

---

# QUANTIZATION BOUNDS FOR WASSERSTEIN METRICS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

The Wasserstein metric is becoming increasingly important in many machine learning applications such as generative modeling, image retrieval and domain adaptation. Despite its appeal, it is often too costly to compute. This has motivated approximation methods like entropy-regularized optimal transport, downsampling, and subsampling, which trade accuracy for computational efficiency. In this paper, we consider the challenge of computing efficient approximations to the Wasserstein metric that also serve as strict upper or lower bounds, as these are essential components of branch-and-bound, A\* path finding, and heuristic search techniques in tasks such as trajectory inference, alignment, and clustering. Focusing on discrete measures on regular grids, our approach involves formulating and exactly solving a Kantorovich problem on a coarse grid using a quantized measure with a tailored cost matrix, followed by an upscaling and correction stage. This is done either in the primal or dual space to obtain valid upper and lower bounds on the Wasserstein metric of the full-resolution inputs. We evaluate our methods on the DOTmark optimal transport images benchmark as well as alignment tasks on volumetric dataset of macromolecules, demonstrating a 10×–100× speedup compared to entropy-regularized OT while keeping the approximation error well below 5% in 2D, and 30% width bounding regions in 3D.

## 1 INTRODUCTION

The Wasserstein metric is a basic tool in machine learning with broad applications in fields as diverse as computer vision, natural language processing, domain adaptation, and computational biology (Arjovsky et al., 2017; Kusner et al., 2015; Courty et al., 2017; Schiebinger et al., 2019). However, a limiting factor to its adoption is computational cost. For example, consider the calculation of the Wasserstein metric between  $N \times N$  images that we treat as discrete measures on a regular grid. Each image has  $N^2$  pixels, so computing the Wasserstein metric involves solving a linear program with  $N^2 \times N^2$  variables and  $\Theta(N^2)$  constraints. This is typically done using a network simplex algorithm whose worst-case runtime in this case is  $O(N^6 \log N)$  (Peyré & Cuturi, 2019). For 3D signals, the computational cost is even worse at  $O(N^9 \log N)$ .

Our work is motivated by recent applications that use Wasserstein metrics as targets for minimization in diverse tasks such as trajectory inference, alignment, and clustering (Schiebinger et al., 2019; Banerjee et al., 2025; Riahi et al., 2023; Singer & Yang, 2024; Rao et al., 2020; Papayiannis et al., 2021). These optimization problems could be solved more efficiently using branch-and-bound, A\* path finding, or other heuristic search techniques. However, this requires fast methods to compute *bounds* on the exact Wasserstein metric. Motivated by this challenge, the goal of this work is to develop fast approximations to the  $p$ -Wasserstein metric that also bound it from above or from below. Our particular focus is on image and volumetric data, but the approach could be generalized to other domains such as point clouds and graphs.

**Our contribution** We propose four methods for computing fast bounds on the Wasserstein metric between measures defined on regular grids: *weighted-cost upper bound*, *min-cost lower bound*, *primal upscaling upper bound*, and *dual upscaling lower bound*. Each method solves an exact optimal transport problem on a quantized (lower resolution) grid with a tailored ground cost and lifts the result to the original resolution. As a baseline for comparisons, we explain how entropy-regularized optimal transport can be used to produce bounds. See Section 3 and the supplementary for the description of the methods and Section 4 for a complexity analysis. For all methods, we developed efficient

JAX-based GPU implementations which we release as free software. To compare the accuracy and running time of the methods, we tested them on 2D images from the DOTmark dataset and on 3D volumetric data from the electron microscopy data bank (EMDB). The datasets and benchmark results are presented in Section 5 and in the appendix.

## 1.1 RELATED WORK

In recent years, many authors have developed fast approximations to the Wasserstein metric. A non-exhaustive list includes entropy-regularized optimal transport (Cuturi, 2013; Altschuler et al., 2018), convolutional Wasserstein (Solomon et al., 2015), sliced-Wasserstein (Deshpande et al., 2018; Kolouri et al., 2019), and linear approximations (Shirdhonkar & Jacobs, 2008; Moosmüller & Cloninger, 2023; Craig & Yu, 2025). The quantization-based approach of Beugnot et al. (2021) is similar in spirit to ours, but is focused on continuous measures and only guarantees the quality of approximation in expectation. Other methods that are related to our approach use hierarchical refinement to obtain a solution to the Kantorovich problem (Mérigot, 2011; Gerber & Maggioni, 2017; Feydy et al., 2021; Chen et al., 2022). However, unlike our approach, these methods do not give bounds on the exact Wasserstein metric.

