
Towards Optimal Transport with Global Invariances

David Alvarez-Melis
CSAIL, MIT

Stefanie Jegelka
CSAIL, MIT

Tommi S. Jaakkola
CSAIL, MIT

Abstract

Many problems in machine learning involve calculating correspondences between sets of objects, such as point clouds or images. Discrete optimal transport provides a natural and successful approach to such tasks whenever the two sets of objects can be represented in the same space, or at least distances between them can be directly evaluated. Unfortunately neither requirement is likely to hold when object representations are learned from data. Indeed, automatically derived representations such as word embeddings are typically fixed only up to some global transformations, for example, reflection or rotation. As a result, pairwise distances across two such instances are ill-defined without specifying their relative transformation. In this work, we propose a general framework for optimal transport in the presence of latent global transformations. We cast the problem as a joint optimization over transport couplings and transformations chosen from a flexible class of invariances, propose algorithms to solve it, and show promising results in various tasks, including a popular unsupervised word translation benchmark.

1 Introduction

Optimal transport (OT) plays dual roles across machine learning applications. First, it provides a well-founded, geometrically driven approach to realizing correspondences between sets of objects such as shapes in different images. Such correspondences can be used for image registration (Haker and Tannenbaum, 2001) or to interpolate between them (Solomon et al., 2015).

More generally, OT extends to problems such as domain adaptation where we wish to transport a set of labeled source points to the realm of the target task (Courty et al., 2017a,b). Second, the transportation problem induces a theoretically well-characterized distance between distributions. This distance is expressed in the form of a transport cost and serves as a natural population difference measure, which can be exploited as a source of feedback in adversarial training (Arjovsky et al., 2017; Bousquet et al., 2017). Our focus in this paper is on the optimal coupling mediating the transport, i.e., realizing the latent correspondences between the objects.

A key limitation of classic OT is that it implicitly assumes that the two sets of objects in question are represented in the same space, or at least that meaningful pairwise distances between them can be computed. This is not always the case, especially when the objects are represented by learned feature vectors. For example, word embedding algorithms operate at the level of inner products or distances between word vectors, so the representations they produce can be arbitrarily rotated, sometimes even for different runs of the same algorithm on the same data. Such global degrees of freedom in the vector representations render direct pairwise distances between objects across the sets meaningless. Indeed, OT is *locally greedy* as it focuses on minimizing individual movement of mass, oblivious to global transformations. As a concrete example, consider two identical sets of points where one set is subjected to a global rotation. The optimal transport coupling evaluated between the resulting sets may no longer recover the correct correspondences (Figure 3).

When the global transformation is known or can be easily estimated, it can be incorporated in the computation of pairwise distances, thereby enabling the use of traditional OT. Unfortunately, only the *type* of underlying transformation (e.g., rotation) is typically known, not the actual realization. In such cases, we would like the optimal transport problem to also find the best latent transformation along with the optimal coupling. In other words, we seek a formulation of OT that remains invariant under global transformations.

In this work, we propose a generalization of discrete optimal transport that incorporates global invariances directly into the optimization problem. While the driving motivation is invariance to rigid transformations (arguably the most common case in practice), our framework is very general and allows various other types of invariances to be encoded. Moreover, our approach unifies previous methods for fusing OT with global transformations such as Procrustes mappings (Grave et al., 2018; Rangarajan et al., 1997; Zhang et al., 2017b), and reveals unexpected connections to the Gromov-Wasserstein distance (Mémoli, 2011).

The main contributions of this work are thus:

- A novel formulation of the discrete optimal transport problem that allows for global geometric invariances to be incorporated into the objective;
- Design and analysis of efficient algorithms for this general class of problems;
- An application of the framework to the problem of unsupervised word translation, yielding performance comparable to state-of-the-art methods at a fraction of their computational cost.

2 Related Work

The general problem of unsupervised estimation of correspondence between sets of features is well-studied and arises in various fields under different names, such as manifold alignment (Wang and Mahadevan, 2009), feature set matching (Grauman and Darrell, 2005) and feature correspondence finding (Torresani et al., 2008). Here, we focus the discussion on related methods that combine soft correspondences (such as those produced by OT) with explicit space alignment.

Perhaps the earliest such approach is by Rangarajan et al. (1997), who derive a framework to establish correspondences between shapes that rejects non-homologies (e.g., rotations) based on an entropy-regularized version of the OT problem. The resulting *Softassign Procrustes Algorithm* proceeds iteratively by alternating between estimating optimal rotations and performing Sinkhorn iterations. This approach, however, only considers rotations, and is tailored to 2D data, where rotations can be easily parametrized.

More recently, Zhang et al. (2017b) propose combining OT with Procrustes alignment to find correspondences between word embedding spaces. They initialize their orthogonal mapping using an adversarial training phase, much like Conneau et al. (2018), and solve the optimization problem with alternating minimization. Our approach, on the other hand, does not rely on neural network initialization, instead leveraging a convexity annealing scheme that leads to smooth

convergence, with little sensitivity to initialization.

Concurrently with our work, Grave et al. (2018) tackle the word embedding alignment task with an approach similar to that of Zhang et al. (2017b), combining Wasserstein distances (an instance of OT) and Procrustes alignment. Their approach differs from Zhang et al. (2017b) in how they scale up optimization, by relying on a stochastic Sinkhorn solver (Genevay et al., 2018), and in how they initialize it, by solving a convex relaxation of the original problem.¹

Although driven by a similar motivation (word embedding alignment) and relying on similar principles (joint optimization of transport coupling and feature mapping) as the work of Zhang et al. (2017b) and Grave et al. (2018), our approach differs from them in several aspects. First, it allows for more general types of invariance classes (characterized as Schatten ℓ_p -norm balls), subsuming orthogonal invariance considered in prior work as a special case. Second, we dispense with the need for any ad-hoc initialization by introducing instead a convexity-annealing approach to optimization. Third, our approach remains robust to the choice of entropy regularization parameter λ .

Compared to methods that directly solve a Procrustes problem from a few known correspondences (Artetxe et al., 2017; Zhang et al., 2016) or by generating pseudo-matches through an initial unsupervised step (Artetxe et al., 2018a; Conneau et al., 2018), OT allows for more flexible correspondence estimation.

A different generalization of the OT problem aimed at overcoming lack of intrinsic correspondence between spaces is the Gromov-Wasserstein (GW) distance (Mémoli, 2011). It has been recently applied to various correspondence problems, including shape interpolation (Peyré et al., 2016; Solomon et al., 2016) and unsupervised word translation (Alvarez-Melis and Jaakkola, 2018). While our framework recovers this distance in certain scenarios (see §4.2), it is best understood as a compromise between the classic formulation of OT that requires the spaces to be fully registered, and the GW distance, which completely forgoes explicit computation of distances across spaces, relying instead on comparison of intra-space similarities. Thus, our approach is best suited to tasks where distances across spaces can be computed, but are meaningful only if made invariant to some latent transformation. A further difference with the usual OT and GW distances is that our approach produces, as intrinsic part of optimization, a global mapping which can be used to map *out-of-sample points* across spaces.

¹Interestingly, this relaxation corresponds to a hybrid of two instances of our framework: optimizing a Frobenius-norm objective (§4.2) over orthogonal matrices (§4.1).

3 Feature Space Correspondences

Notation. We denote vectors and matrices with bold font (e.g., \mathbf{x} , \mathbf{X}), their entries without it (x_i , X_{ij}), and sets as X , Y . We use super-indices to enumerate vectors, and subindices to denote their entries. For matrices \mathbf{A} , \mathbf{B} , let $\langle \mathbf{A}, \mathbf{B} \rangle = \sum_{i,j} [\mathbf{A}]_{ij} [\mathbf{B}]_{ij}$ denote their Frobenius inner product, and $\|\cdot\|_*$ the nuclear (trace) norm. We denote by \mathcal{O}_n the set of orthogonal matrices of order n , and finally, let $\llbracket n \rrbracket \triangleq \{1, \dots, n\}$.

3.1 Problem Setting

Suppose we are given two sets of examples, $X = \{\mathbf{x}^{(i)}\}_{i=1}^n$ and $Y = \{\mathbf{y}^{(j)}\}_{j=1}^m$, drawn from potentially distinct feature spaces $\mathcal{X} \subset \mathbb{R}^{d_x}$ and $\mathcal{Y} \subset \mathbb{R}^{d_y}$. Our goal is to learn correspondences between X and Y , in the challenging case where no prior instance-wise correspondences are known—i.e., the problem is fully unsupervised—and the spaces \mathcal{X} and \mathcal{Y} are unregistered—i.e., the *global* correspondence between them is unknown too.

