on subspace projections," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[133] C. Bonet, T. Vayer, N. Courty, F. Septier, and L. Drumetz, "Subspace detours meet gromov–wasserstein," *Algorithms*, vol. 14, no. 12, p. 366, 2021.

[134] M. Petrovich, C. Liang, R. Sato, Y. Liu, Y.-H. H. Tsai, L. Zhu, Y. Yang, R. Salakhutdinov, and M. Yamada, "Feature-robust optimal transport for high-dimensional data," in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2022, Grenoble, France, September 19–23, 2022, Proceedings, Part V*, pp. 291–307, Springer, 2023.

[135] C. Meng, Y. Ke, J. Zhang, M. Zhang, W. Zhong, and P. Ma, "Large-scale optimal transport map estimation using projection pursuit," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[136] J. Zhang, P. Ma, W. Zhong, and C. Meng, "Projection-based techniques for high-dimensional optimal transport problems," *WIREs Computational Statistics*, 2022.

[137] G. Beugnot, A. Genevay, K. Greenewald, and J. Solomon, "Improving approximate optimal transport distances using quanti-

[138] P. O. Aboagye, Y. Zheng, M. Yeh, J. Wang, Z. Zhuang, H. Chen, L. Wang, W. Zhang, and J. Phillips, "Quantized wasserstein procrustes alignment of word embedding spaces," in *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pp. 200–214, 2022.

[139] Z. Goldfeld and K. Greenewald, "Gaussian-smoothed optimal transport: Metric structure and statistical efficiency," in *International Conference on Artificial Intelligence and Statistics*, pp. 3327–3337, PMLR, 2020.

[140] Y. Ding and J. Niles-Weed, "Asymptotics of smoothed wasserstein distances in the small noise regime," in *Advances in Neural Information Processing Systems*, 2022.

[141] A. Rakotomamonjy and R. Liva, "Differentially private sliced wasserstein distance," in *International Conference on Machine Learning*, pp. 8810–8820, PMLR, 2021.

[142] Z. Goldfeld, K. Greenewald, and K. Kato, "Asymptotic guarantees for generative modeling based on the smooth wasserstein distance," *Advances in neural information processing systems*, vol. 33, pp. 2527–2539, 2020.

[143] Z. Goldfeld, K. Greenewald, J. Niles-Weed, and Y. Polyanskiy, "Convergence of smoothed empirical measures with applications to entropy estimation," *IEEE Transactions on Information Theory*, vol. 66, no. 7, pp. 4368–4391, 2020.

[144] M. Scetbon, M. Cuturi, and G. Peyré, "Low-rank sinkhorn factorization," in *International Conference on Machine Learning*, pp. 9344–9354, PMLR, 2021.

[145] W. Liu, C. Zhang, N. Zheng, and H. Qian, "Approximating optimal transport via low-rank and sparse factorization," *arXiv preprint arXiv:2111.06546*, 2021.

[146] M. Klatt, C. Tameling, and A. Munk, "Empirical regularized optimal transport: Statistical theory and applications," *SIAM Journal on Mathematics of Data Science*, vol. 2, no. 2, pp. 419–443, 2020.

[147] Y. Xie, X. Wang, R. Wang, and H. Zha, "A fast proximal point method for computing exact wasserstein distance," in *Uncertainty in artificial intelligence*, pp. 433–453, PMLR, 2020.

[148] B. Schmitzer, "Stabilized sparse scaling algorithms for entropy regularized transport problems," *SIAM Journal on Scientific Computing*, vol. 41, no. 3, pp. A1443–A1481, 2019.

[149] A. Forrow, J.-C. Hütter, M. Nitzan, P. Rigollet, G. Schiebinger, and J. Weed, "Statistical optimal transport via factored couplings," in *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2454–2465, PMLR, 2019.

[150] C.-H. Lin, M. Azabou, and E. L. Dyer, "Making transport more robust and interpretable by moving data through a small number of anchor points," *Proceedings of machine learning research*, vol. 139, p. 6631, 2021.

[151] M. Scetbon *et al.*, "Low-rank optimal transport: Approximation, statistics and debiasing," in *Advances in Neural Information Processing Systems*, 2022.

