



Approximation of Wasserstein distance with Transshipment

Nicolas Papadakis

► To cite this version:

| Nicolas Papadakis. Approximation of Wasserstein distance with Transshipment. 2019. hal-01997598

HAL Id: hal-01997598

<https://hal.science/hal-01997598>

Preprint submitted on 22 Jul 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Approximation of Wasserstein distance with Transshipment

Nicolas Papadakis¹

¹CNRS, Univ. Bordeaux, IMB, F-33400 Talence, France

Abstract

An algorithm for approximating the p -Wasserstein distance between histograms defined on unstructured discrete grids is presented. It is based on the computation of a barycenter constrained to be supported on a low dimensional subspace, which corresponds to a transshipment problem. A multi-scale strategy is also considered. The method provides sparse transport matrices and can be applied to large and non structured data.

1 Introduction

The computation of optimal transportation between two discrete normalized measures μ_x and μ_y defined from $x \in \Omega \subset \mathbb{R}^d$ to $[0; 1]$ remains a challenging problem when an accurate discretization of the domain Ω is considered. Optimal transportation requires to define a ground distance between points $x, y \in \mathbb{R}^d$ of the domain Ω . This ground metric is then used to measure how much it costs to move $\mu_x(x)$ to $\mu_y(y)$. Ground distances $\|x - y\|^p$ are here considered for $p \geq 1$, which leads to the p -Wasserstein distance between μ_x and μ_y (see [45, 36, 33] for more detailed introductions). Such distances give robust metrics in retrieval applications for $1 \leq p < 2$ [35, 31, 32, 21]. The underlying sparse transport matrix is also a useful tool for interpolation and transfer purposes [34, 17, 42].

Computing Wasserstein distances The computation of Wasserstein distances is only explicit for $d = 1$. A standard approach to estimate Wasserstein distances when $d > 1$ consists in pre-computing a cost matrix $\|x - y\|^p$ for $x, y \in \Omega$ and then estimating the whole transport matrix which dimension grows quadratically with $n = |\Omega|$. Linear programming or transportation simplex [25, 19] can be applied to estimate p -Wasserstein distance for $p \geq 1$ but they are limited in practice to low dimensional problems, i.e. small values of n , for complexity and storage issues. The implementation provided in [11] nevertheless allows to tackle problems of interesting dimensions, i.e. Ω discretized with more than $n = 10^4$ points. By exploiting the sparsity of the transport map, multi-scale strategy [28] or grid refinement [38] can deal with larger problems by solving iteratively sparse low dimensional problems with linear programming. More efficient algorithms can be considered in the specific case of 1-Wasserstein distances [23].

Approximation of Wasserstein distance For large scale problems, approximated sparse transport matrices and Wasserstein distances can be obtained by considering successive one dimensional problems with the so-called sliced Wasserstein distance [34, 10]. The entropic regularization of the transport map proposed in [14] is another relevant way to deal with problems of high dimension. Given a regularization parameter $\gamma > 0$, it approximates the true Wasserstein distance (that corresponds in this setting to $\gamma = 0$) with the well-known Sinkhorn algorithm. The estimated transport matrices are nevertheless dense and should be truncated with care for interpolation purposes. When data are discretized on a uniform grid, the estimation of the distance can be obtained through iterative convolutions [42] with a kernel $K_p = \exp(-\|x - y\|^p/\gamma)$. This leads to very fast algorithm

for the p -Wasserstein distance as the Kernel is separable when considering the L^p norm $||.||_p$, and only convolutions and storage of d one-dimensional kernels are necessary. For small values of γ , numerical instabilities nevertheless arise and dedicated attention must be given to the implementation by considering for instance decaying values of γ , stabilization in log domain or multi-resolution approaches [39]. The numerical convergence is also reduced with low values of γ , and overrelaxation [37, 43] or greedy coordinate descents [3] approaches have been proposed to tackle this issue. When the kernel K_p is not separable, it must be carefully truncated to zero to store sparse matrices. If not considering more complex and adaptive truncations [39], it limits the possible amplitude of the transport, which may be annoying in case of large displacements. Other regularizations can then be of interest [16, 9, 41], namely when it is suitable to recover sparse transport matrices.

However, as mentioned before, this kind of techniques only leads to efficient implementations for densities discretized on structured grids. Moreover, if large displacements are involved in the data, the Wasserstein distance can not be accurately approximated, since numerical instabilities arises with $||x - y||^p/\gamma$ when γ goes to 0.

Wasserstein barycenters and Transshipment The p -Wasserstein barycenter, as introduced in [1], can be used to approximate the p -Wasserstein distance between measures involving a large number n of dirac masses. As proposed in [46] for clustering problems, the discrete barycenter between 2 discrete measures can be parameterized with a weighted sum of $\kappa \ll n$ diracs to obtain a low dimensional problem that corresponds to the transshipment problem of n resources with κ intermediate locations. The sum of the distances between each data and the barycenter then gives an approximation of the effective distance, as it has been proposed for 1-Wasserstein distances [5]. Statistical properties of such a method have been later studied in [18, 30], where it has been underlined the robustness of this low rank regularization of the transport matrix to data outliers. These ideas have been extended to discrete approximation of barycenter between continuous measures in [12].

Content The use of low dimensional barycenters is the point of view adopted in this note to propose fast approximation of Wasserstein distances involving sparse transport matrices. As in [46], at each iteration of the presented algorithm, the κ locations of the barycenter are updated and a linear program of dimension $2\kappa n$ is solved. Compared to the dimension n^2 for the classical Wasserstein distance, the overall complexity of the algorithm becomes attractive for high dimensional data. In [46] and [18, 30], this low dimensional barycenter problem is respectively solved with the Alternating Direction Method of Multipliers or Sinkhorn iterations [15]. The barycenter problem between two discrete densities is here seen as a transshipment problem and solved up to numerical accuracy with an efficient network simplex graph algorithm based on the work of [11]. Then κ transportation subproblems are solved in parallel to recover a sparse transport matrix. Following [39, 28, 24], this leads to the design of a multi-scale barycenter scheme, to iteratively refine the transport matrix and the associated approximation of the Wasserstein distance.