This task can be thought of as consisting of two sub-problems: (i) finding a global alignment of spaces \mathcal{X} and \mathcal{Y} , e.g., via a mapping $T : \mathcal{X} \rightarrow \mathcal{Y}$ such that $\|T(\mathbf{x}) - \mathbf{y}\|$ is small for every correspondence pair (\mathbf{x}, \mathbf{y}) , and (ii) finding correspondences between the items in X and Y , via an assignment $\mathcal{A} : \llbracket n \rrbracket \mapsto \llbracket m \rrbracket$ such that $\mathbf{x}^{(i)} \rightarrow \mathbf{y}^{(j)}$ if these points are in correspondence. Individually, these two problems are well-studied and understood. Below, we briefly discuss popular approaches to solve them. In Section 4, we show how to combine them to enforce invariances in the optimal transport problem, leading to a flexible class of problems that can be solved efficiently.

3.2 Space alignment from paired samples

Assume for now² that $m = n$ and $d_x = d_y$. Given paired samples from the two domains, consider matrices $\mathbf{X} \in \mathbb{R}^{d \times n}$ and $\mathbf{Y} \in \mathbb{R}^{d \times n}$, whose columns correspond to these paired elements. The problem of finding the best mapping T that maps the target samples to the source ones can be cast as

$$\min_{T \in \mathcal{F}} \|\mathbf{X} - T(\mathbf{Y})\|^2$$

where \mathcal{F} is some class of functions and $\|\cdot\|$ is typically taken to be the Frobenius norm $\|\mathbf{A}\|_F = \sqrt{\sum_{i,j} |a_{ij}|^2}$. Naturally, the choice of space \mathcal{F} will determine the difficulty of finding T as well as the quality of the alignment implied by it.

²We discuss how to relax these in Appendix.

The classic *Orthogonal Procrustes problem* restricts \mathcal{F} to be rigid (rotation and reflection) transformations—i.e., orthogonal matrices:

$$\min_{\mathbf{P} \in \mathcal{O}_n} \|\mathbf{X} - \mathbf{P}\mathbf{Y}\|_F^2 \quad (1)$$

Despite its simplicity, Procrustes analysis is a powerful tool used in various applications, from statistical shape analysis (Goodall, 1991) to market research and others (Gower and Dijkstra, 2004). Its main advantage is that it has a closed-form solution in terms of a singular value decomposition (SVD) (Schönmann, 1966). Namely, given an SVD, say $\mathbf{U}\Sigma\mathbf{V}^\top$, of $\mathbf{X}\mathbf{Y}^\top$, the orthogonal matrix minimizing problem (1) is $\mathbf{P}^* = \mathbf{U}\mathbf{V}^\top$, which is a direct consequence of a well-known approximation property of the SVD:

Lemma 3.1. *If $\mathbf{A} \in \mathbb{R}^{n \times m}$ and $\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^\top$ is an SVD of \mathbf{A} , then $\operatorname{argmax}_{\mathbf{P} \in \mathcal{O}_n} \langle \mathbf{P}, \mathbf{A} \rangle = \mathbf{U}\mathbf{V}^\top$.*

Proof. This is a particular case of Lemma 4.2. All other proofs are provided in the Appendix. \square

We emphasize that the Procrustes problem (1) crucially requires the columns of \mathbf{X} and \mathbf{Y} be paired, making it an intrinsically *supervised* approach. Thus, its application to the problem of feature alignment requires either an—ideally small—initial set of true paired examples or a method to generate them.

3.3 Correspondences between aligned spaces

Optimal transport is an appealing approach for finding correspondences between feature vectors in a fully unsupervised way. Besides strong theoretical foundations and fast algorithms, it is ideal for this setting, as it is fully unsupervised, inferring correspondences by leveraging the geometry of these collections.

The OT problem considers two measures μ and ν over spaces \mathcal{X} and \mathcal{Y} , and a transport cost $c : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^+$. It seeks to minimize the cost of *transporting* measure μ onto ν . In its original discrete formulation, μ and ν are taken as empirical distributions:

$$\mu = \sum_{i=1}^n p_i \delta_{\mathbf{x}^{(i)}}, \quad \nu = \sum_{j=1}^m q_j \delta_{\mathbf{y}^{(j)}} \quad (2)$$

so all relevant pairwise costs can be represented as a matrix $\mathbf{C} \in \mathbb{R}^{n \times m}$. Monge’s formulation seeks a transport map $T : \mathcal{X} \rightarrow \mathcal{Y}$ that assigns each point $\mathbf{x}^{(i)}$ to a single point $\mathbf{y}^{(j)}$, and which pushes the mass of μ toward that of ν (i.e., $T_{\#}\mu = \nu$), while minimizing:

$$\min_T \left\{ \sum_{i=1}^n c(\mathbf{x}^{(i)}, T(\mathbf{x}^{(i)})) : T_{\#}\mu = \nu \right\} \quad (3)$$

Since the solution to this problem is not guaranteed to exist, a relaxation of this problem by Kantorovich considers instead “soft” assignments defined in terms of probabilistic *couplings* $\Gamma \in \mathbb{R}_+^{n \times m}$ whose marginals recover μ and ν . Formally, Kantorovich’s formulation seeks Γ in the transportation polytope:

$$\Pi(\mathbf{p}, \mathbf{q}) = \{\Gamma \in \mathbb{R}_+^{n \times m} \mid \Gamma \mathbf{1} = \mathbf{p}, \Gamma^\top \mathbf{1} = \mathbf{q}\} \quad (4)$$

that solves

$$\min_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \langle \Gamma, \mathbf{C} \rangle. \quad (5)$$

The discrete optimal transport (DOT) problem is a linear program. Practical algorithms to solve it include Orlin’s algorithm and interior point methods, both of which have $O(n^3 \log n)$ complexity (Pele and Werman, 2009). However, adding an entropy regularization term leads to much more efficient algorithms:

$$\min_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \langle \Gamma, \mathbf{C} \rangle - \frac{1}{\lambda} H(\Gamma). \quad (6)$$

This is a strictly convex optimization problem, whose solution has the form $\Gamma^* = \text{diag}(\mathbf{u}) \mathbf{K} \text{diag}(\mathbf{v})$, with $\mathbf{K} = e^{-\frac{\mathbf{C}}{\lambda}}$ where the exponential is computed entry-wise (Peyré and Cuturi, 2018), and can be obtained efficiently via the Sinkhorn-Knopp algorithm, an iterative matrix-scaling procedure (Cuturi, 2013). Besides computational motivation, introducing this regularization often leads to better empirical performance in applications where having denser correspondences is beneficial, e.g., in settings where the support points correspond to noisy features (Alvarez-Melis et al., 2018).

Unfortunately, naive application of DOT often fails in applications where the two spaces \mathcal{X} and \mathcal{Y} are not *registered*: i.e., when there is no a priori meaningful notion of distance between them. As stated before, this is likely to be the case when the feature spaces are learned. In such cases, even if the feature spaces are of the same dimensionality, the naive choice of cost function $c(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|$ is flawed, as there is no guarantee that these two spaces are coherent, i.e., that their coordinate axes are in correspondence. To remedy this situation, we seek to make optimal transport invariant to global transformations of the space. In the next section, we do so for a general class of invariances.

4 Transporting with Global Geometric Invariances

Our goal in this work is to extend the DOT problem (5) to enforce invariance with respect to certain classes of transformations. Formally, we assume there exists an unknown function f in a pre-specified class \mathcal{F} which characterizes the global correspondence between

spaces \mathcal{X} and \mathcal{Y} , i.e., for which $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{x} = f(\mathbf{y}), \mathbf{y} \in \mathcal{Y}\}$. Naturally, the choice of \mathcal{F} should be informed by the application domain.