[152] J. Gasteiger, M. Lienen, and S. Günnemann, "Scalable optimal transport in high dimensions for graph distances, embedding alignment, and more," in *International Conference on Machine Learning*, pp. 5616–5627, PMLR, 2021.

[153] L. Chizat, P. Roussillon, F. Léger, F.-X. Vialard, and G. Peyré, "Faster wasserstein distance estimation with the sinkhorn divergence," *Advances in Neural Information Processing Systems*, vol. 33, pp. 2257–2269, 2020.

[154] A. Thibault, L. Chizat, C. H. Dossal, and N. Papadakis, "Over-relaxed sinkhorn-knopp algorithm for regularized optimal transport," in *NIPS Workshop on Optimal Transport & Machine Learning (OTML'17)*, 2017.

[155] J. Thornton and M. Cuturi, "Rethinking initialization of the sinkhorn algorithm," in *International Conference on Artificial Intelligence and Statistics*, pp. 8682–8698, PMLR, 2023.

[156] D. Alvarez-Melis, T. S. Jaakkola, and S. Jegelka, "Structured optimal transport," *International Conference on Artificial Intelligence and Statistics, AISTATS*, pp. 1771–1780, 2018.

[157] O. Pele and M. Werman, "Fast and robust earth mover's distances," in *2009 IEEE 12th international conference on computer vision*, pp. 460–467, IEEE, 2009.

[158] F. Elvander, I. Haasler, A. Jakobsson, and J. Karlsson, "Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion," *Signal Processing*, vol. 171, p. 107474, 2020.

[159] I. Haasler, A. Ringh, Y. Chen, and J. Karlsson, "Multimarginal optimal transport with a tree-structured cost and the schrodinger bridge problem," *SIAM Journal on Control and Optimization*, vol. 59, no. 4, pp. 2428–2453, 2021.

[160] Y. Chen, C. Li, and Z. Lu, "Computing wasserstein-p distance between images with linear cost," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 519–528, 2022.

[161] J. Franklin and J. Lorenz, "On the scaling of multidimensional matrices," *Linear Algebra and its applications*, vol. 114, pp. 717–735, 1989.

[162] C. Williams and M. Seeger, "Using the nyström method to speed up kernel machines," *Advances in neural information processing systems*, vol. 13, 2000.

[163] J. Wang, H. T. Shen, J. Song, and J. Ji, "Hashing for similarity search: A survey," *arXiv preprint arXiv:1408.2927*, 2014.

[164] S. Gerber and M. Maggioni, "Multiscale strategies for computing optimal transport," *Journal of Machine Learning Research*, vol. 18, no. 72, pp. 1–32, 2017.

[165] Q. Mérigot, "A multiscale approach to optimal transport," in *Computer Graphics Forum*, vol. 30, pp. 1584–1592, 2011.

[166] A. J. Blumberg, M. Carriere, M. A. Mandell, R. Rabadan, and S. Villar, "Mrec: a fast and versatile framework for aligning and matching point clouds with applications to single cell molecular data," *arXiv preprint arXiv:2001.01666*, 2020.

[167] S. Chowdhury, D. Miller, and T. Needham, "Quantized gromov-wasserstein," in *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part III 21*, pp. 811–827, Springer, 2021.

[168] F. Mémoli, A. Sidiropoulos, and K. Singhal, "Sketching and clustering metric measure spaces," *arXiv preprint arXiv:1801.00551*, 2018.

[169] F. Parés, D. G. Gasulla, A. Vilalta, J. Moreno, E. Ayguadé, J. Labarta, U. Cortés, and T. Suzumura, "Fluid communities: A competitive, scalable and diverse community detection algorithm," in *Complex Networks & Their Applications VI: Proceedings of Complex Networks 2017 (The Sixth International Conference on Complex Networks and Their Applications)*, pp. 229–240, Springer, 2018.

[170] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer networks and ISDN systems*, vol. 30, no. 1-7, pp. 107–117, 1998.

[171] M. Sommerfeld, J. Schrieber, Y. Zemel, and A. Munk, "Optimal transport: Fast probabilistic approximation with exact solvers.," *J. Mach. Learn. Res.*, vol. 20, no. 105, pp. 1–23, 2019.

[172] K. Fatras, *Deep learning and optimal transport: learning from one another*. PhD thesis, Université de Bretagne Sud, 2021.