The theoretical computer science community has recently provided improved bounds for graph based algorithms solving approximate transportation problems up to ϵ additive or $1 + \epsilon$ multiplicative errors. This namely includes the optimal transportation problem [8], its entropic regularization [2] or the transshipment problem [6]. As for the sliced method [34], the presented approach does not have such theoretical guarantees but very good performances are observed in practice. More precisely, approximate p -Wasserstein distances between cloud points of 10^5 elements are obtained in a few minutes without involving prohibitive memory storage issues. The proposed empirical algorithm can also be directly applied to non structured data. Thanks to the multi-scale refinement approach, a sparse transport map is provided which can be of interest for interpolation purposes.

Outline Section 2 details how to approximate the Wasserstein distance when computing the barycenter between two discrete measures. The strategy of [46] is recalled in the general context of p -Wasserstein distance. A multi-scale algorithm for recovering sparse transport matrices is finally presented. The performances of the algorithm are discussed in section 3 through extensive experiments realized on the benchmark [40]. Numerical results show that this whole empirical process is efficient, namely when one of the two input data is spatially regular.

2 Approximate Wasserstein distance from barycenter estimation

Let μ_x and μ_y be two discrete measures defined on $\Omega \subset \mathbb{R}^d$, $d \geq 1$: $\mu_x = \sum_{i=1}^m w_i^x \delta_{x_i}$ and $\mu_y = \sum_{j=1}^n w_j^y \delta_{y_j}$. These measures are supported at positions $\{x_i\}_{i=1}^m$ and $\{y_j\}_{j=1}^n$, $x_i, y_j \in \Omega \subset \mathbb{R}^d$. They have normalized positive weights vectors $w^x \in \mathcal{S}_m$ and $w^y \in \mathcal{S}_n$, where \mathcal{S}_n is the simplex of size n defined as $\mathcal{S}_n = \{w \in \mathbb{R}_+^n, \text{ s.t. } \sum_{i=1}^n w_i = 1\}$. For $p \geq 1$, let $c^{xy} \in \mathbb{R}_+^{m \times n}$ be the ground cost matrix over Ω defined as $c_{ij}^{xy} = \|x_i - y_j\|^p$, which corresponds to the power p of the distance related to a given norm on \mathbb{R}^d . Then the p -Wasserstein distance between discrete measures μ_x and μ_y is

$$W_p^p(\mu_x, \mu_y) = \min_{\gamma \in \mathcal{P}(w^x, w^y)} \langle \gamma, c^{xy} \rangle := \sum_{ij} \gamma_{ij} c_{ij}^{xy}, \quad (1)$$

with the set of admissible transport matrices

$$\mathcal{P}(w^x, w^y) = \{\gamma \in \mathbb{R}_+^{m \times n}, \text{ s.t. } \gamma 1_n = w^x, \gamma^\top 1_m = w^y\}, \quad (2)$$

and where 1_n the vector full of ones in \mathbb{R}^n . This problem can be efficiently solved with linear programming. It can also be formulated through a directed graph containing $m + n$ nodes and mn vertices. The final transport matrix γ is very sparse in practice (at most $m + n - 1$ non null entries) but the involved complexity and memory storage scale with the product of data dimensions nm . For latter purpose, Algorithm 1 details the function estimating the distance $W_p^p(\mu_x, \mu_y)$.

Algorithm 1 Estimate p -Wasserstein distance $W_p^p(\mu_x, \mu_y)$

- 1: **procedure** WP(x, w^x, y, w^y, p)
 - 2: Set $c_{ij}^{xy} = \|x_i - y_j\|^p$
 - 3: Solve problem (1) under the constraints (2) to get γ
 - 4: Set $W = \langle \gamma, c^{xy} \rangle$
 - 5: **return** W, γ
-

2.1 Interpolation and barycenters

Let γ^{xy} be an optimal transport matrix solution of (1) and μ_t be the following interpolation between measures μ_x and μ_y for $t \in [0; 1]$:

$$\mu_t = \sum_{ij} \gamma_{ij}^{xy} \delta_{x_i + t(y_j - x_i)} := \sum_k w_k^t \delta_{z_k}. \quad (3)$$

This interpolation is the discrete analogue [33] to the geodesic between μ_x and μ_y defined by the McCann's interpolation [26]. Discrete versions of some results in [36] can now be expressed.

Proposition 1 For $p \geq 1$, γ^{xy} a solution of (1), μ_t defined in (3) and $t \in [0; 1]$, the following relations hold:

$$W_p(\mu_x, \mu_t) = t W_p(\mu_x, \mu_y) \quad W_p(\mu_y, \mu_t) = (1 - t) W_p(\mu_x, \mu_y), \quad (4)$$

so that

$$W_p(\mu_x, \mu_y) = W_p(\mu_x, \mu_t) + W_p(\mu_y, \mu_t). \quad (5)$$

Proof. Observing that

$$W_p^p(\mu_x, \mu_t) = \min_{\gamma \in \mathcal{P}(w^x, w^t)} \sum_{ik} \gamma_{ik} \|x_i - z_k\|^p \leq \sum_{ij} \gamma_{ij}^{xy} \|t(x_i - y_j)\|^p,$$

leads to the upper bound $W_p(\mu_x, \mu_t) \leq tW_p(\mu_x, \mu_y)$, $\forall t \in [0; 1]$. It can be shown in the same way that $W_p(\mu_y, \mu_t) \leq (1-t)W_p(\mu_x, \mu_y)$. Since W_p is a distance, the triangle inequality $W_p(\mu_x, \mu_y) \leq W_p(\mu_x, \mu_t) + W_p(\mu_t, \mu_y)$ involves that the previous relations are in fact equalities. \square

Following [1], it can be shown that the mid interpolation $\mu_{1/2}$ is solution of the p -Wasserstein barycenter problem between μ_x and μ_y with weights $(1/2, 1/2)$.