Some works dedicated to accelerating similarity search and retrieval tasks, derive computationally efficient lower bounds for the Wasserstein metric. Yang et al. (2022) discusses several linear-time lower bounds for pruning search spaces, providing loose bounds with no tightness guarantees. Tree-based embeddings like Indyk & Thaper (2003); Backurs et al. (2020) provide bounds by embedding the metric into hierarchical structures, though at the cost of multiplicative  $O(\log n)$  approximation factor. Another class of bounds arises from relaxing the transport constraints, as seen in the *relaxed word mover’s distance* (Kusner et al., 2015; Atasu & Mittelholzer, 2019). While all these methods successfully provide valid lower bounds essential for large-scale retrieval, they prioritize computational speed over the tightness of the bounds. See Montesuma et al. (2025) for a recent survey of computation and approximation methods.

## 2 BACKGROUND

**Notation** We denote the non-negative real numbers by  $\mathbb{R}_+$  and the set of integers  $\{1, \dots, n\}$  by  $[n]$ . The tensor product is denoted by  $\otimes$  whereas pointwise multiplication and division are denoted by  $\odot$  and  $\oslash$ , respectively. The  $L^p$  norm of a vector is  $\|\cdot\|_p$ . The all-ones column vector is  $\mathbf{1}_n \in \mathbb{R}^n$ . The standard vector inner product is denoted by  $\langle \cdot, \cdot \rangle$  and we use the same notation for the inner products of matrices. The support of a matrix  $A \in \mathbb{R}^{n \times m}$  is the set of indices of nonzero elements  $\text{supp } A = \{(i, j) \in [n] \times [m] \mid A_{i,j} \neq 0\}$ . The cardinality of a set  $S$  is denoted by  $\#S$ .  $\delta_x$  is the Dirac delta function at point  $x$ . Complete list of notations used in the paper is provided in Table 7.

### 2.1 OPTIMAL TRANSPORT

In the following, we give a quick review of basic concepts from optimal transport and refer the reader to Peyré & Cuturi (2019) for a more thorough introduction. Consider two discrete probability measures  $\mu, \nu$  with point masses at  $\mathcal{X} = \{x_1, \dots, x_n\}$  and  $\mathcal{Y} = \{y_1, \dots, y_m\}$  respectively. We can express the measures as a sum of Dirac delta functions,

$$\mu = \sum_{i=1}^n \mu_i \delta_{x_i}, \quad \nu = \sum_{j=1}^m \nu_j \delta_{y_j}. \quad (1)$$

We identify the measures with their non-negative coefficient vectors  $\boldsymbol{\mu} \in \mathbb{R}_+^n, \boldsymbol{\nu} \in \mathbb{R}_+^m$ . Since  $\mu$  and  $\nu$  are probability measures, we must have  $\boldsymbol{\mu} \in \Sigma_n$  and  $\boldsymbol{\nu} \in \Sigma_m$  where

$$\Sigma_N := \{(p_1, \dots, p_N) \in \mathbb{R}_+^N : p_1 + \dots + p_N = 1\} \quad (2)$$

is the probability simplex. The set of *coupling* matrices between  $\mu$  and  $\nu$  is

$$\Pi(\boldsymbol{\mu}, \boldsymbol{\nu}) := \{\pi \in \mathbb{R}_+^{n \times m} \mid \pi \mathbf{1}_m = \boldsymbol{\mu}, \pi^\top \mathbf{1}_n = \boldsymbol{\nu}\}. \quad (3)$$

Each coupling can be viewed as a transport plan between  $\mu$  and  $\nu$  where  $\pi_{ij}$  is the amount of mass transported from  $x_i$  to  $y_j$ . The marginal constraints  $\pi \mathbf{1}_m = \boldsymbol{\mu}$  mean that the entire source measure

108  $\mu$  is transported, whereas the constraints  $\pi^\top \mathbf{1}_n = \nu$  mean that this transport results in the target  
 109 measure  $\nu$ . In particular, the set  $\Pi(\mu, \nu)$  always contains the trivial coupling  $\pi_\otimes := \mu \otimes \nu$  that  
 110 distributes every point mass in  $\mathcal{X}$  to all point masses in  $\mathcal{Y}$  proportionately to  $\nu$ . Let  $C \in \mathbb{R}_+^{n \times m}$   
 111 be a ground-cost matrix, where  $C_{ij}$  represents the cost of transporting a unit mass from  $x_i \in \mathcal{X}$   
 112 to  $y_j \in \mathcal{Y}$ . The Kantorovich problem is the minimization problem that seeks a cost-minimizing  
 113 coupling between  $\mu$  and  $\nu$ ,

$$114 \quad L_C(\mu, \nu) := \min_{\pi \in \Pi(\mu, \nu)} \langle \pi, C \rangle. \quad (4)$$