[173] D. Alvarez-Melis, N. Fusi, L. Mackey, and T. Wagner, "Budget-constrained bounds for mini-batch estimation of optimal transport," *arXiv preprint arXiv:2210.13630*, 2022.

[174] K. Fatras, Y. Zine, R. Flamary, R. Gribonval, and N. Courty, "Learning with minibatch wasserstein: asymptotic and gradient properties," in *International Conference on Artificial Intelligence and Statistics*, pp. 2131–2141, PMLR, 2020.

[175] K. Fatras, T. Séjourné, R. Flamary, and N. Courty, "Unbalanced minibatch optimal transport; applications to domain adaptation," in *International Conference on Machine Learning*, pp. 3186–3197, PMLR, 2021.

[176] G. Peyré, L. Chizat, F.-X. Vialard, and J. Solomon, "Quantum entropic regularization of matrix-valued optimal transport," *European Journal of Applied Mathematics*, vol. 30, no. 6, pp. 1079–1102, 2019.

[177] T. Lehmann, M.-K. Von Renesse, A. Sambale, and A. Uschmajew, "A note on overrelaxation in the sinkhorn algorithm," *Optimization Letters*, pp. 1–12, 2021.

[178] B. Tang, N. C. Henderson, and R. Varadhan, "Accelerating fixed-point algorithms in statistics and data science: A state-of-art review," *Journal of Data Science*, pp. 1–26, 2022.

[179] D. An, N. Lei, X. Xu, and X. Gu, "Efficient optimal transport algorithm by accelerated gradient descent," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, pp. 10119–10128, 2022.

[180] P. Dvurechensky, A. Gasnikov, and A. Kroshnin, "Computational optimal transport: Complexity by accelerated gradient descent is better than by sinkhorn's algorithm," in *International conference on machine learning*, pp. 1367–1376, PMLR, 2018.

[181] M. Cuturi and G. Peyré, "A smoothed dual approach for variational wasserstein problems," *SIAM Journal on Imaging Sciences*, vol. 9, no. 1, pp. 320–343, 2016.

[182] A. Korotin, L. Li, A. Genevay, J. M. Solomon, A. Filippov, and E. Burnaev, "Do neural optimal transport solvers work? a continuous wasserstein-2 benchmark," *Advances in Neural Information Processing Systems*, vol. 34, pp. 14593–14605, 2021.

[183] L. Li, A. Genevay, M. Yurochkin, and J. M. Solomon, "Continuous regularized wasserstein barycenters," *Advances in Neural Information Processing Systems*, vol. 33, pp. 17755–17765, 2020.

[184] V. Seguy, B. B. Damodaran, R. Flamary, N. Courty, A. Rolet, and M. Blondel, "Large-scale optimal transport and mapping estimation," in *ICLR 2018-International Conference on Learning Representations*, pp. 1–15, 2018.

[185] A. Makkuva, A. Taghvaei, S. Oh, and J. Lee, "Optimal transport mapping via input convex neural networks," in *International Conference on Machine Learning*, pp. 6672–6681, PMLR, 2020.

[186] L. Béthune, A. González-Sanz, F. Mamalet, and M. Serrurier, "The many faces of 1-lipschitz neural networks," *arXiv e-prints*, pp. arXiv–2104, 2021.

[187] L. Meunier, B. J. Delattre, A. Araujo, and A. Allauzen, "A dynamical system perspective for lipschitz neural networks," in *International Conference on Machine Learning*, pp. 15484–15500, PMLR, 2022.

[188] B. Amos, L. Xu, and J. Z. Kolter, "Input convex neural networks," in *International Conference on Machine Learning*, pp. 146–155, PMLR, 2017.

[189] A. Vacher and F.-X. Vialard, "Convex transport potential selection with semi-dual criterion," *arXiv preprint arXiv:2112.07275*, 2021.

[190] A. Korotin, V. Egiazarian, A. Asadulaev, A. Safin, and E. Burnaev, "Wasserstein-2 generative networks," in *International Conference on Learning Representations*, 2020.

[191] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.

[192] S. G. Krantz and H. R. Parks, *The Implicit Function Theorem: History, Theory, and Applications*. Springer Science & Business Media, 2012.

[193] S. Gould, R. Hartley, and D. Campbell, "Deep declarative networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 8, pp. 3988–4004, 2021.