Proposition 2 *For $p \geq 1$ and the interpolation $\mu_{1/2}$ defined in (3), it holds that*

$$\frac{1}{2}W_p^p(\mu_x, \mu_{1/2}) + \frac{1}{2}W_p^p(\mu_y, \mu_{1/2}) = \left(\frac{1}{2}W_p(\mu_x, \mu_y)\right)^p \quad (6)$$

and $\mu_{1/2}$ is a solution of the p -Wasserstein barycenter problem:

$$\mu_{1/2} \in \operatorname{argmin}_{\mu} \frac{1}{2}W_p^p(\mu_x, \mu) + \frac{1}{2}W_p^p(\mu_y, \mu). \quad (7)$$

Proof. From (4), it can first be noticed that:

$$\frac{1}{2}(W_p^p(\mu_x, \mu_{1/2}) + W_p^p(\mu_y, \mu_{1/2})) = \frac{1}{2} \left(2 \left(\frac{1}{2}W_p(\mu_x, \mu_y) \right)^p \right) = \left(\frac{1}{2}W_p(\mu_x, \mu_y) \right)^p. \quad (8)$$

Since W_p is a distance and the function $|\cdot|^p$, is convex for $p \geq 1$, it can next be observed that $\forall \mu$:

$$\left(\frac{1}{2}W_p(\mu_x, \mu_y) \right)^p \leq \left(\frac{1}{2}(W_p(\mu_x, \mu) + W_p(\mu_y, \mu)) \right)^p \leq \frac{1}{2}(W_p^p(\mu_x, \mu) + W_p^p(\mu_y, \mu)). \quad (9)$$

Combing relations (8) and (9) implies that $\mu_{1/2}$ is a solution of the barycenter problem (7). \square

Existence (and uniqueness for $p > 1$) of Wasserstein barycenters have been deeply studied in [1, 20]. As stated in the following proposition, the p -Wasserstein distance can be obtained for any $p \geq 1$ through the resolution of the barycenter problem (7).

Proposition 3 *Let $\tilde{\mu}$ be a solution of the p -Wasserstein barycenter problem (7), then*

$$W_p(\mu_x, \mu_y) = W_p(\mu_x, \tilde{\mu}) + W_p(\mu_y, \tilde{\mu}) \quad (10)$$

and $W_p(\mu_x, \tilde{\mu}) = W_p(\mu_y, \tilde{\mu})$ for $p > 1$.

Proof. Assuming, without loss of generality, that $W_p(\mu_y, \tilde{\mu}) = \alpha W_p(\mu_x, \tilde{\mu})$, with $\alpha \in [0; 1]$, it holds from Proposition 2 that $\tilde{\mu}$ is a solution of (7) iff $2^{p-1}(1 + \alpha^p)W_p^p(\mu_x, \tilde{\mu}) = W_p^p(\mu_x, \mu_y)$. Assuming that $W_p(\mu_x, \tilde{\mu}) + W_p(\mu_y, \tilde{\mu}) = (1 + \alpha)W_p(\mu_x, \tilde{\mu}) > W_p(\mu_x, \mu_y)$, it involves that

$$\begin{aligned} 2^{p-1}(1 + \alpha^p) &< (1 + \alpha)^p < (1 + \alpha)2^{p-1} \\ \alpha^p &< \alpha, \end{aligned} \quad (11)$$

which is impossible $\forall \alpha \in [0; 1]$ and $p \geq 1$. Thus relation (10) holds by contradiction and $2^{p-1}(1 + \alpha^p) = (1 + \alpha)^p$. With the very same arguments than in (11), it leads to $\alpha = 1$ for $p > 1$. \square

2.2 Barycenter computation

A discrete barycenter $\mu_{1/2}$ between μ_x and μ_y can be obtained by solving (7) with the distance (1). As done in [15] with an additional entropic regularization, the mid barycenter $\mu_{1/2}$ between μ_x and μ_y can be constrained to be supported on a set of κ dirac masses, i.e. $\mu_{1/2} = \sum_{k=1}^{\kappa} w_k^z \delta_{z_k}$, with positions $z_k \in \mathbb{R}^d$ and weights $w^z \in \mathcal{S}_{\kappa}$. The barycenter problem can be rewritten as:

$$\min_{\substack{\{z_k\}_{k=1}^{\kappa} \in \Omega^{\kappa} \\ w^z \in \mathcal{S}_{\kappa}}} W_p^p(\mu_x, \mu_z) + W_p^p(\mu_y, \mu_z) = \min_{\substack{\{z_k\}_{k=1}^{\kappa} \in \Omega^{\kappa} \\ (\gamma^x, \gamma^y) \in \tilde{\mathcal{P}}(w^x, w^y)}} \langle \gamma^x, c^{xz} \rangle + \langle \gamma^y, c^{yz} \rangle, \quad (12)$$

with the admissible set of matrices

$$\tilde{\mathcal{P}}(w^x, w^y) = \{\gamma^x \in \mathbb{R}_+^{m \times \kappa}, \gamma^y \in \mathbb{R}_+^{n \times \kappa}, \text{ s.t. } \gamma^x \mathbf{1}_{\kappa} = w^x, \gamma^y \mathbf{1}_{\kappa} = w^y, (\gamma^x)^{\top} \mathbf{1}_m - (\gamma^y)^{\top} \mathbf{1}_n = \mathbf{0}_{\kappa}\}$$

and the cost matrices $c_{ik}^{xz} = \|x_i - z_k\|^p$ and $c_{jk}^{yz} = \|y_j - z_k\|^p$. Notice that the weight vector of the barycenter $w^z \in \mathbb{R}^{\kappa}$ is implicitly included in the set of constraints: $(\gamma^x)^{\top} \mathbf{1}_m = (\gamma^y)^{\top} \mathbf{1}_n = w^z$. The positions z_k here act as intermediate locations where the mass has to transit from x to y . Following [46] and as illustrated in Figure 1, the idea behind this modeling is to consider a limited number of transshipment locations κ to speed up the computation.




Figure 1: Illustration of transshipment with $\kappa = 3$ intermediate locations.