In this setting, given collections $\{\mathbf{x}^{(i)}\}_{i=1}^n$ and $\{\mathbf{y}^{(j)}\}_{j=1}^m$ and associated empirical measures μ, ν , we seek to simultaneously find the best global transformation of the space (within \mathcal{F}) and the best local correspondences between the two collections, as defined by the optimal transport problem. In other words, we seek to jointly optimize $f \in \mathcal{F}$ and $\Gamma \in \Pi(\mathbf{p}, \mathbf{q})$ to minimize the transportation cost between the two collections. Formally, given an invariance set \mathcal{F} , for $f \in \mathcal{F}$ let $f(\mathbf{Y})$ denote the matrix of size $d \times m$ whose columns are $f(\mathbf{y}^{(j)})$. The problem we wish to solve is

$$\min_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \min_{f \in \mathcal{F}} \langle \Gamma, C(\mathbf{X}, f(\mathbf{Y})) \rangle. \quad (7)$$

where $[C(\mathbf{X}, f(\mathbf{Y}))]_{ij} = c(\mathbf{x}^{(i)}, f(\mathbf{y}^{(j)}))$. In this work, we focus on invariances defined by linear operators with bounded norm:

$$\mathcal{F}_p \triangleq \{\mathbf{P} \in \mathbb{R}^{d \times d} \mid \|\mathbf{P}\|_p \leq k_p\}, \quad (8)$$

where $\|\cdot\|_p$ is the Schatten ℓ_p -norm, that is, $\|\mathbf{P}\|_p = \|\sigma(\mathbf{P})\|_p$ where $\sigma(\mathbf{P})$ is a vector containing the singular values of \mathbf{P} . In addition, k_p is a norm- and problem-dependent constant.³ This choice of invariance sets follows both modeling and computational motivations. As for the former, Schatten norms allow for immediate interpretation of the elements of \mathcal{F}_p in terms of their spectral properties (Fig. 1). For example, choosing $p = 1$ encourages solutions with sparse spectra (e.g., projections, useful when the support of one of the two distributions is known to be contained in a lower-dimensional subspace), while $p = \infty$ instead seeks solutions with uniform spectra (e.g., unitary matrices, thus enforcing invariance to rigid transformations). Intermediate values of p interpolate between these two extremes. Interestingly, the choice $p = 2$ recovers a recent popular generalization of the optimal transport problem motivated by a similar goal: the Gromov-Wasserstein distance (Mémoli, 2011), as we show in Section 4.2. Thus, the proposed Schatten invariance framework offers significant flexibility. In terms of computation, Schatten norms exhibit various desirable properties, such as unitary invariance, submultiplicativity, and easy characterization via duality, all of which play an important role in deriving efficient optimization algorithms below.

Here, we formulate the problem for the case where the ground metric c is the squared euclidean distance, i.e., $c(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2^2$, which is arguably the most common choice in practice. With this choice of ground

³In the most common case, k_p would be chosen to ensure the identity mapping is contained in this set.

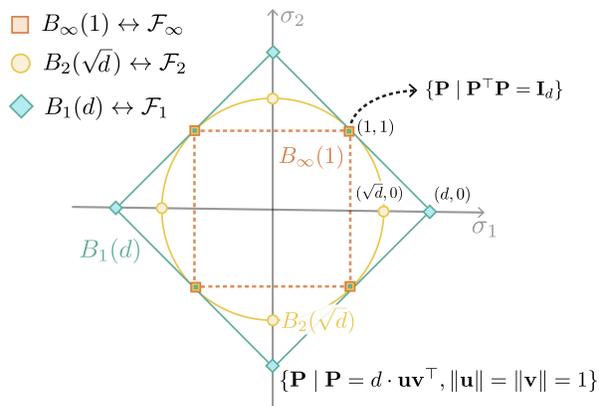


Figure 1: **Schatten-norm invariance classes.** The depicted ℓ_p -norm balls in singular value space correspond to matrix invariance classes \mathcal{F}_p . The radius is chosen so as to include the identity matrix ($\boldsymbol{\sigma} = [1, 1]$). For linear objectives, solutions when optimizing over \mathcal{F}_∞ and \mathcal{F}_1 can be found on the extreme points of their respective constraint spaces: orthogonal matrices for the former and rank-one matrices for the latter.

metric, let \mathbf{u}, \mathbf{v} be vectors with entries $u_i = \|\mathbf{x}^{(i)}\|_2^2$, and $v_j = \|\mathbf{P}\mathbf{y}^{(j)}\|_2^2$ respectively. Then, it is easy to show that (7) becomes:

$$\max_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \max_{\mathbf{P} \in \mathcal{F}} 2\langle \Gamma, \mathbf{X}^\top \mathbf{P} \mathbf{Y} \rangle - \langle \mathbf{u}, \mathbf{p} \rangle - \langle \mathbf{v}, \mathbf{q} \rangle \quad (9)$$

This objective has a clear interpretation. The first term, which can be equivalently written as $\langle \mathbf{X}\Gamma, \mathbf{P}\mathbf{Y} \rangle$, measures agreement between $\mathbf{X}\Gamma$, the source points mapped according to the barycentric mapping implied by Γ , and $\mathbf{P}\mathbf{Y}$, the target points mapped according to \mathbf{P} . The other two terms, which can be interpreted as empirical expectations $\hat{\mathbb{E}}_{\mathbf{x} \sim \mu} \|\mathbf{x}\|_2^2$ and $\hat{\mathbb{E}}_{\mathbf{y} \sim \nu} \|\mathbf{P}\mathbf{y}\|_2^2$, act as a counterbalance, normalizing the objective and preventing artificial maximization of the similarity term by arbitrary scaling of the mapped vectors.

In general, Problem (9) is not jointly concave on \mathbf{P} and Γ , but it is concave in either variable if the other one is fixed. Hence, we can solve it via alternating maximization on \mathbf{P} and Γ . Since only the first term depends on Γ , solving for this variable for a fixed \mathbf{P} is a usual OT problem, for which we discuss optimization in Section 4.3. On the other hand, for a fixed Γ the problem is a concave maximization over a compact and convex set, which can be solved efficiently with Frank-Wolfe-type algorithms since projecting onto Schatten norm balls is tractable (Jaggi, 2013). While this approach provides a tractable way to solve problem (9) in general, we show that under conditions that often hold in practice, optimization is much simpler. This simplification relies on eliminating the dependence on \mathbf{P} of the third term in problem (9):

Lemma 4.1. *Under either of the conditions*

1. $\forall \mathbf{P} \in \mathcal{F}$, \mathbf{P} is angle-preserving (i.e., $\forall \mathbf{x}, \mathbf{y} \langle \mathbf{P}\mathbf{x}, \mathbf{P}\mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{y} \rangle$), or
2. $\exists k \geq 0 : \|\mathbf{P}\|_F = k \quad \forall \mathbf{P} \in \mathcal{F}$ and the matrix \mathbf{Y} is ν -whitened (i.e., $\mathbf{Y} \text{diag}(\mathbf{q})^2 \mathbf{Y}^\top = \mathbf{I}_d$),

Problem (9) is equivalent to

$$\max_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \max_{\mathbf{P} \in \mathcal{F}} \langle \mathbf{X}\Gamma\mathbf{Y}^\top, \mathbf{P} \rangle. \quad (10)$$

Condition (1) above is reasonable as it guarantees \mathbf{P} preserves geometric relations across spaces. On the other hand, whitening is a common pre-processing step in feature learning (Hyvärinen and Oja, 2000) and correspondence problems (Artetxe et al., 2018b).

The following generalization of Lemma 3.1 shows that when optimizing over Schatten ℓ_p -norm balls, the inner problem in (10) admits a closed form solution.

Lemma 4.2. *Let \mathbf{M} be a matrix with SVD $\mathbf{M} = \mathbf{U}\Sigma\mathbf{V}^\top$ and let $\Sigma = \text{diag}(\boldsymbol{\sigma})$, then*

$$\operatorname{argmax}_{\mathbf{P} : \|\mathbf{P}\|_p \leq k} \langle \mathbf{P}, \mathbf{M} \rangle = \mathbf{U} \text{diag}(\mathbf{s}) \mathbf{V}^\top \quad (11)$$

where \mathbf{s} is such that $\|\mathbf{s}\|_p \leq k$ and attains $\mathbf{s}^\top \boldsymbol{\sigma} = k \|\boldsymbol{\sigma}\|_q$, for $\|\cdot\|_q$ the dual norm of $\|\cdot\|_p$.

Therefore, the inner problem in (10) boils down to maximization of support functions of vector-valued ℓ_p balls, which can be done in closed form for any $p \geq 1$ by choosing $s_i \propto \sigma_i^{q-1}$ (Jaggi, 2013). This, in turn, greatly simplifies the alternating optimization approach. For a fixed Γ , we can use Lemma 4.2 to obtain a closed-form solution \mathbf{P}^* . On the other hand, for a fixed \mathbf{P} , optimizing for Γ is a classic discrete optimal transport problem with cost matrix $\tilde{\mathbf{C}} = -\mathbf{X}^\top \mathbf{P} \mathbf{Y}$,⁴ which can be solved with off-the-shelf OT algorithms.

Next, we investigate what Lemma 4.2 implies for two salient cases, $p = \infty$ and $p = 2$. The case $p = 1$ is discussed in the Appendix. Then, in Section 4.3, we discuss optimization of the general problem in detail.