[194] S. Bai, J. Z. Kolter, and V. Koltun, "Deep equilibrium models," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[195] M. Eisenberger, A. Toker, L. Leal-Taixé, F. Bernard, and D. Cremers, "A unified framework for implicit sinkhorn differentiation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 509–518, 2022.

[196] D. Campbell, L. Liu, and S. Gould, "Solving the blind perspective-n-point problem end-to-end with robust differentiable geometric optimization," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pp. 244–261, Springer, 2020.

[197] C. Zhang, Y. Cai, G. Lin, and C. Shen, "Deepemd: Few-shot image classification with differentiable earth mover's distance and structured classifiers," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12203–12213, 2020.

[198] S. Gould, D. Campbell, I. Ben-Shabat, C. H. Koneputugodage, and Z. Xu, "Exploiting problem structure in deep declarative networks: Two case studies," *AAAI Workshop on Optimal Transport and Structured Data Modeling*, 2022.

[199] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert, "Efficient and modular implicit differentiation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 5230–5242, 2022.

[200] J. Stanczuk, C. Etmann, L. M. Kreusser, and C.-B. Schönlieb, "Wasserstein gans work because they fail (to approximate the wasserstein distance)," *arXiv preprint arXiv:2103.01678*, 2021.

[201] Y. Chen, Y. Shi, and B. Zhang, "Optimal control via neural networks: A convex approach," in *International Conference on Learning Representations*, 2019.

[202] J. Richter-Powell, J. Lorraine, and B. Amos, "Input convex gradient networks," *arXiv preprint arXiv:2111.12187*, 2021.

[203] A. Genevay, M. Cuturi, G. Peyré, and F. Bach, "Stochastic optimization for large-scale optimal transport," *Advances in neural information processing systems*, vol. 29, 2016.

[204] K. Atasu, T. Parnell, C. Dünner, M. Sifalakis, H. Pozidis, V. Vasileiadis, M. Vlachos, C. Berrospi, and A. Labbi, "Linear-

complexity relaxed word mover's distance with gpu acceleration," in *2017 IEEE International Conference on Big Data (Big Data)*, pp. 889–896, IEEE, 2017.

[205] M. Scetbon, G. Peyré, and M. Cuturi, "Linear-time gromov wasserstein distances using low rank couplings and costs," in *International Conference on Machine Learning*, pp. 19347–19365, PMLR, 2022.

[206] N. Nakagawa, R. Togo, T. Ogawa, and M. Haseyama, "Gromov-wasserstein autoencoders," *arXiv preprint arXiv:2209.07007*, 2022.

[207] J. Wu, T. Zhang, Z. Zhang, F. Wu, and Y. Zhang, "Motion-modulated temporal fragment alignment network for few-shot action recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9151–9160, 2022.

[208] N. Shlezinger, Y. C. Eldar, and S. P. Boyd, "Model-based deep learning: On the intersection of deep learning and optimization," *IEEE Access*, vol. 10, pp. 115384–115398, 2022.

[209] T. Chen, X. Chen, W. Chen, Z. Wang, H. Heaton, J. Liu, and W. Yin, "Learning to optimize: A primer and a benchmark," *The Journal of Machine Learning Research*, vol. 23, no. 1, pp. 8562–8620, 2022.

[210] B. Amos, "Tutorial on amortized optimization for learning to optimize over continuous domains," *arXiv preprint arXiv:2202.00665*, 2022.

[211] B. Amos, S. Cohen, G. Luise, and I. Redko, "Meta optimal transport," *arXiv preprint arXiv:2206.05262*, 2022.

[212] J. Lacombe, J. Digne, N. Courty, and N. Bonneel, "Learning to generate wasserstein barycenters," *Journal of Mathematical Imaging and Vision*, vol. 65, no. 2, pp. 354–370, 2023.

[213] Y. Dong, Y. Gao, R. Peng, I. Razenshteyn, and S. Sawlani, "A study of performance of optimal transport," *arXiv preprint arXiv:2005.01182*, 2020.

[214] J. Schrieber, D. Schuhmacher, and C. Gottschlich, "Dotmark–a benchmark for discrete optimal transport," *IEEE Access*, vol. 5, pp. 271–282, 2016.