The problem (12) is separately convex with respect its variables γ^x , γ^y and z_k . For $p > 1$, the coupling terms are differentiable and alternate minimization over transport matrices (γ^x, γ^y) and dirac positions z_k converges [44] to a saddle point. Following (10), the Wasserstein distance $W_p(\mu_x, \mu_y)$ can be approximated with $\tilde{W}_p(\mu_x, \mu_y)$ obtained from transport matrices γ^x and γ^y solutions of the problem (12) as

$$W_p(\mu_x, \mu_y) \leq \tilde{W}_p(\mu_x, \mu_y) = W_p(\mu_x, \mu_{1/2}) + W_p(\mu_y, \mu_{1/2}) = \langle \gamma^x, c^{xz} \rangle^{1/p} + \langle \gamma^y, c^{yz} \rangle^{1/p}.$$

The alternate optimization steps for solving (12) are now detailed and the process is summed up in Algorithm 2.

Remark 1 The norm $\|\cdot\|_2$ is usually taken as reference for the p -Wasserstein distance. From the equivalence of norms in finite dimensions and as can be done with entropic regularization to make the problem more tractable numerically, the L^p norm $\|x_i - y_j\|_p^p = \sum_{s=1}^d (x_i^s - y_j^s)^p$ is here considered for computing the p -Wasserstein distance.

Algorithm 2 Estimate Approximate p -Wasserstein $\tilde{W}_p^p(\mu_x, \mu_y)$ through transshipment

```

1: procedure BARWP( $x, w^x, y, w^y, p, \kappa$ )
2:    $\epsilon = 10^{-3}$ 
3:   Initialize  $\kappa$  positions  $z_k$  randomly from the  $m$  and  $n$  positions  $x_i$  and  $y_j$ 
4:   repeat
5:     Set  $c_{ik}^{xz} = \|x_i - z_k\|^p$  and  $c_{jk}^{yz} = \|y_j - z_k\|^p$ 
6:     Get  $\gamma^x$  and  $\gamma^y$  by solving transshipment problem (15) under the constraints (16)
7:     Set  $\tilde{z}_k = z_k$ 
8:     Update positions  $z_k$  from  $\gamma^x$  and  $\gamma^y$  by following Sec. 2.2.1
9:   until  $\|z_k - \tilde{z}_k\| / \|\tilde{z}_k\| < \epsilon$ 
10:  Set  $\tilde{W} = (\langle c^{xz}, \gamma^x \rangle^{1/p} + \langle c^{yz}, \gamma^y \rangle^{1/p})^p$ 
11:  return  $\tilde{W}, \gamma^x, \gamma^y$ 

```

2.2.1 Update of positions

To update positions $z_k \in \mathbb{R}^d$, the problem (12) is solved for fixed transport matrices γ^x and γ^y . For $p \geq 1$, this leads to d convex problems, that can be solved in parallel for each space dimension $s=1 \dots d$. Dimension indexes s are thus omitted in the following and the problem writes

$$\min_{\{z_k\}_{k=1}^\kappa} \sum_{i=1}^M \sum_{k=1}^\kappa \gamma_{ik}^x |x_i - z_k|^p + \sum_{j=1}^N \sum_{k=1}^\kappa \gamma_{jk}^y |y_j - z_k|^p, \quad (13)$$

Different strategies are considered according to p . Some locations z_k may become useless if $\sum_i \gamma_{ik}^x = \sum_j \gamma_{jk}^y = 0$. In this case the corresponding z_k are removed and κ is decreased. When considering large scale problems and few transshipment locations κ , this almost never happens.

Case W_2 . For $p = 2$, there exists an explicit update formula of the dirac positions to find the unique minimizer of (13) with respect to z :

$$z_k = \frac{((\gamma^x)^\top x + (\gamma^y)^\top y)_k}{((\gamma^x)^\top \mathbf{1}_m + (\gamma^y)^\top \mathbf{1}_n)_k}. \quad (14)$$

This step acts like the cluster position update in a κ -mean algorithm. It realizes for each z_k a weighted mean of the positions x_i 's and y_j 's according to γ_{ik}^x and γ_{jk}^y .

Case $W_p, p > 2$. For $p > 2$ the problem (13) admits a unique minimizer and is twice differentiable. Newton's method can be considered to approximate the solution:

$$z_k^{\ell+1} = z_k^\ell - \frac{\left(\left(\gamma^x \otimes |c^{xz^\ell}|^{p-2} \otimes c^{xz^\ell} \right)^\top \mathbf{1}_m + \left(\gamma^y \otimes |c^{yz^\ell}|^{p-2} \otimes c^{yz^\ell} \right)^\top \mathbf{1}_n \right)_k}{(p-1) \left((\gamma^x \otimes |c^{xz^\ell}|^{p-2})^\top \mathbf{1}_m + (\gamma^y \otimes |c^{yz^\ell}|^{p-2})^\top \mathbf{1}_n \right)_k},$$

where $c_{ik}^{xz^\ell} = z_k^\ell - x_i$ and $c_{jk}^{yz^\ell} = z_k^\ell - y_j$, while \otimes denotes the elementwise product between matrices and the power $p-2$ is also element-wise.

Case W_1 . When $p = 1$, a global optimum of (13) can be obtained by taking each z_k as a weighted median of the positions x_i and y_j with respect to the weights γ_{ik}^x and $(\gamma^y)_{jk}$. This operation mainly requires to sort the value of the x_i 's and y_j 's along each dimension. Notice that alternate minimization on problem (12) may not converge when $p = 1$.

Case $W_p, p \in]1; 2[$. The problem does not admits a second derivative. An iterative scheme is then considered by decomposing $|z_k - x_i|^p = |z_k - x_i|^{p-1} |z_k - x_i|$ and solving successive weighted median problems between the positions x_i and y_j with the weights $\gamma_{ik}^x |z_k^\ell - x_i|^{p-1}$ and $\gamma_{jk}^y |z_k^\ell - y_j|^{p-1}$.