4.1 The case $p = \infty$

The Schatten ℓ_∞ -norm is the spectral norm $\|A\|_\infty = \sigma_{\max}(A)$. To guarantee that the identity is contained in \mathcal{F}_∞ , we take $k_\infty \triangleq 1$. Note that combining either condition in Lemma 4.1 with this implies that $\mathcal{F}_\infty = \mathcal{O}_n$. Therefore, this choice of norm naturally encodes invariance to rotations and reflections. The dual characterization of Schatten norms implies that

$$\max_{\mathbf{P} \in \mathcal{F}_\infty} \langle \mathbf{X}\Gamma\mathbf{Y}^\top, \mathbf{P} \rangle = \|\mathbf{X}\Gamma\mathbf{Y}^\top\|_* \quad (12)$$

⁴This is of course equivalent to solving the original problem (7), whose cost matrix has a simpler interpretation.

so that (10) becomes a single-block problem:

$$\max_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \|\mathbf{X}\Gamma\mathbf{Y}^\top\|_* \quad (13)$$

Albeit succinct, this alternative representation of the problem is not easier to solve. Despite having eliminated \mathbf{P} , the objective is now non-convex with respect to Γ (maximization of a convex function). Nevertheless, this formulation offers an interesting geometric interpretation. When \mathbf{p}, \mathbf{q} are uniform distributions, then $\hat{\mathbf{X}} \triangleq \mathbf{X}\Gamma$ is a matrix whose columns correspond to the those of \mathbf{X} transported to \mathcal{Y} according to the optimal barycentric mapping. Hence, $\hat{\mathbf{X}}\mathbf{Y}^\top$ is the (shifted) cross-covariance matrix of the features in \mathcal{X} and \mathcal{Y} space, i.e., $[\hat{\mathbf{X}}\mathbf{Y}^\top]_{ij} = \text{cov}(\hat{x}_i, y_j)$, and its norm indicates the strength of correlation of these features. Therefore, problem (13) essentially seeks a transport coupling that maximizes the correlation of feature dimensions after transportation. We leave exploration of direct techniques to optimize (13) for future work. Here instead we rely on the generic alternating minimization scheme described in the previous section.

4.2 The case $p = 2$

The Schatten ℓ_2 -norm is the Frobenius norm. Since $\|\mathbf{I}_d\|_F = \sqrt{d}$, we take $\mathcal{F}_2 = \left\{ \mathbf{P} \mid \|\mathbf{P}\|_F = \sqrt{d} \right\}$. As before, we use Schatten norm duality to note that

$$\max_{\mathbf{P} \in \mathcal{F}_2} \langle \mathbf{X}\Gamma\mathbf{Y}^\top, \mathbf{P} \rangle = \sqrt{d} \|\mathbf{X}\Gamma\mathbf{Y}^\top\|_F,$$

whereupon problem (10) now becomes

$$\max_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \|\mathbf{X}\Gamma\mathbf{Y}^\top\|_F, \quad (14)$$

with similar intuition to Problem (13), but for a different metric. However, this subtle difference has important consequences, such as the following connection.

Lemma 4.3. *Consider the Gromov-Wasserstein problem for discrete measures μ and ν (Peyré et al., 2016):*

$$\min_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \sum_{i,j,k,l} L(\mathbf{C}_{ik}^x, \mathbf{C}_{jl}^y) \Gamma_{ij} \Gamma_{kl}, \quad (15)$$

where $(\mathbf{C}^x, \mathbf{p})$ and $(\mathbf{C}^y, \mathbf{q})$ are (intra-space) measured similarity matrices and L is a loss function. For the choice of cosine similarity and squared loss $L(a, b) = \frac{1}{2}|a - b|^2$, Problems (15) and (14) are equivalent.

4.3 Optimization

We solve problem (10) with alternating maximization on Γ and \mathbf{P} . For a fixed Γ , Lemma 4.2 shows a closed-form solution \mathbf{P}^* at the cost of an $d \times d$ SVD, i.e., $O(d^3)$. For a fixed \mathbf{P} , the optimal Γ^* can be obtained

exactly in $O(N^3 \log N)$ time ($N = \max\{m, n\}$) via linear programming, or ϵ -approximately via the Sinkhorn algorithm in $O(N^2 \log N \epsilon^{-3})$ time (Altschuler et al., 2017). The latter would correspond to solving an entropy-regularized version of the problem:

$$\max_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \max_{\mathbf{P} \in \mathcal{F}} \langle \Gamma, \mathbf{X}^\top \mathbf{P} \mathbf{Y} \rangle + \lambda H(\Gamma). \quad (16)$$

The Sinkhorn algorithm can be applied even to the original (non-regularized) problem (10) by relying on inexact alternating minimization methods that allow for approximate solution of intermediate steps (Eckstein and Yao, 2017; Mokhtari et al., 2015). Besides providing an alternative algorithmic approach, this observation could be used to prove convergence rates for problem (10). Here, we instead focus on optimizing the regularized formulation (16).

Alternating optimization methods for non-convex objectives are known to be sensitive to initialization (Hardt, 2014; Jain et al., 2013). Indeed, a key component of fully unsupervised approaches to feature alignment is finding good quality initializations. For example, for unsupervised word embedding alignment, state-of-the-art methods rely on additional—often heuristic—steps to generate good initial solutions, such as adversarially-trained neural networks (Conneau et al., 2018; Zhang et al., 2017a,b), which themselves are often very sensitive to initialization, sometimes failing completely on the same problem for different random restarts (Artetxe et al., 2018b).

Neither Problem (10) nor (16) is jointly concave in Γ and \mathbf{P} , thus facing in principle a similar challenge in terms of sensitivity to initialization. However, in (16) the strength of the entropic regularization controls the extent of non-concavity: strong regularization leads to a more concave objective, while $\lambda \rightarrow 0$ leads to increasingly more non-concave objective. We propose to leverage this observation to alleviate sensitivity to initialization by using an annealing scheme on the regularization term. Starting from a large value of λ , we decay this value in each iteration by setting $\lambda_t = \alpha \times \lambda_{t-1}$ with $\alpha \in (0, 1)$, until a minimum value $\underline{\lambda}$ is reached. We stop the method when the value of the objective converges. The advantage of this annealing approach is that it avoids ad-hoc initialization, and eliminates the need for hyperparameter tuning on λ , since any sufficiently large choice of λ_0 achieves the same objective. In *all* our experiments, we use the same parameter values $\lambda_0 = 1$ and decay $\alpha = 0.95$.

The method described so far, summarized as Algorithm 1 in the supplement, is used in our first set of experiments. However, to scale up to very large problems, such as the motivating application of word embedding alignment, we propose an incremental approach, described in detail in Appendix F.

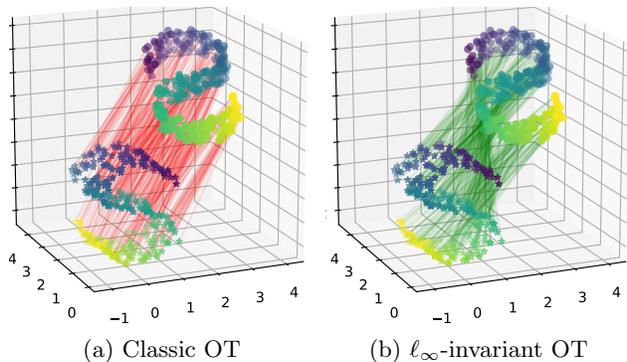


Figure 2: Optimal couplings for the synthetic point cloud dataset with underlying orthogonal (\mathcal{F}_∞) invariance; green (red) edges denote (in)correct matchings.

5 Experiments

5.1 Synthetic Datasets

First, we test our approach in a controlled setting with known underlying invariance. We generate a point cloud in \mathbb{R}^3 (the *source*), and then apply a transformation \mathbf{P} randomly sampled from one of the families \mathcal{F}_p to generate a *target* point cloud. The goal is thus to recover the true correspondences between source and target points. We generate a discrete matching ψ from a coupling Γ as $\psi(i) = \operatorname{argmax}_j \Gamma_{ij}$, and compute its accuracy with respect to the known true point-wise correspondences. As expected, when the true latent transformation is orthogonal, endowing OT with ℓ_∞ invariance allows it to recover the correct matching between the point clouds, while the classic (invariance-agnostic) formulation does not, greedily matching based on proximity instead (Figure 2).