[215] A. Korotin, A. Kolesov, and E. Burnaev, "Kantorovich strikes back! wasserstein gans are not optimal transport?," in *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.

[216] D. Alvarez-Melis, Y. Mroueh, and T. Jaakkola, "Unsupervised hierarchy matching with optimal transport over hyperbolic spaces," in *International Conference on Artificial Intelligence and Statistics*, pp. 1606–1617, PMLR, 2020.

[217] S. P. Singh and M. Jaggi, "Model fusion via optimal transport," *Advances in Neural Information Processing Systems*, vol. 33, pp. 22045–22055, 2020.

[218] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE signal processing magazine*, vol. 37, no. 3, pp. 50–60, 2020.

[219] A. Rakotomamonjy, M. Vono, H. J. M. Ruiz, and L. Ralaivola, "Personalised federated learning on heterogeneous feature spaces," *arXiv preprint arXiv:2301.11447*, 2023.

[220] M. Cuturi, L. Meng-Papaxanthos, Y. Tian, C. Bunne, G. Davis, and O. Teboul, "Optimal transport tools (ott): A jax toolbox for all things wasserstein," *arXiv preprint arXiv:2201.12324*, 2022.

[221] J. Feydy, A. Glaunès, B. Charlier, and M. Bronstein, "Fast geometric learning with symbolic matrices," *Advances in Neural Information Processing Systems*, vol. 33, pp. 14448–14462, 2020.

[222] M. Cuturi and G. Peyré, "Semidual regularized optimal transport," *SIAM Review*, vol. 60, no. 4, pp. 941–965, 2018.

[223] L. G. Hanin, "Kantorovich-rubinstein norm and its application in the theory of lipschitz spaces," *Proceedings of the American Mathematical Society*, vol. 115, no. 2, pp. 345–352, 1992.

[224] K. Saenko, B. Kulis, M. Fritz, and T. Darrell, "Adapting visual category models to new domains," in *Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part IV 11*, pp. 213–226, Springer, 2010.

[225] J. Wang, C. Lan, C. Liu, Y. Ouyang, T. Qin, W. Lu, Y. Chen, W. Zeng, and P. Yu, "Generalizing to unseen domains: A survey on domain generalization," *IEEE Transactions on Knowledge and Data Engineering*, pp. 1–1, 2022.

[226] J. Yang, K. Zhou, Y. Li, and Z. Liu, "Generalized out-of-distribution detection: A survey," *arXiv preprint arXiv:2110.11334*, 2021.

[227] I. Redko, E. Morvant, A. Habrard, M. Sebban, and Y. Bennani, "A survey on domain adaptation theory: learning bounds and theoretical guarantees," *arXiv preprint arXiv:2004.11829*, 2020.

[228] I. Redko, A. Habrard, and M. Sebban, "Theoretical analysis of domain adaptation with optimal transport," in *Proceedings of European Conference on Machine Learning and Knowledge Discovery in Databases: ECML PKDD 2017*, pp. 737–753, Springer, 2017.

[229] M. Rostami, "Lifelong domain adaptation via consolidated internal distribution," *Advances in neural information processing systems*, vol. 34, pp. 11172–11183, 2021.

[230] V. Titouan, I. Redko, R. Flamary, and N. Courty, "Co-optimal transport," *Advances in neural information processing systems*, vol. 33, pp. 17559–17570, 2020.

[231] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, "A comprehensive survey on transfer learning," *Proceedings of the IEEE*, vol. 109, no. 1, pp. 43–76, 2020.

[232] D. Bolya, R. Mittapalli, and J. Hoffman, "Scalable diverse model selection for accessible transfer learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 19301–19312, 2021.

[233] Y. Tan, Y. Li, and S.-L. Huang, "Otce: A transferability metric for cross-domain cross-task representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15779–15788, 2021.

[234] A. T. Tran, C. V. Nguyen, and T. Hassner, "Transferability and hardness of supervised classification tasks," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1395–1405, 2019.

[235] Y. Bao, Y. Li, S.-L. Huang, L. Zhang, L. Zheng, A. Zamir, and L. Guibas, "An information-theoretic approach to transferability in task transfer learning," in *2019 IEEE international conference on image processing (ICIP)*, pp. 2309–2313, IEEE, 2019.