2.2.2 Update of transport matrices through Transshipment

For fixed z_k , problem (12) can be solved with classic linear programming optimization tools. The interesting point is that this barycenter problem is a transshipment problem with κ intermediate locations. It can therefore be formulated in terms of a directed graph with $m + \kappa + n$ vertices (i.e. x_i, z_k and y_j) and $(m\kappa + \kappa n)$ edges e_{ik} and e_{kj} that correspond to the transport matrices γ_{ik}^x and γ_{jk}^y . The following cost function is then minimized

$$(\gamma^x, (\gamma^y)^\top) \in \underset{e=\{\{e_{ik}\}, \{e_{kj}\}\}}{\operatorname{argmin}} \sum_{i=1}^m \sum_{k=1}^{\kappa} e_{ik} c_{ik}^{xz} + \sum_{k=1}^{\kappa} \sum_{j=1}^n e_{kj} c_{jk}^{yz}, \quad (15)$$

under the set of constraint $\tilde{\mathcal{P}}(w^x, w^y)$ that translates into:

$$\left\{ \begin{array}{ll} e_{ik}, e_{kj} \geq 0 & i = 1 \dots m, k = 1 \dots \kappa, j = 1 \dots n \\ \sum_{k=1}^{\kappa} e_{ik} = w_i^x & i = 1 \dots m \\ \sum_{j=1}^n e_{kj} - \sum_{i=1}^m e_{ik} = 0 & k = 1 \dots \kappa \\ -\sum_{j=1}^n e_{kj} = -w_j^y & j = 1 \dots n \end{array} \right. \quad (16)$$

This problem can be efficiently solved with the network simplex algorithm [29]. An extension of the original non sparse implementation proposed in [11] is here considered.

2.2.3 Discussion

It is well known [1, 4] that if μ_x and μ_y are respectively supported by m and n dirac masses, then there exists a barycenter supported by up to $m + n + 1$ dirac masses. Hence, the approximation of the barycenter from a set of $\kappa < \min(m, n)$ dirac masses seems to be a bad choice at first sight. Notice however that, as pointed out in [18], a barycenter supported by a low dimensional space allows the approximate distance to be more robust to data outliers. Next, contrary to entropic regularization of OT, the obtained approximate transport map is here sparse, which allows efficient storage and can be directly used for interpolation purposes without any complex post-processing like sharpening.

Finally, it is worth noting that the approximation depends on the regularity of the data. As underlined in [1], when computing the barycenter of a set of densities μ_i in the continuous case, if one of the input data μ_i is absolutely continuous with respect to the Lebesgue measure, so does the barycenter. Such observation gives an interesting insight of the experiments presented in this paper: when at least one of the two data μ_x or μ_y is smooth, then the obtained approximate Wasserstein distance \tilde{W}_p can be very close to the true one W_p for small values of κ . Similar behaviour can be observed with semi-discrete optimal transport models where voronoi cells act like barycenters [27]. This point is illustrated in Figure 2 with the comparison of relative errors between true and approximate distances obtained with increasing values of κ for different scenarios involving “random” or “smooth” 2D data of dimensions $m = n = 10^4$. When at least one “regular” data is involved, then a small relative error $(\tilde{W}_p - W_p)/W_p < 10^{-3}$ (i.e. 0.1%) is observed for $\kappa/n = 0.015$.




Figure 2: Accuracy of the approximate Wasserstein distance estimated by solving the transshipment problem (12) for increasing values of κ . Scenarios involving either smooth and/or random data of dimension $m = n = 10^5$ have been considered.

2.3 Transport refinement and Multi-scale approach

As illustrated in Figure 2, the barycenter approach is not sufficient to produce an accurate approximation of the p -Wasserstein distance between any data. First notice that the approximation \tilde{W}_p of the Wasserstein distance can be easily improved using the following result that directly considers the transport matrix between μ_x and μ_y .

Proposition 4 *Let $\mu_z = \sum_{k=1}^{\kappa} w_k^z \delta_{z_k}$ and γ^x (resp. γ^y) be an optimal transport matrix from μ_x (resp. μ_y) to μ_z . Let also $\hat{\gamma}$ be the transport matrix between μ_x and μ_y defined as $\hat{\gamma} = \gamma^x D(\gamma^y)^\top$, with the rescaling $D^{-1} = \text{diag}(w^z)$ given by the weights $w_k^z = \sum_i \gamma_{ik}^x = \sum_j \gamma_{jk}^y$. Then the following relation holds*

$$\left(\sum_{ij} \hat{\gamma}_{ij} \|x_i - y_j\|^p \right)^{1/p} \leq W_p(\mu_x, \mu_z) + W_p(\mu_y, \mu_z). \quad (17)$$

Proof. The matrix $\hat{\gamma}_{ij} = \sum_k \gamma_{ik}^x \gamma_{jk}^y / w_k^z$ is an admissible transport matrix between μ_x and μ_y , since $\sum_j \hat{\gamma}_{ij} = \sum_k \gamma_{ik}^x = w_i^x$ and $\sum_i \hat{\gamma}_{ij} = \sum_k \gamma_{jk}^y = w_j^y$. Following the proof of Minkowski's inequality:

$$\begin{aligned} \sum_{ij} \hat{\gamma}_{ij} \|x_i - y_j\|^p &= \sum_{ijk} \gamma_{ik}^x \gamma_{jk}^y / w_k^z \|x_i - y_j\|^{p-1} \|x_i - y_j\| \\ &\leq \sum_{ijk} (\gamma_{ik}^x \gamma_{jk}^y / w_k^z)^{(p-1)/p+1/p} \|x_i - y_j\|^{p-1} (\|x_i - z_k\| + \|y_j - z_k\|) \\ &\leq \left(\sum_{ijk} \gamma_{ik}^x \gamma_{jk}^y / w_k^z \|x_i - y_j\|^p \right)^{(p-1)/p} \left(\left(\sum_{ijk} \gamma_{ik}^x \gamma_{jk}^y / w_k^z \|x_i - z_k\|^p \right)^{1/p} + \left(\sum_{ijk} \gamma_{ik}^x \gamma_{jk}^y / w_k^z \|y_j - z_k\|^p \right)^{1/p} \right) \\ &\leq \left(\sum_{ij} \hat{\gamma}_{ij} \|x_i - y_j\|^p \right)^{(p-1)/p} \left(\left(\sum_{ik} \gamma_{ik}^x \|x_i - z_k\|^p \right)^{1/p} + \left(\sum_{jk} \gamma_{jk}^y \|y_j - z_k\|^p \right)^{1/p} \right), \end{aligned}$$

then gives (17). Notice that the proof of the triangle inequality [13] is here also obtained, since $W_p(\mu_x, \mu_y) \leq (\sum_{ij} \hat{\gamma}_{ij} \|x_i - y_j\|^p)^{1/p}$. \square