Next, we investigate the effect of noise. As before, we generate point clouds with two types of invariances (\mathcal{F}_2 and \mathcal{F}_∞), but now add a Gaussian noise term with variance σ to the target points. We compare the performance of the following versions of OT: classic EMD (5), the entropy-regularized formulation (6) solved via the SINKHORN algorithm, our ℓ_2 and ℓ_∞ -invariant versions, and an ORACLE which solves a entropy-regularized problem *without the transformation applied*, i.e. only adding noise. Figure 3 (top) shows the matching accuracy (mean and one standard deviation over 5 repetitions). As expected, ℓ_∞ -OT is better than ℓ_2 -OT at recovering the correspondences when the true transformation is orthogonal, and vice versa; and there is a loss of accuracy caused by the estimation of \mathbf{P} in the ℓ_2 case, as shown by the gap between our methods and ORACLE OT. But surprisingly, both invariance methods outperform the oracle in the ℓ_∞ case, which we attribute to the added freedom of choosing \mathbf{P} to overcome noise, combined with

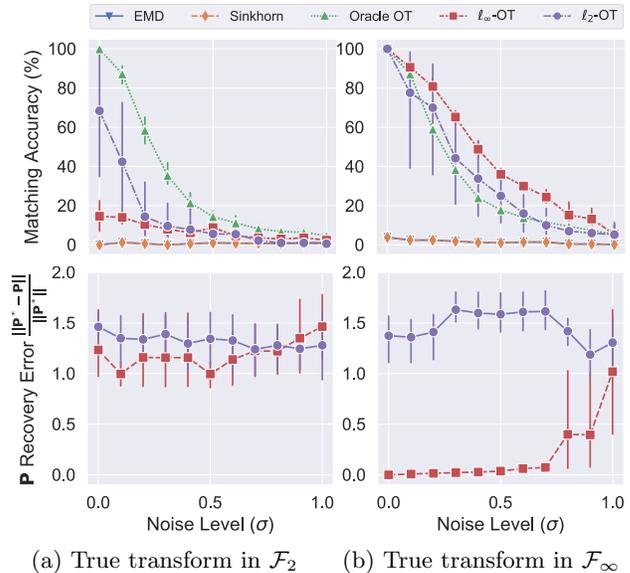


Figure 3: Results for the synthetic datasets. **Top**: matching accuracy for the computed coupling $\hat{\Gamma}$. **Bottom**: error in recovering \mathbf{P} . The plots show mean values and one s.d. error bars over 5 repetitions.

the ease of optimizing over \mathcal{F}_∞ compared to \mathcal{F}_2 . This hypothesis is supported by the overall higher error in recovering \mathbf{P} in the latter case (Fig. 3, bottom).

5.2 Unsupervised word translation

Automatic word translation has a long tradition (under the name *bilingual lexical induction*) in computational linguistics (Fung, 1995; Rapp, 1995), and has seen a recent revival due to the success of fully unsupervised methods (Artetxe et al., 2018a; Conneau et al., 2018). Most such methods cast the problem as feature alignment between sets of word embeddings, motivated by the observation that these possess similar geometry across languages (Mikolov et al., 2013). Though their relational structure is similar, the absolute position of these vectors is irrelevant. Indeed, word embedding algorithms are naturally interpreted as metric recovery methods (Hashimoto et al., 2016), making these vectors intrinsically invariant to angle (or distance) preserving transformations. This observation suggests inducing invariance to orthogonal transformations, as described in Section 4.1.

Most current unsupervised methods circumvent this issue by resorting to ad-hoc normalization, joint re-embedding, or by estimating a complex mapping between the two spaces with adversarial training. These methods require careful initialization and post-mapping refinements, such as mitigating the effect of frequent words on neighborhoods, and are often hard to tune properly (Artetxe et al., 2018a).

	Supervision	Time	EN-ES		EN-FR		EN-DE		EN-IT		EN-RU	
			→	←	→	←	→	←	→	←	→	←
PROCRUSTES	5K words	3	77.6	77.2	74.9	75.9	68.4	67.7	73.9	73.8	47.2	58.2
PROCRUSTES + CSLS	5K words	3	81.2	82.3	81.2	82.2	73.6	71.9	76.3	75.5	51.7	63.7
ADV + CSLS	None	643	75.7	79.7	77.8	71.2	70.1	66.4	72.4	71.2	37.1	48.1
ADV + CSLS + REFINE	None	957	81.7	83.3	82.3	82.1	74.0	72.2	77.4	76.1	44.0	59.1
WASSERSTEIN + CSLS	None	–	82.8	84.1	82.6	82.9	75.4	73.3	–	–	43.7	59.1
GROMOV-WASSERSTEIN	None	37	81.7	80.4	81.3	78.9	71.9	72.8	78.9	75.2	45.1	43.7
SELF-LEARN + CSLS	None	476	82.3	84.7	82.3	83.6	75.1	74.3	79.2	79.8	48.9	65.9
ℓ_∞ -INVAROT + CSLS	None	70	81.3	81.8	82.9	81.6	73.8	71.1	77.7	77.7	41.7	55.4

Table 1: Accuracy on the word translation task. “Time” refers to average runtime in minutes per problem (language pair) on the same CPU machine. We report results for our ℓ_∞ -OT method *without* relying on the iterative refinement step of Conneau et al. (2018), so it is more appropriately compared to their ADV + CSLS.

5.2.1 Dataset and Methods

The dataset of Conneau et al. (2018) consists of fast-Text word embeddings and parallel dictionaries for 110 language pairs. Here, we focus on the five language pairs for which they report results: English to Spanish, French, Italian, German and Russian. As they do, we include a strong semi-supervised baseline of solving a PROCRUSTES problem directly using the available (5K) cross-lingual embedding pairs. In addition, we compare against various unsupervised methods: SELFLEARN (Artetxe et al., 2018a), GROMOV-WASSERSTEIN (Alvarez-Melis and Jaakkola, 2018), ADV (Conneau et al., 2018) and WASSERSTEIN (Grave et al., 2018).⁵ The code for the last of these was not available, so we report results from their paper (which excludes EN-IT), and omit runtime. Whenever nearest neighbor search is required, we use Cross-domain Similarity Local Scaling (CSLS) (Conneau et al., 2018).

5.2.2 Results

Optimization Dynamics. Figure 4 shows a typical run of our algorithm for ℓ_∞ -OT, exhibiting a common pattern: little progress at the beginning (during which \mathbf{P} is being aggressively adjusted), followed by a steep decline in the objective (during which both \mathbf{P} and Γ are increasingly modified in each step), after which convergence is reached. Note how the value of the optimization objective (left) and the accuracy in the translation task (right) are strongly correlated, particularly when compared against adversarial networks (Conneau et al., 2018). This is crucial because accuracy (shown here for expository purposes) *is not available* during the actual task, so model selection and early stopping are made based solely on the unsu-

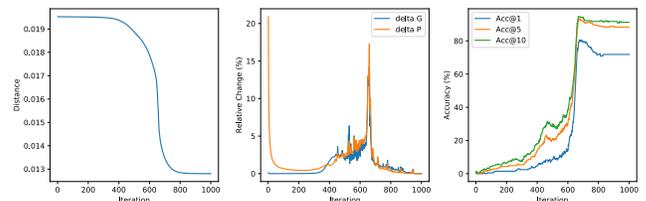


Figure 4: Optimization dynamics of ℓ_∞ -OT on the word translation task (EN→IT, 5K vocab). Left to right: objective, change in \mathbf{P} and Γ from previous iteration, and translation accuracy (Precision@K).

pervised objective. In addition, note that except for a small adjustment at the end of training, our method does not risk degradation by over-training, as is often the case for adversarial training alternatives.

Performance. The accuracy results (Table 1) show that our general framework performs on par with state-of-the-art approaches tailored this task, at a fraction of the computational cost, and yields an objective which is more faithful to the true metric of interest.

6 Discussion

We proposed a general formulation of optimal transport that accounts for global invariances in the underlying feature spaces, unifying various existing approaches to deal with such invariances. The problem allows for very efficient algorithms in two cases often found in practice. The experiments show that this framework provides a fast, principled and robust alternative to state-of-the-art methods for unsupervised word translation, delivering comparable performance. These results suggest that OT with invariances is a viable alternative to adversarial methods that infer correspondences from complex, often underdetermined, neural network maps.

⁵Despite its relevance, we do not include the OT-based method of Zhang et al. (2017b) in the comparison because their implementation required use of proprietary software.

Acknowledgements

The authors would like to thank Youssef Mroueh and Suvrit Sra for helpful discussions and remarks. The work was partially supported by the MIT-IBM project on adversarial learning of multi-modal and structured data, and by NSF CAREER award 1553284.