115 This is a linear optimization problem with linear constraints. It admits a dual program,

$$116 \quad L_C(\mu, \nu) = \max_{(\mathbf{f}, \mathbf{g}) \in \mathcal{R}(C)} \langle \mathbf{f}, \mu \rangle + \langle \mathbf{g}, \nu \rangle, \quad (5)$$

117 where

$$118 \quad \mathcal{R}(C) := \{(\mathbf{f}, \mathbf{g}) \in \mathbb{R}^n \times \mathbb{R}^m \mid \forall i, j : f_i + g_j \leq C_{ij}\} \quad (6)$$

119 is the set of admissible dual potentials, also known as Kantorovich potentials. Given a distance metric  
 120  $\rho : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$ , for any  $p \geq 1$  the *Wasserstein- $p$*  metric is a metric over the space of discrete  
 121 probability measures with point masses at  $\mathcal{X} = \{x_1, \dots, x_n\}$ , defined as  $\mathcal{W}_p(\mu, \nu) := L_C(\mu, \nu)^{\frac{1}{p}}$   
 122 where  $C_{ij} = \rho(x_i, y_j)^p$ .

## 123 2.2 MEASURE COARSENING

124 Consider a  $d$ -dimensional regular grid  $\mathcal{X} = [N]^d$  and suppose that  $N = n\kappa$  for some integers  $n, \kappa$ .  
 125 One may subdivide the grid along the axes into a set  $\mathbb{X}(\kappa) := \{X_1, \dots, X_{n\kappa}\}$  of non-overlapping  
 126 hypercubes of cardinality  $\kappa^d$  that cover the entire grid  $\mathcal{X}$ , such that  $N = n\kappa$ . We define the coarse  
 127 grid  $\tilde{\mathcal{X}} := \{\tilde{x}_1, \dots, \tilde{x}_{n\kappa}\}$  as the set of all hypercube centers, with  $\tilde{x}_k = \text{mean}(X_k)$ . Given a discrete  
 128 measure  $\mu$  over the grid  $\mathcal{X}$ , it induces a coarsened discrete measure by aggregating the mass of  
 129 each hypercube and placing it at its center point. This coarsening corresponds to the `SumPool` and  
 130 `AvgPool` operations on the measure and the coordinates, respectively, with size and stride  $\kappa$

$$131 \quad \text{SumPool}(\mu; \mathbb{X}(\kappa))_k := \sum_{x \in X_k} \mu(x) \quad \text{AvgPool}(\mathcal{X}; \mathbb{X}(\kappa))_k := \frac{1}{\#X_k} \sum_{x \in X_k} x. \quad (7)$$

## 132 3 METHODS

133 In this section, we describe several algorithms for computing bounds of the Wasserstein distance  
 134  $\mathcal{W}_p(\mu, \nu)$  on a regular grid. First, in Section 3.1 we explain how entropic optimal transport can be  
 135 used to obtain strict bounds. In the following subsections we introduce several novel algorithms  
 136 for fast bounds on the Wasserstein metric that are based on quantization (or downscaling) of the  
 137 inputs onto a coarse grid and sum-pooling the measures. These four methods have the following  
 138 structure: First, they build a particular cost matrix for the coarse grid that takes the original measures  
 139 into account. This is followed by a correction stage that upscales the solution on the coarse grid to  
 140 a solution on the original grid and corrects the marginals using iterative proportional-fitting. The  
 141 upscaling and correction stage is done separately in the primal and dual spaces, to obtain upper and  
 142 lower bounds (respectively). Finally, to guarantee the correctness of the bounds without relying on the  
 143 convergence of the proportional fitting procedure, we introduce an additional total variation correction  
 144 term. The four quantization-based bounds, assuming the ground cost is Lipschitz continuous, are  
 145 equal to the exact Wasserstein metric up to an additive  $O(\kappa\sqrt{d})$  term, as detailed in [Appendix D](#).