From this proposition, the approximation $\hat{W}_p(\mu_x, \mu_y)$ corresponding to a barycenter $\mu_{1/2} = \sum_{k=1}^{\kappa} w_k^z \delta_{z_k}$ solution of (12) with transport matrices γ^x and γ^y is defined as

$$W_p(\mu_x, \mu_y) \leq \hat{W}_p(\mu_x, \mu_y) = \langle \gamma^x \text{diag}(w^z)^{-1} (\gamma^y)^\top, c^{xy} \rangle^{1/p} \leq \tilde{W}_p(\mu_x, \mu_y). \quad (18)$$

For matching or interpolation purposes, $\hat{\gamma} = \gamma^x (\text{diag}(w^z))^{-1} (\gamma^y)^\top$ gives a sparse approximation of the optimal transport matrix. However, as illustrated in Figure 3, such approach maps all locations x_i and y_j that transit by z_k , resulting in a poor block transport matrix for small values of κ .




Figure 3: Illustration of the block transport matrices $\hat{\gamma}$ estimated from data μ_x and μ_y through transshipment for increasing values of κ .

In order to get a sparser and more accurate approximation of the optimal transport matrix and associated distance, it is necessary to “untie” the links between locations passing through z_k . To do so, a solution is to consider κ optimal transportation sub-problems, by refining the transport of the mass transshipped through z_k . The barycenter approach then acts as a clustering and for each intermediate location z_k , the p -Wasserstein distance between the following partial discrete densities is computed:

$$\mu_x^k = \sum_{i=1}^m \gamma_{ik}^x \delta_{x_i} \quad \mu_y^k = \sum_{j=1}^n \gamma_{jk}^y \delta_{y_j}, \quad (19)$$

where the number of active dimensions are expected to be reduced: $m_k = \#\{\gamma_{ik}^x > 0\} < m$ and $n_k = \#\{\gamma_{jk}^y > 0\} < n$. The multi-scale approach is then performed as follows. The Wasserstein distance $W_p(\mu_x^k, \mu_y^k)$ is estimated exactly with network simplex [11] if the sum of m_k and n_k is small enough. Otherwise the barycenter approach is recursively applied to the subproblem. From numerical experiments, the threshold $n_k + m_k < 2000$ has been chosen to reach the best compromise between numerical accuracy and computational cost. The whole process is illustrated in Figure 4 and detailed in Algorithm 3.




Figure 4: Illustration of transportation obtained from refinement of transshipment with $\kappa = 4$ and $\kappa = 8$ intermediate locations. edges indicate there is a mass transport (i.e. $\gamma_{ij} > 0$) between locations x_i (in blue) and y_j (in red). Optimal Transportation is here recovered when refining the transshipment solution obtained with $\kappa = 8$.

Algorithm 3 Multi-scale estimation of approximate p -Wasserstein distance $\hat{W}_p^p(\mu_x, \mu_y)$

```

1: procedure APPROXWP( $x, w^x, y, w^y, p, \kappa$ )
2:    $(\tilde{W}, \gamma^x, \gamma^y) = \text{BARWP}(x, w^x, y, w^y, p, \kappa)$  ▷ Clustering with Algo. 2
3:   Initialize  $\hat{W} = 0, \hat{\gamma} = \mathbf{0}_{m \times n}$ 
4:   for  $k = 1$  to  $\kappa$  do ▷ Can be done in parallel
5:     Set  $\mathcal{I}^k = \{i \mid \gamma_{ik}^x > 0\}, \mathcal{J}^k = \{j \mid \gamma_{jk}^y > 0\}$ 
6:     Set  $m_k = \#\mathcal{I}^k, n_k = \#\mathcal{J}^k$ 
7:     Set  $x^k = x_{\mathcal{I}^k}, y^k = y_{\mathcal{J}^k}, w^{x^k} = \gamma_{\mathcal{I}^k}^x, w^{y^k} = \gamma_{\mathcal{J}^k}^y$ 
8:     if  $m_k + n_k < T$  then
9:        $(W^k, \gamma^k) = \text{WP}(x^k, w^{x^k}, y^k, w^{y^k}, p)$  ▷ Compute exact  $W_p^p$  with Algo. 1
10:    else
11:       $(W^k, \gamma^k) = \text{APPROXWP}(x^k, w^{x^k}, y^k, w^{y^k}, p, \kappa)$  ▷ Approximation of  $W_p^p$ 
12:       $\hat{W} = \hat{W} + W^k$ 
13:       $\hat{\gamma}_{\mathcal{I}^k \mathcal{J}^k} = \gamma^k$ 
14:  return  $\hat{W}, \hat{\gamma}$ 

```

3 Experiments

True distance 2–Wasserstein distances (1) are here compared with the approximated ones computed with the multi-scale procedure of Algorithm 3. Exact and approximate distances are respectively obtained with the C++ network simplex implementation of [11] based on the graph library LEMON [22] and the proposed transshipment extension¹. The experiments have been realized on a standard Macbook with a processor Intel Core i7 2,2 GHz and 16 Go of RAM.

Accuracy In order to study the performance of the proposed approximation \hat{W}_p , the 32×32 and 64×64 and 128×128 images of the Benchmark [40] have been considered. For each image size, this data set contains 10 classes of different densities, and each class contains 10 images. The exact and approximate distances have been computed between all (i.e. ≈ 5000) possibles pairs of images. This has been done for different values of κ and a threshold of $T = 2000$ in Algorithm 3. For each experiment, the mean and median relative errors between approximate and exact methods are computed and presented in Table 1.