References

- Altschuler, J., J. Weed, and P. Rigollet (2017). “Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration”. In: *Advances in Neural Information Processing Systems*, pp. 1961–1971.
- Alvarez-Melis, D. and T. S. Jaakkola (2018). “Gromov-Wasserstein Alignment of Word Embedding Spaces”. In: *Empirical Methods in Natural Language Processing*.
- Alvarez-Melis, D., T. S. Jaakkola, and S. Jegelka (2018). “Structured Optimal Transport”. In: *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Arjovsky, M., S. Chintala, and L. Bottou (2017). “Wasserstein Generative Adversarial Networks”. In: *Proceedings of the 34th International Conference on Machine Learning*. Ed. by D. Precup and Y. W. Teh. Vol. 70. PMLR, pp. 214–223.
- Artetxe, M., G. Labaka, and E. Agirre (2017). “Learning bilingual word embeddings with (almost) no bilingual data”. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1, pp. 451–462.
- Artetxe, M., G. Labaka, and E. Agirre (2018a). “A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, pp. 789–798.
- Artetxe, M., G. Labaka, E. Agirre, and K. Cho (2018b). “Unsupervised Neural Machine Translation”. In: *International Conference on Learning Representations*.
- Bousquet, O., S. Gelly, I. Tolstikhin, C.-J. Simon-Gabriel, and B. Schoelkopf (2017). “From optimal transport to generative modeling: the VEGAN cookbook”. In: pp. 1–15.
- Conneau, A., G. Lample, M. Ranzato, L. Denoyer, and H. Jégou (2018). “Word Translation Without Parallel Data”. In: *International Conference on Learning Representations*, pp. 1–13.
- Courty, N., R. Flamary, A. Habrard, and A. Rakotomamonjy (2017a). “Joint Distribution Optimal Transportation for Domain Adaptation”. In: *Neural Information Processing Systems*.
- Courty, N., R. Flamary, D. Tuia, and A. Rakotomamonjy (2017b). “Optimal Transport for Domain Adaptation”. In: *Pattern Analysis and Machine Intelligence, IEEE Transactions on X.X*, pp. 1–14.
- Cuturi, M. (2013). “Sinkhorn distances: Lightspeed computation of optimal transport”. In: *Advances in Neural Information Processing Systems*, pp. 2292–2300.
- Eckstein, J. and W. Yao (2017). “Approximate ADMM algorithms derived from Lagrangian splitting”. In: *Computational Optimization and Applications* 68.2, pp. 363–405.
- Fung, P. (1995). “Compiling bilingual lexicon entries from a non-parallel English-Chinese corpus”. In: *Third Workshop on Very Large Corpora*.
- Genevay, A., G. Peyre, and M. Cuturi (2018). “Learning Generative Models with Sinkhorn Divergences”. In: *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*. Ed. by A. Storkey and F. Perez-Cruz. Vol. 84. Proceedings of Machine Learning Research. Playa Blanca, Lanzarote, Canary Islands: PMLR, pp. 1608–1617.
- Goodall, C. (1991). “Procrustes Methods in the Statistical Analysis of Shape”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 53.2, pp. 285–339.
- Gower, J. C. and G. B. Dijkstra (2004). *Procrustes problems*. Vol. 30. Oxford University Press on Demand.
- Grauman, K. and T. Darrell (2005). “The pyramid match kernel: Discriminative classification with sets of image features”. In: *Proceedings of the IEEE International Conference on Computer Vision*.
- Grave, E., A. Joulin, and Q. Berthet (2018). “Unsupervised Alignment of Embeddings with Wasserstein Procrustes”. In: *arXiv e-prints*.
- Haker, S. and A. Tannenbaum (2001). “Optimal mass transport and image registration”. In: *Proceedings - IEEE Workshop on Variational and Level Set Methods in Computer Vision, VLSM 2001*, pp. 29–36.
- Hardt, M. (2014). “Understanding alternating minimization for matrix completion”. In: *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE, pp. 651–660.

- Hashimoto, T. B., D. Alvarez-Melis, and T. S. Jaakkola (2016). “Word Embeddings as Metric Recovery in Semantic Spaces”. In: *Transactions of the Association for Computational Linguistics* 4, pp. 273–286.
- Hyvärinen, A. and E. Oja (2000). “Independent component analysis: algorithms and applications”. In: *Neural networks* 13.4-5, pp. 411–430.
- Jaggi, M. (2013). “Revisiting {Frank-Wolfe}: Projection-free sparse convex optimization”. In: *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*.
- Jain, P., P. Netrapalli, and S. Sanghavi (2013). “Low-rank matrix completion using alternating minimization”. In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, pp. 665–674.
- Mémoli, F. (2011). “Gromov–Wasserstein distances and the metric approach to object matching”. In: *Foundations of computational mathematics* 11.4, pp. 417–487.
- Mikolov, T., Q. V. Le, and I. Sutskever (2013). “Exploiting Similarities among Languages for Machine Translation”. In: *arXiv preprint arXiv:1309.4168v1*, pp. 1–10.
- Mokhtari, A., W. Shi, Q. Ling, and A. Ribeiro (2015). “Decentralized quadratically approximated alternating direction method of multipliers”. In: *Signal and Information Processing (GlobalSIP), 2015 IEEE Global Conference on*. IEEE, pp. 795–799.
- Park, H. (1991). “A parallel algorithm for the unbalanced orthogonal Procrustes problem”. In: *Parallel Computing* 17.8, pp. 913–923.
- Pele, O. and M. Werman (2009). “Fast and robust earth mover’s distances”. In: *2009 IEEE 12th International Conference on Computer Vision*. IEEE, pp. 460–467.
- Peyré, G. and M. Cuturi (2018). *Computational Optimal Transport*. Tech. rep.
- Peyré, G., M. Cuturi, and J. Solomon (2016). “Gromov-Wasserstein averaging of kernel and distance matrices”. In: *International Conference on Machine Learning*, pp. 2664–2672.
- Rangarajan, A., H. Chui, and F. L. Bookstein (1997). “The Softassign Procrustes Matching Algorithm”. In: *Lecture Notes in Computer Science* 1230, pp. 29–42.
- Rapp, R. (1995). “Identifying word translations in non-parallel texts”. In: *Proceedings of the 33rd annual meeting on Association for Computational Linguistics*. Association for Computational Linguistics, pp. 320–322.
- Schönemann, P. H. (1966). “A generalized solution of the orthogonal procrustes problem”. In: *Psychometrika* 31.1, pp. 1–10.
- Solomon, J., F. De Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas (2015). “Convolutional Wasserstein Distances : Efficient Optimal Transportation on Geometric Domains”. In: *ACM Transactions of Graphics (TOG)* 34, p. 66.
- Solomon, J., G. Peyré, V. G. Kim, and S. Sra (2016). “Entropic metric alignment for correspondence problems”. In: *ACM Transactions on Graphics (TOG)* 35.4, p. 72.
- Torresani, L., V. Kolmogorov, and C. Rother (2008). “Feature correspondence via graph matching: Models and global optimization”. In: *European Conference on Computer Vision*. Springer, pp. 596–609.
- Viklands, T. (2006). “Algorithms for the weighted orthogonal Procrustes problem and other least squares problems”. PhD thesis. Datavetenskap.
- Wang, C. and S. Mahadevan (2009). “Manifold alignment without correspondence”. In: *IJCAI International Joint Conference on Artificial Intelligence*, pp. 1273–1278.
- Zhang, M., Y. Liu, H. Luan, and M. Sun (2017a). “Adversarial training for unsupervised bilingual lexicon induction”. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1, pp. 1959–1970.
- Zhang, M., Y. Liu, H. Luan, and M. Sun (2017b). “Earth Mover’s Distance Minimization for Unsupervised Bilingual Lexicon Induction”. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1934–1945.
- Zhang, Y., D. Gaddy, R. Barzilay, and T. Jaakkola (2016). “Ten Pairs to Tag – Multilingual POS Tagging via Coarse Mapping between Embeddings”. In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. San Diego, California: Association for Computational Linguistics, pp. 1307–1317.

A Dealing with different samples sizes and dimensions

In the general case where we consider collections of potentially different size ($n \neq m$) of vectors of potentially different dimensionality ($d_x \neq d_y$), we have that \mathbf{X} and \mathbf{Y} are matrices of size $d_x \times m$ and $d_y \times n$, respectively, while \mathbf{P} is of size $d_x \times d_y$.

Note that for the transportation problem in (10), i.e.,

$$\max_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \langle \Gamma, \mathbf{X}^\top \mathbf{P} \mathbf{Y} \rangle,$$

the dimensions d_x, d_y are irrelevant. While the transportation coupling Γ is now of size $m \times n$, the problem is equally meaningful as before, i.e., DOT is well formulated and solved analogously in this case ($n \neq m$) as in the case of equal-sized marginals.

On the other hand, the transformation problem in (10), namely

$$\max_{\mathbf{P} \in \mathcal{F}} \langle \mathbf{X} \Gamma \mathbf{Y}^\top, \mathbf{P} \rangle,$$

is oblivious to n and m . However, when $d_x \neq d_y$, the problem no longer admits a closed form solution in general, so in this case this step requires optimization too. However, there exist various iterative algorithms to solve this problem—known as *unbalanced Procrustes*—efficiently (Gower and Dijkstra, 2004; Park, 1991; Viklands, 2006).

B The case $p = 1$

Recall that the Schatten ℓ_1 -norm is the nuclear norm $\|A\|_* = \sum_{i=1}^n \sigma_i(A)$. Therefore, the invariance set of interest is now

$$\mathcal{F}_1 = \{\mathbf{P} \mid \|\mathbf{P}\|_* = d\}, \quad (17)$$

which, as before, contains the identity matrix. Note that adding either condition in Lemma 4.1 yields, again, the set of orthonormal matrices.⁶ Therefore, in the case one wants to rely on Lemma 4.2 to solve the problem efficiently, this choice of invariance ends up being equivalent to the $p = \infty$ case described in Section 4.1. However, we remark that this equivalence is a consequence of the simplifying assumptions, and that one could still solve this problem with the Frank-Wolfe approach described in Section 4.3, in which case the two cases $p = \infty$ and $p = 1$ would indeed lead to different solutions.