[236] C. Nguyen, T. Hassner, M. Seeger, and C. Archambeau, "Leep: A new measure to evaluate transferability of learned representations," in *International Conference on Machine Learning*, pp. 7294–7305, PMLR, 2020.

[237] A. R. Zamir, A. Sax, W. Shen, L. J. Guibas, J. Malik, and S. Savarese, "Taskonomy: Disentangling task transfer learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3712–3722, 2018.

[238] K. Saito, K. Watanabe, Y. Ushiku, and T. Harada, "Maximum classifier discrepancy for unsupervised domain adaptation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3723–3732, 2018.

[239] J. Shen, Y. Qu, W. Zhang, and Y. Yu, "Wasserstein distance guided representation learning for domain adaptation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 2018.

[240] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The journal of machine learning research*, vol. 17, no. 1, pp. 2096–2030, 2016.

[241] I. Redko, N. Courty, R. Flamary, and D. Tuia, "Optimal transport for multi-source domain adaptation under target shift," in *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 849–858, PMLR, 2019.

[242] Y. Mansour, M. Mohri, and A. Rostamizadeh, "Domain adaptation: Learning bounds and algorithms," in *22nd Conference on Learning Theory, COLT 2009*, 2009.

[243] S. Shankar, V. Piratla, S. Chakrabarti, S. Chaudhuri, P. Jyothi, and S. Sarawagi, "Generalizing across domains via cross-gradient training," in *International Conference on Learning Representations*, 2018.

[244] K. Zhou, Y. Yang, T. Hospedales, and T. Xiang, "Learning to generate novel domains for domain generalization," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pp. 561–578, Springer, 2020.

[245] P. Panareda Busto and J. Gall, "Open set domain adaptation," in *Proceedings of the IEEE international conference on computer vision*, pp. 754–763, 2017.

[246] R. Xu, P. Liu, Y. Zhang, F. Cai, J. Wang, S. Liang, H. Ying, and J. Yin, "Joint partial optimal transport for open set domain adaptation.," in *IJCAI*, pp. 2540–2546, 2020.

[247] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.

[248] B. B. Damodaran, B. Kellenberger, R. Flamary, D. Tuia, and N. 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.

[249] W. Wang, H. Xu, G. Wang, W. Wang, and L. Carin, "Zero-shot recognition via optimal transport," in *Proceedings of the IEEE/CVF*

Winter Conference on Applications of Computer Vision, pp. 3471–3481, 2021.

[250] M. De Lange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, and T. Tuytelaars, "A continual learning survey: Defying forgetting in classification tasks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 7, pp. 3366–3385, 2021.

[251] S. Yuan, K. Bai, L. Chen, Y. Zhang, C. Tao, C. Li, G. Wang, R. Henao, L. Carin, and A. A. Amazon, "Weakly supervised cross-domain alignment with optimal transport," in *BMVC*, 2020.

[252] Q. Li, Z. Wen, Z. Wu, S. Hu, N. Wang, Y. Li, X. Liu, and B. He, "A survey on federated learning systems: vision, hype and reality for data privacy and protection," *IEEE Transactions on Knowledge and Data Engineering*, 2021.

[253] H. Yan, Y. Guo, and C. Yang, "Augmented self-labeling for source-free unsupervised domain adaptation," in *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021.

[254] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232, 2017.

[255] S. Kulinski and D. I. Inouye, "Towards explaining distribution shifts," *arXiv preprint arXiv:2210.10275*, 2022.

[256] S. Kulinski and D. I. Inouye, "Towards explaining image-based distribution shifts," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4788–4792, 2022.

[257] N. Courty, R. Flamary, A. Habrard, and A. Rakotomamonjy, "Joint distribution optimal transportation for domain adaptation," *Advances in neural information processing systems*, vol. 30, 2017.

[258] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.

[259] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in Neural Information Processing Systems*, vol. 27, 2014.

[260] O. Bousquet, S. Gelly, I. Tolstikhin, C.-J. Simon-Gabriel, and B. Schölkopf, "From optimal transport to generative modeling: the vegan cookbook," *stat*, vol. 1050, p. 22, 2017.

[261] H. Petzka, A. Fischer, and D. Lukovnikov, "On the regularization of wasserstein gans," in *International Conference on Learning Representations*, 2018.