	$\kappa = 4$		$\kappa = 16$	
	Mean	Median	Mean	Median
$n = 32 \times 32 = 1024$	2.70%	2.18%	1.61%	0.90%
$n = 64 \times 64 = 4096$	3.56%	2.61%	1.31%	0.81%
$n = 128 \times 128 = 16384$	3.49%	2.34%	1.38%	0.82%

Table 1: Mean and median relative errors between approximate and true EMD on the Benchmark [40] for 32×32 , 64×64 and 128×128 images and different values of κ .

A more detailed presentation is then given in Figure 5, with the mean relative errors over intra and inter classes experiments in the case of 128×128 images. As can be observed, significant errors are obtained when unstructured random data (classes 4 and 10) are involved. In all other cases, the relative errors are very low.

¹The code is available at <https://www.math.u-bordeaux.fr/~npapadakis/GOTMI/codes.php>.





















										
	0.4%	0.4%	0.4%	0.4%	0.3%	0.3%	0.4%	0.5%	0.5%	0.4%
	0.4%	1.3%	0.9%	0.8%	0.7%	0.6%	1.4%	1.2%	1.2%	0.6%
	0.4%	0.9%	1.2%	1.6%	0.6%	0.5%	1.5%	1.1%	1.4%	1.0%
	0.4%	0.8%	1.6%	2.7%	0.8%	0.6%	2.0%	1.1%	1.6%	3.8%
	0.3%	0.7%	0.6%	0.8%	0.4%	0.6%	0.4%	1.0%	0.4%	0.2%
	0.3%	0.6%	0.5%	0.6%	0.6%	0.6%	1.8%	1.1%	1.5%	1.2%
	0.4%	1.4%	1.5%	2.0%	0.4%	1.8%	1.8%	1.5%	1.5%	1.7%
	0.5%	1.2%	1.1%	1.1%	1.0%	1.1%	1.5%	0.7%	1.1%	1.2%
	0.5%	1.2%	1.4%	1.6%	0.4%	1.5%	1.5%	1.1%	2.1%	6.1%
	0.4%	0.6%	1.0%	3.8%	0.2%	1.2%	1.7%	1.2%	6.1%	9.7%

Figure 5: Detailed mean relative errors for intra and inter classes tests on images of size 128×128 .

Computational cost In Figure 6, the running time for computing an approximate 2–Wasserstein distance, with Algorithm 3 and $\kappa = 16$, is compared with the C++ network simplex implementation and its sparse multi-threaded extension proposed in [11]. The running times become asymptotically very interesting with the multi-threaded extension, but due to memory storage, these methods can not handle dimensions n larger than 3.10^4 on the considered computer. As transshipment involves problems of size $n\kappa$, it can be applied to data containing more dirac masses and thus deal with one additional order of magnitude ($n = 2.10^5$). With the proposed full C++ implementation of the transshipment problem, the κ sub-problems are solved successively. The provided Matlab interface calling the C++ code through a simple parfor loop with 4 workers is thus much faster. Optimal transshipment matrices being dense for small values of κ , it is counter-productive to consider sparse optimized implementation in the multi-scale framework.

4 Conclusion

This paper presents an empirical method for approximating Wasserstein distances. It is based on existing concepts used in parallel works [18, 5, 30]. The contribution is to provide an efficient multi-scale implementation able to deal with unstructured point clouds while providing sparse transport matrices. Numerical experiments demonstrate the accuracy of the computed approximate distance, while the involved computational cost are improved with respect to the literature. As a perspective, several intermediary transshipment levels could be considered, in relation to branched transport [7] It would also be of interest to add constraints encouraging an homogeneous repartition of the number of points transiting by each location z_k of the barycenter, or at least a more uniform distribution of the barycenter weights w_k^z . The κ multi-scale sub-problems that can be solved in parallel would have similar dimensions, and theoretical guarantees on the overall running time could be given.




Figure 6: Comparison of running times for computing 2–Wasserstein distances for different values of n : C++ network simplex [11], its multi-threaded extension and the proposed multi-scale approximation (matlab C++ mex), that is the only one being able to handle high scale problems.

Acknowledgements This study has been carried out with financial support from the French State, managed by the French National Research Agency (ANR GOTMI) (ANR-16-CE33-0010-01). The project has also received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 777826.

References

- [1] M. Agueh and G. Carlier. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- [2] J. Altschuler, F. Bach, A. Rudi, and J. Weed. Approximating the Quadratic Transportation Metric in Near-Linear Time. *arXiv preprint arXiv:1810.10046*, 2018.
- [3] J. Altschuler, J. Weed, and P. Rigollet. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. *arXiv e-prints arXiv 1705.09634*, 2017.
- [4] E. Anderes, S. Borgwardt, and J. Miller. Discrete Wasserstein Barycenters: Optimal Transport for Discrete Data. *arXiv preprint arXiv:1507.07218*, July 2015.
- [5] G. Auricchio, F. Bassetti, S. Gualandi, and M. Veneroni. Computing kantorovich-wasserstein distances on d -dimensional histograms using $(d+1)$ -partite graphs. *arXiv preprint arXiv:1805.07416*, 2018.
- [6] R. Becker, A. Karrenbauer, S. Krinninger, and C. Lenzen. Near-Optimal Approximate Shortest Paths and Transshipment in Distributed and Streaming Models. *arXiv preprint arXiv:1607.05127*, 2016.
- [7] M. Bernot, V. Caselles, and J.-M. Morel. *Optimal transportation networks: models and theory*. Springer, 2008.
- [8] J. Blanchet, A. Jambulapati, C. Kent, and A. Sidford. Towards Optimal Running Times for Optimal Transport. *arXiv preprint arXiv:1810.07717*, 2018.