⁶The intersection of the Schatten ℓ_2 and ℓ_∞ norm balls, defined in terms of that of the ℓ_2 and ℓ_∞ vector norm balls, occurs in the extreme points of the latter (see Fig. 1).

C Further Extensions

The framework proposed here can be further extended by considering other transformations that can be easily incorporated into the Procrustes problem framework. For example, scaling and translation can be added on top of orthogonal Procrustes and still yield a closed form solution (Gower and Dijkstra, 2004).

D Proofs

Lemma 4.1. *If any of the following conditions holds;*

1. $\forall \mathbf{P} \in \mathcal{F}$, \mathbf{P} is angle-preserving
2. $\exists k \geq 0 : \|\mathbf{P}\|_F = k \quad \forall \mathbf{P} \in \mathcal{F}$ and the matrix \mathbf{Y} is ν -whitened (i.e., $\mathbf{Y} \text{diag}(\mathbf{q})^2 \mathbf{Y}' = \mathbf{I}_d$).

then problem (9) is equivalent to

$$\max_{\Gamma \in \Pi(\mu, \mathbf{q})} \max_{\mathbf{P} \in \mathcal{F}} \langle \Gamma, \mathbf{X}' \mathbf{P} \mathbf{Y} \rangle = \max_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \max_{\mathbf{P} \in \mathcal{F}} \langle \mathbf{X} \Gamma \mathbf{Y}', \mathbf{P} \rangle \quad (18)$$

Proof. Suppose (1) holds, i.e., $\langle \mathbf{P} \mathbf{x}, \mathbf{P} \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{y} \rangle$ for every $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. Then, in particular $\|\mathbf{P} \mathbf{y}\|_2 = \|\mathbf{y}\|_2$ for every $\mathbf{y}^{(j)}$, and therefore:

$$\langle \mathbf{v}, \mathbf{q} \rangle = \sum_{j=1}^m \|\mathbf{P} \mathbf{y}^{(j)}\|_2 = \|\mathbf{y}^{(j)}\|_2$$

and therefore only the first term in (10) depends on \mathbf{P} or Γ , from which the conclusion follows. On the other hand, suppose (2) holds, and let $\tilde{\mathbf{Y}} = \mathbf{Y} \text{diag}(\mathbf{q})$, so that $\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}' = \mathbf{I}_d$. We have:

$$\begin{aligned} \langle \mathbf{v}, \mathbf{q} \rangle &= \sum_{i=1}^m q_j \|\mathbf{P} \mathbf{y}^{(j)}\|_2^2 \\ &= \sum_{j=1}^m \|\mathbf{P} \mathbf{y}^{(j)}\|_2^2 q_j \\ &= \|\mathbf{P} \tilde{\mathbf{Y}}\|_2^2 \\ &= \langle \mathbf{P} \tilde{\mathbf{Y}}, \mathbf{P} \tilde{\mathbf{Y}} \rangle = \langle \mathbf{P}, \mathbf{P} \tilde{\mathbf{Y}} \tilde{\mathbf{Y}}' \rangle = \|\mathbf{P}\|_F^2 = k^2, \end{aligned}$$

that is, $\langle \mathbf{v}, \mathbf{q} \rangle$ again does not depend on \mathbf{P} . This concludes the proof. \square

Lemma 4.2. *Let \mathbf{M} be a matrix with SVD decomposition $\mathbf{M} = \mathbf{U} \Sigma \mathbf{V}'$ and let $\Sigma = \text{diag}(\boldsymbol{\sigma})$, then*

$$\operatorname{argmax}_{\mathbf{P}: \|\mathbf{P}\|_p \leq k} \langle \mathbf{P}, \mathbf{M} \rangle = \mathbf{U} \text{diag}(\mathbf{s}) \mathbf{V}' \quad (19)$$

where \mathbf{s} is such that $\|\mathbf{s}\|_p \leq k$ and attains $\mathbf{s}' \boldsymbol{\sigma} = k \|\boldsymbol{\sigma}\|_q$, for $\|\cdot\|_q$ the dual norm of $\|\cdot\|_p$.

Proof. Suppose \mathbf{P} is such that $\|\mathbf{P}\|_p \leq k$, and let $\mathbf{U}_\mathbf{P} \text{diag}(\mathbf{s}) \mathbf{V}'_\mathbf{P}$ be its singular value decomposition. This implies that $\|\mathbf{s}\|_p = \|\mathbf{P}\| \leq k$. In addition,

$$\begin{aligned} \langle \mathbf{P}, \mathbf{M} \rangle &= \langle \mathbf{P}, \mathbf{U}\Sigma\mathbf{V}' \rangle \\ &= \langle \mathbf{U}'\mathbf{P}\mathbf{V}, \Sigma \rangle \\ &= \sum_{i=1}^d [\mathbf{U}'\mathbf{P}\mathbf{V}]_{ii} \sigma_i(\mathbf{M}) \\ &= \sum_{i=1}^d \mathbf{u}_i \mathbf{P} \mathbf{v}_i \sigma_i(\mathbf{M}) \leq \sum_{i=1}^d s_i \sigma_i(\mathbf{M}) = \langle \mathbf{s}, \boldsymbol{\sigma} \rangle \end{aligned}$$

Here, the inequality holds because, by definition of the SVD decomposition, for every i it must hold that $\|\mathbf{u}_i\|_2 = \|\mathbf{v}_i\|_2 = 1$ and

$$\mathbf{u}_i \mathbf{P} \mathbf{v}_i \leq \sup_{\substack{\mathbf{u} \perp \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_{i-1}\} \\ \mathbf{v} \perp \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_{i-1}\}}} \frac{\mathbf{u}' \mathbf{P} \mathbf{v}}{\|\mathbf{u}\| \|\mathbf{v}\|} \leq \sigma_i(\mathbf{P}) = s_i \quad (20)$$

Therefore:

$$\begin{aligned} \sup_{\mathbf{P}: \|\mathbf{P}\|_p \leq k} \langle \mathbf{P}, \mathbf{M} \rangle &\leq \sup_{\mathbf{s}: \|\mathbf{s}\|_p \leq k} \langle \mathbf{s}, \boldsymbol{\sigma} \rangle \\ &= k \sup_{\mathbf{s}: \|\mathbf{s}\|_p \leq 1} \langle \mathbf{s}, \boldsymbol{\sigma} \rangle = k \|\boldsymbol{\sigma}\|_q \end{aligned}$$

where the last equality follows from the definition of dual norm for vectors.

Conversely, take any vector \mathbf{s} with $\|\mathbf{s}\|_p = k$, and define $\tilde{\mathbf{P}}(\mathbf{s}) = \mathbf{U} \text{diag}(\mathbf{s}) \mathbf{V}'$. Clearly, $\|\tilde{\mathbf{P}}(\mathbf{s})\|_p = k$, so the supremum must satisfy:

$$\begin{aligned} \sup_{\mathbf{P}} \langle \mathbf{P}, \mathbf{M} \rangle &\geq \sup_{\mathbf{s}: \|\mathbf{s}\|_p \leq k} \langle \tilde{\mathbf{P}}(\mathbf{s}), \mathbf{M} \rangle \\ &= \sup_{\mathbf{s}: \|\mathbf{s}\|_p \leq k} \langle \mathbf{U} \text{diag}(\mathbf{s}) \mathbf{V}', \mathbf{U}\Sigma\mathbf{V}' \rangle \\ &= \sup_{\mathbf{s}: \|\mathbf{s}\|_p \leq k} \langle \text{diag}(\mathbf{s}), \Sigma \rangle = k \|\boldsymbol{\sigma}\|_q \end{aligned}$$

Therefore, we conclude that the optimal value of (19) is exactly $k \|\boldsymbol{\sigma}\|_q$.

Furthermore, (20) holds with equality if and only if $(\mathbf{u}_i, \mathbf{v}_i)$ coincide with the left and right singular vectors of \mathbf{P} . Thus, any \mathbf{P} maximizing (19) must have the form $\mathbf{P} = \mathbf{U} \text{diag}(\mathbf{s}) \mathbf{V}'$, with $\|\mathbf{s}\|_p \leq k$ and $\langle \mathbf{s}, \boldsymbol{\sigma} \rangle = k \|\boldsymbol{\sigma}\|_q$, as stated. \square

Lemma 4.3. Consider the Gromov-Wasserstein problem for discrete measures μ and ν (Peyré et al., 2016):

$$\min_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \sum_{i,j,k,l} L(\mathbf{C}_{ik}^x, \mathbf{C}_{jl}^y) \Gamma_{ij} \Gamma_{kl}, \quad (21)$$

where $(\mathbf{C}^x, \mathbf{p})$ and $(\mathbf{C}^y, \mathbf{q})$ are (intra-space) measured similarity matrices and L is a loss function. For the

choice of cosine similarity and squared loss $L(a, b) = \frac{1}{2}|a - b|^2$, Problems (15) and (14) are equivalent.