[262] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," *Advances in neural information processing systems*, vol. 30, 2017.

[263] M. Xu, "Towards generalized implementation of wasserstein distance in gans," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 10514–10522, 2021.

[264] A. Genevay, G. Peyré, and M. Cuturi, "Learning generative models with sinkhorn divergences," in *International Conference on Artificial Intelligence and Statistics*, pp. 1608–1617, PMLR, 2018.

[265] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[266] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *International Conference on Machine Learning*, pp. 2256–2265, PMLR, 2015.

[267] Y. Song, C. Durkan, I. Murray, and S. Ermon, "Maximum likelihood training of score-based diffusion models," *Advances in Neural Information Processing Systems*, vol. 34, pp. 1415–1428, 2021.

[268] H. Wu, J. Köhler, and F. Noé, "Score-based Generative Modeling Secretly Minimizes the Wasserstein Distance," *Advances in Neural Information Processing Systems*, vol. 2020-December, 10 2022.

[269] G. Dantzig and M. Thapa, *Linear Programming 1*. New York: Springer-Verlag, 1997.

[270] G. Dantzig and M. Thapa, *Linear Programming 2*. New York: Springer-Verlag, 2003.

[271] H. Bradley and Magnanti, *Applied Mathematical Programming*. 1977.

[272] S. J. Wright, *Primal-Dual Interior-Point Methods*. Society for Industrial and Applied Mathematics, 1 1997.

[273] D. Kelly, B. Comm, G. O'Neill, F. Gaines, and D. O'Connor, "The Minimum Cost Flow Problem and The Network Simplex Solution Method," 1991.

[274] J. B. Orlin, "A polynomial time primal network simplex algorithm for minimum cost flows," *Mathematical Programming, Series B*, vol. 78, pp. 109–129, 8 1997.

[275] L. Chen, R. Kyng, Y. P. Liu, R. Peng, M. P. Gutenberg, and S. Sachdeva, "Maximum flow and minimum-cost flow in almost-linear time," in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 612–623, IEEE, 2022.

Abdelwahed Khamis, PhD Dr. Abdelwahed Khamis is a research scientist at the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Data61 in Australia. Abdelwahed completed his Ph.D. in Computer Science and Engineering from UNSW, Sydney in 2020. His current research interests include AI on the edge, RF and device-free sensing, and optimal transport.
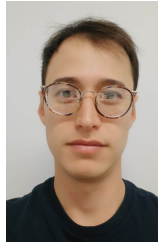
Russell Tsuchida, PhD Dr. Russell Tsuchida is a research scientist at the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Data61. Russell received his PhD from The University of Queensland (2020). His research interests include kernel methods, Gaussian processes, probabilistic machine learning and neural networks.

Mohamed Tarek, PhD Dr. Mohamed Tarek is a senior product engineer at Pumas-AI Inc. and a research affiliate at the University of Sydney Business School. Mohamed received his BSc in Mechanical Engineering from Alfaisal University (2015), MSc in Industrial and Systems Engineering from Korea Advanced Institute of Science and Technology (2017) and PhD in Computer Science from University of New South Wales (2022). His current research interests include: statistical learning of hierarchical models, Bayesian inference and applied mathematical optimization.

Vivien Rolland, PhD Dr. Vivien Rolland received a BSc degree and a MSc degree from the University Claude Bernard Lyon 1, France in 2005 and 2007, respectively, and a PhD degree from the University of Vienna, Austria, in 2012. He is a Senior Research Scientist and Team Leader at the Commonwealth Scientific and Industrial Research Organisation, Canberra, Australia. He is interested in how machine learning and computer vision can accelerate the rate of innovation in the Agrifood sector.

Lars Petersson, PhD Dr. Lars Petersson is a Senior Principal Research Scientist leading the Imaging and Computer Vision Group in Data61, CSIRO, Australia. His research interests are spanning machine learning and computer vision across a number of domains, from the smallest of microscopy scales to the largest of astronomical scales. Previously, he was a Principal Researcher in NICTA's computer vision research group leading projects such as Smart Cars, AutoMap, and Distributed Large Scale Vision. He received his PhD in March 2002 from KTH, Stockholm, Sweden, where he also received his Master's degree in Engineering Physics.