- [9] M. Blondel, V. Seguy, and A. Rolet. Smooth and Sparse Optimal Transport. *arXiv preprint arXiv:1710.06276*, 2017.
- [10] N. Bonneel, J. Rabin, G. Peyré, and H. Pfister. Sliced and radon wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51(1):22–45, 2015.
- [11] N. Bonneel, M. van de Panne, S. Paris, and W. Heidrich. Displacement interpolation using lagrangian mass transport. *ACM Trans. Graph.*, 30(6):158:1–158:12, 2011.
- [12] S. Clatici, E. Chien, and J. Solomon. Stochastic Wasserstein Barycenters. *arXiv preprint arXiv:1802.05757*, 2018.
- [13] P. Clement and W. Desch. An elementary proof of the triangle inequality for the wasserstein metric. *Proceedings of the American Mathematical Society*, 136(1):333–339, 2008.
- [14] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Conference on Neural Information Processing Systems (NIPS’13)*, pages 2292–2300, 2013.
- [15] M. Cuturi and A. Doucet. Fast computation of wasserstein barycenters. In *International Conference on Machine Learning (ICML’14)*, pages 685–693, 2014.
- [16] A. Dessein, N. Papadakis, and J.-L. Rouas. Regularized optimal transport and the rot mover’s distance. *Journal of Machine Learning Research*, 19(15):1–53, 2018.
- [17] S. Ferradans, N. Papadakis, G. Peyré, and J.-F. Aujol. Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences*, 7(3):1853–1882, 2014.
- [18] A. Forrow, J.-C. Hütter, M. Nitzan, P. Rigollet, G. Schiebinger, and J. Weed. Statistical Optimal Transport via Factored Couplings. *arXiv preprint arXiv:1806.07348*, 2018.
- [19] C. Gottschlich and D. Schuhmacher. The shortlist method for fast computation of the earth mover’s distance and finding optimal solutions to transportation problems. *PLoS ONE*, 9(10):1–10, 10 2014.
- [20] T. L. Gouic and J.-M. Loubes. Existence and consistency of wasserstein barycenters. *arXiv preprint arXiv:1506.04153*, 2015.
- [21] T. Hurtut, Y. Gousseau, and F. Schmitt. Adaptive image retrieval based on the spatial organization of colors. *Computer Vision and Image Understanding*, 112(2):101–113, 2008.
- [22] LEMON. Library for efficient modeling and optimization in networks. <http://lemon.cs.elte.hu/trac/lemon>, 2010.
- [23] W. Li, E. K. Ryu, S. Osher, W. Yin, and W. Gangbo. A parallel method for earth mover’s distance. *Journal of Scientific Computing*, 75(1):182–197, 2018.
- [24] J. Liu, W. Yin, W. Li, and Y. Tin Chow. Multilevel Optimal Transport: a Fast Approximation of Wasserstein-1 distances. *arXiv preprint arXiv:1810.00118*, 2018.
- [25] D. G. Luenberger and Y. Ye. *Linear and Nonlinear Programming*. Springer Publishing Company, Incorporated, 2015.
- [26] R. J. McCann. A convexity principle for interacting gases. *Advances in Mathematics*, 128(1):153 – 179, 1997.
- [27] Q. Mérigot. A multiscale approach to optimal transport. In *Computer Graphics Forum*, volume 30, pages 1583–1592. Wiley Online Library, 2011.

- [28] A. M. Oberman and Y. Ruan. An efficient linear programming method for Optimal Transportation. *arXiv preprint arXiv:1509.03668*, 2015.
- [29] J. B. Orlin, S. A. Plotkin, and É. Tardos. Polynomial dual network simplex algorithms. *Mathematical Programming*, 60(1):255–276, Jun 1993.
- [30] F.-P. Paty and M. Cuturi. Subspace Robust Wasserstein Distances. *arXiv preprint arXiv:1901.08949*, 2019.
- [31] O. Pele and M. Werman. A linear time histogram metric for improved sift matching. In *European Conference on Computer Vision (ECCV’08)*, pages 495–508, 2008.
- [32] O. Pele and M. Werman. Fast and robust earth mover’s distances. In *IEEE International Conference on Computer Vision (ICCV’09)*, pages 460–467, 2009.
- [33] G. Peyré and M. Cuturi. Computational Optimal Transport. *arXiv preprint arXiv:1803.00567*, 2018.
- [34] J. Rabin, G. Peyré, J. Delon, and M. Bernot. Wasserstein barycenter and its application to texture mixing. In *Scale Space and Variational Methods in Computer Vision*, pages 435–446, 2012.
- [35] Y. Rubner, C. Tomasi, and L. Guibas. A metric for distributions with applications to image databases. In *IEEE Int. Conference on Computer Vision (ICCV’98)*, pages 59–66, 1998.
- [36] F. Santambrogio. Optimal transport for applied mathematicians. *Progress in Nonlinear Differential Equations and their applications*, 87, 2015.
- [37] M. A. Schmitz, M. Heitz, N. Bonneel, F. M. Ngolè Mboula, D. Coeurjolly, M. Cuturi, G. Peyré, and J.-L. Starck. Wasserstein Dictionary Learning: Optimal Transport-based unsupervised non-linear dictionary learning. *arXiv preprint arXiv:1708.01955*, 2017.
- [38] B. Schmitzer. A sparse multiscale algorithm for dense optimal transport. *Journal of Mathematical Imaging and Vision*, 56(2):238–259, 2016.
- [39] B. Schmitzer. Stabilized Sparse Scaling Algorithms for Entropy Regularized Transport Problems. *arXiv preprint arXiv:1610.06519*, 2016.
- [40] J. Schrieber, D. Schuhmacher, and C. Gottschlich. DOTmark - A Benchmark for Discrete Optimal Transport. *arXiv preprint arXiv:1610.03368*, 2016.
- [41] V. Seguy, B. Bhushan Damodaran, R. Flamary, N. Courty, A. Rolet, and M. Blondel. Large-Scale Optimal Transport and Mapping Estimation. *arXiv preprint arXiv:1711.02283*, 2017.
- [42] J. Solomon, F. de Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. In *ACM Transactions on Graphics (SIGGRAPH’15)*, 2015.
- [43] A. Thibault, L. Chizat, C. Dossal, and N. Papadakis. Overrelaxed Sinkhorn-Knopp Algorithm for Regularized Optimal Transport. *arXiv preprint arXiv:1711.01851*, 2017.
- [44] P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109(3):475–494, 2001.
- [45] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Bus. Media, 2008.
- [46] J. Ye, P. Wu, J. Z. Wang, and J. Li. Fast discrete distribution clustering using wasserstein barycenter with sparse support. *IEEE Trans. on Signal Processing*, 65(9):2317–2332, 2017.