### 146 3.1 BOUNDS BASED ON ENTROPY REGULARIZED OT

147 Entropy regularized optimal transport, also known as Sinkhorn distance ([Cuturi, 2013](#)), adds an  
 148 entropy term  $\mathcal{H}(\pi) = -\sum_{x,y} \pi(x,y) \log \pi(x,y)$  to the primal:

$$149 \quad L_C^\varepsilon(\mu, \nu) := \min_{\pi \in \Pi(\mu, \nu)} \langle \pi, C \rangle - \varepsilon \mathcal{H}(\pi). \quad (8)$$

150 This makes the problem strongly convex and solvable using Sinkhorn iterations ([Knopp & Sinkhorn, 1967](#)).  
 151 In this section we will show how computing entropy regularized OT can be used to construct  
 152 upper and lower bounds on the exact Wasserstein distance.

162 **Lower bound** The dual form of the Sinkhorn distance is

$$163 L_C^\varepsilon(\mu, \nu) := \max_{\mathbf{f} \in \mathbb{R}^n, \mathbf{g} \in \mathbb{R}^m} \langle \mathbf{f}, \boldsymbol{\mu} \rangle + \langle \mathbf{g}, \boldsymbol{\nu} \rangle - \varepsilon \langle e^{\mathbf{f}/\varepsilon}, K e^{\mathbf{g}/\varepsilon} \rangle \quad (9)$$

164 where  $K_{ij} := e^{-C_{ij}/\varepsilon}$  is the Gibbs kernel. The algorithmic solution, defined by the use of a finite  
165 number of iterations  $t$  that achieves some stopping criteria, is known to satisfy a lower bound.  
166 Summarized in the following proposition:

167 **Proposition 3.1.** *Let  $\hat{\mathbf{f}}_\varepsilon^{(t)}, \hat{\mathbf{g}}_\varepsilon^{(t)}$  be the iterations of the Sinkhorn distance algorithm in step  $t \in \mathbb{N}$ .*

$$168 \langle \hat{\mathbf{f}}_\varepsilon^{(t)}, \boldsymbol{\mu} \rangle + \langle \hat{\mathbf{g}}_\varepsilon^{(t)}, \boldsymbol{\nu} \rangle \leq L_C(\mu, \nu) \quad (10)$$

169 as soon as  $t \geq 1$ .

170 This follows directly from [Peyré & Cuturi \(2019, Propositions 4.5,4.8\)](#), so we can define

$$171 \underline{\mathcal{W}}_{p,\varepsilon}(\mu, \nu) := (\langle \hat{\mathbf{f}}_\varepsilon^{(t)}, \boldsymbol{\mu} \rangle + \langle \hat{\mathbf{g}}_\varepsilon^{(t)}, \boldsymbol{\nu} \rangle)^{\frac{1}{p}} \leq \mathcal{W}_p(\mu, \nu) \text{ for all } p \geq 1.$$

172 **Upper bound** Consider  $d$ -dimensional regular grids with side length  $N \in \mathbb{N}$ ,  $\mathcal{X} = \mathcal{Y} = [N]^d$ , with  
173 discrete measures  $\mu, \nu \in \Sigma_{N^d}$ . Although the converged regularized optimal coupling

$$174 \pi_\varepsilon^* = \arg \min_{\pi \in \Pi(\mu, \nu)} \langle \pi, C \rangle - \varepsilon \mathcal{H}(\pi) \quad (11)$$

175 defines an upper bound on the optimal transport  $\langle \pi_\varepsilon^*, C \rangle \geq L_C(\mu, \nu)$ , the algorithmic solution  $\hat{\pi}_\varepsilon^{(t)}$   
176 does not, since the marginals  $\hat{\boldsymbol{\mu}}_\varepsilon^{(t)} = \hat{\pi}_\varepsilon^{(t)} \mathbf{1}_N$ ,  $\hat{\boldsymbol{\nu}}_\varepsilon^{(t)} = (\hat{\pi}_\varepsilon^{(t)})^\top \mathbf{1}_N$  do not identify with the couplings  
177  $\mu, \nu$ . We bound the effect of this difference using the weighted total variation. For some  $x_0 \in \mathcal{X}$ ,  
178 using distance weights  $\mathbf{w} = \{\rho(x_0, x)^p\}_{x \in \mathcal{X}}$  the Wasserstein- $p$  distance is controlled by *weighted*  
179 total variation (TV) ([Villani, 2009](#)),

$$180 \mathcal{TV}_p^{\mathbf{w}}(\mu, \nu) := 2^{1-\frac{1}{p}} \langle \mathbf{w}, |\boldsymbol{\mu} - \boldsymbol{\nu}| \rangle^{\frac{1}{p}} \geq \mathcal{W}_p(\mu, \nu). \quad (12)$$