Proof. For the choice of cosine metric, and assuming without loss of generality that the columns of \mathbf{X} and \mathbf{Y} are normalized, the similarity matrices are given by $\mathbf{C}^x = \mathbf{X}^\top \mathbf{X}$ and $\mathbf{C}^y = \mathbf{Y}^\top \mathbf{Y}$. In addition, let L be the ℓ_2 loss, i.e., $L(a, b) = |a - b|^2$. Then the objective in problem (21) becomes:

$$\begin{aligned} \mathcal{L}(\Gamma) &= \sum_{i,j,k,l} (\mathbf{C}_{ik}^x - \mathbf{C}_{jl}^y)^2 \Gamma_{ij} \Gamma_{kl} \\ &= \sum_{i,j,k,l} (\mathbf{C}_{ik}^x)^2 \Gamma_{ij} \Gamma_{kl} - 2 \sum_{i,j,k,l} (\mathbf{C}_{ik}^x \mathbf{C}_{jl}^y) \Gamma_{ij} \Gamma_{kl} \\ &\quad + \sum_{i,j,k,l} (\mathbf{C}_{jl}^y)^2 \Gamma_{ij} \Gamma_{kl} \end{aligned}$$

Since $\Gamma \in \Pi(\mathbf{p}, \mathbf{q})$, the first of these terms becomes

$$\sum_{i,k} (\mathbf{C}_{ik}^x)^2 \sum_{j,l} \Gamma_{ij} \Gamma_{jl} = \sum_{i,k} (\mathbf{C}_{ik}^x)^2 \mathbf{p}_i \mathbf{p}_k = \mathbf{p}^\top (\mathbf{C}^x)^2 \mathbf{p}$$

where \mathbf{p} is the vector of probabilities in empirical distribution μ , and the last equation follows from the definition of the transportation polytope. Crucially, this term does not depend on Γ anymore. Analogously, the last term in $\mathcal{L}(\Gamma)$ does not depend on Γ either, so

$$\arg\min_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \mathcal{L}(\Gamma) = \arg\max_{\Gamma \in \Pi(\mathbf{p}, \mathbf{q})} \sum_{i,j,k,l} (\mathbf{C}_{ik}^x \mathbf{C}_{jl}^y) \Gamma_{ij} \Gamma_{kl} \quad (22)$$

On the other hand, consider problem (14). The objective it seeks to maximize is

$$\begin{aligned} \|\mathbf{X}\Gamma\mathbf{Y}^\top\|_F^2 &= \langle \mathbf{X}\Gamma\mathbf{Y}^\top, \mathbf{X}\Gamma\mathbf{Y}^\top \rangle \\ &= \langle \mathbf{X}^\top \mathbf{X} \Gamma, \Gamma \mathbf{Y} \mathbf{Y}^\top \rangle \\ &= \sum_{i=1}^n \sum_{l=1}^m [\mathbf{X}^\top \mathbf{X} \Gamma]_{il} [\Gamma \mathbf{Y} \mathbf{Y}^\top]_{il} \\ &= \sum_{i=1}^n \sum_{l=1}^m [\mathbf{C}^x \Gamma]_{il} [\Gamma \mathbf{C}^y]_{il} \\ &= \sum_{i=1}^n \sum_{l=1}^m \left(\sum_{k=1}^n \mathbf{C}_{ik}^x \Gamma_{kl} \right) \left(\sum_{j=1}^m \Gamma_{ij} \mathbf{C}_{jl}^y \right) \\ &= \sum_{i=1}^n \sum_{l=1}^m \sum_{k=1}^n \sum_{j=1}^m \mathbf{C}_{ik}^x \Gamma_{kl} \Gamma_{ij} \mathbf{C}_{jl}^y \end{aligned}$$

which is exactly the objective in (22). Hence, Problems (14) and (21) are indeed equivalent. \square

E The Algorithm

Algorithm 1 Optimal Transport with Invariances

Inputs:

- Data matrices and histograms $(\mathbf{X}, \mathbf{p}), (\mathbf{Y}, \mathbf{q})$
- Order of invariance p and radius k_p
- Initial/final entropy regularization λ_0 and $\underline{\lambda}$, decay rate η

```

// Initialize feasible transformation in  $\mathcal{F}_p$ 
 $\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}^\top \leftarrow \text{SVD}(\text{RANDOMMATRIX}(d \times d))$ 
 $\boldsymbol{\sigma} \leftarrow \text{diag}(\mathbf{\Sigma})$ 
 $\mathbf{s} \leftarrow k_p \cdot \boldsymbol{\sigma} / \|\boldsymbol{\sigma}\|_p$ 
 $\mathbf{P} = \mathbf{U} \text{diag}(\mathbf{s}) \mathbf{V}^\top$ 
 $\lambda \leftarrow \lambda_0$ 
while not converged do
  // Compute distances w.r.t. current mapping  $\mathbf{P}$ 
   $\mathbf{C}_\mathbf{P} \leftarrow \text{PAIRWISEDISTANCES}(\mathbf{X}, \mathbf{P}\mathbf{Y})$ 
  // Solve regularized OT via Sinkhorn iterations
   $\mathbf{b} \leftarrow \mathbb{1}, \mathbf{K} \leftarrow \exp\{-\mathbf{C}_\mathbf{P}/\lambda\}$ 
  while not converged do
     $\mathbf{a} \leftarrow \mathbf{p} \oslash \mathbf{K}\mathbf{b}$ 
     $\mathbf{b} \leftarrow \mathbf{q} \oslash \mathbf{K}^\top \mathbf{a}$ 
  end while
   $\mathbf{\Gamma} \leftarrow \text{diag}(\mathbf{a}) \mathbf{K} \text{diag}(\mathbf{b})$ 
  // Solve generalized Procrustes problem
   $\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}^\top \leftarrow \text{SVD}(\mathbf{X}\mathbf{\Gamma}\mathbf{Y}^\top)$ 
   $\boldsymbol{\sigma} \leftarrow \text{diag}(\mathbf{\Sigma})$ 
   $q \leftarrow \frac{p}{p-1}$ 
   $\mathbf{s} \leftarrow k_p \cdot \boldsymbol{\sigma}^{q-1} / \|\boldsymbol{\sigma}^{q-1}\|_p$ 
   $\mathbf{P} = \mathbf{U} \text{diag}(\mathbf{s}) \mathbf{V}^\top$ 
  // Anneal entropy regularization
   $\lambda \leftarrow \max\{\lambda * \eta, \underline{\lambda}\}$ 
end while
return  $\mathbf{\Gamma}, \mathbf{P}$ 

```

F Solving very large problems

While direct application of Algorithm 1 leads to high-quality solutions for small and mid-sized problems, scaling up to very large sets of points—e.g., hundreds of thousands of word embeddings in the word translation application—can be prohibitive.

We address this issue by dividing the problem into two phases. In the first stage, we solve a smaller problem (by taking a subsample of k points on each domain thus leading to smaller $\mathbf{\Gamma}$ and faster OT solution, but same size of \mathbf{P}). Once the first phase reaches convergence, we use the solution \mathbf{P}^* of the first stage to initialize the full-size problem. Note that while this might resemble other approaches that also consider a reduced set of points in their initialization step (Conneau et al., 2018; Grave et al., 2018), a crucial dif-

ference is that here we rely on the same optimization problem (16) in both stages, although with different problem sizes.

We experimented with various choices of parameter k , and observed that the algorithm is remarkably robust to the choice of this parameter. We conjecture that the ordering in which word embeddings are provided (higher-frequency words first, in every language) helps ensure that the solution of the initial problem of reduced size is consistent with the full-size problem.⁷

While the end performance is consistent regardless of the choice of sub-sample size k , there is naturally a trade-off in run time of the two stages. While solving a smaller initial problem is obviously faster, we observed that in such cases the second stage required more iterations to converge, suggesting that the initial \mathbf{P}^* fed into the second stage was of lower quality (further from the optimal for the full-size problem). In the results presented in Section 5.2, we take k as large as possible while keeping the time-per-iteration reasonable: $k = 5000$.

Note that this strategy of *bootstrapping* solutions of smaller problems can be applied repeatedly, to increasingly grow the problem size over multiple stages. While we did not require to do so in our experiments, it might be an appealing approach for solving extremely large problems.

⁷This, in fact, points to an issue mostly ignored in previous work on this task: the order of the word embeddings *leaks* important—albeit noisy—correspondence information, which various methods presented as ‘fully-unsupervised’ seem to rely on one way or another, yet rarely acknowledge it.