181 Here  $|\cdot|$  is the element-wise absolute value. We define the TV-corrected upper bound by

$$182 \overline{\mathcal{W}}_{p,\varepsilon}(\mu, \nu) := \langle \hat{\pi}_\varepsilon^{(t)}, C \rangle^{\frac{1}{p}} + \Delta_{\hat{\boldsymbol{\mu}}_\varepsilon^{(t)}} + \Delta_{\hat{\boldsymbol{\nu}}_\varepsilon^{(t)}}, \quad (13)$$

183 where  $\Delta_{\hat{\boldsymbol{\mu}}_\varepsilon^{(t)}} = \mathcal{TV}_p^{\mathbf{w}}(\hat{\boldsymbol{\mu}}_\varepsilon^{(t)}, \mu)$ ,  $\Delta_{\hat{\boldsymbol{\nu}}_\varepsilon^{(t)}} = \mathcal{TV}_p^{\mathbf{w}}(\nu, \hat{\boldsymbol{\nu}}_\varepsilon^{(t)})$  are the marginal corrections with weights  
184  $\mathbf{w} = \{\rho(\bar{x}, x_i)\}_i$  taken from the center of the measure  $\bar{x} = \text{mean}(\mathcal{X})$ . Using the triangle inequality  
185 for the Wasserstein metric, we can write

186 **Lemma 3.2.** *Let  $\mu, \hat{\mu}, \nu, \hat{\nu} \in \Sigma_{N^d}$  discrete measures on  $\mathcal{X} = [N]^d$ , and  $\hat{\pi} \in \Pi(\hat{\mu}, \hat{\nu})$  is a coupling  
187 between  $\hat{\mu}$  and  $\hat{\nu}$ . For  $p \geq 1$ ,  $\mathcal{W}_p(\mu, \nu) \leq \langle \hat{\pi}, C \rangle^{\frac{1}{p}} + \Delta_{\hat{\mu}} + \Delta_{\hat{\nu}}$ .*

188 Combining this with (13) shows that  $\mathcal{W}_p(\mu, \nu) \leq \overline{\mathcal{W}}_{p,\varepsilon}(\mu, \nu)$ . See illustration in [Figure 1](#).

189 **Proposition 3.3.** *Let  $\mu, \hat{\mu}, \nu, \hat{\nu} \in \Sigma_n$  be discrete measures in  $\mathcal{X}$  and  $\xi > 0$ , satisfying the convergence  
190 criteria  $\|\hat{\mu} - \mu\|_1 + \|\nu - \hat{\nu}\|_1 < \xi$ . The sum of the marginal corrections is bounded,*

$$191 \Delta_{\hat{\mu}} + \Delta_{\hat{\nu}} < 2^{2-\frac{2}{p}} \xi^{\frac{1}{p}} r \quad (14)$$

192 where  $r := \max_{x \in \mathcal{X}} \{\rho(\bar{x}, x)\}$  is radius of  $\mathcal{X}$ .

193 Proofs for [Lemma 3.2](#) and [Proposition 3.3](#) are provided in [Appendix C.1](#).

### 204 3.2 WEIGHTED-COST UPPER BOUND

205 In this subsection we consider an approach to bound Wasserstein distance by downscaling the grid  $\mathcal{X}$   
206 to  $\tilde{\mathcal{X}} = \tilde{\mathcal{Y}}$  and the measures to  $\tilde{\mu}, \tilde{\nu} \in \Sigma_{n^d}$  using regular hypercubes as described in [Section 2.2](#). We  
207 define the marginally weighted coarse cost

$$208 \bar{C}_{k\ell} := \frac{1}{\mu(X_k)\nu(Y_\ell)} \sum_{x \in X_k, y \in Y_\ell} \rho(x, y)^p \mu(x)\nu(y). \quad (15)$$

209 It follows that  $\bar{C} = \text{SumPool}(C \odot \pi_\otimes; \mathbb{X}(\kappa)^2) \otimes \text{SumPool}(\pi_\otimes; \mathbb{X}(\kappa)^2)$ , which can be used  
210 efficiently for small enough grids. Then we compute the optimal coupling for the marginally  
211 weighted coarse cost using network simplex solver ([Bonnel et al., 2011](#)), defining an upper bound

$$212 \overline{\mathcal{W}}_p^\otimes := L_{\bar{C}}(\tilde{\mu}, \tilde{\nu})^{\frac{1}{p}}. \quad (16)$$

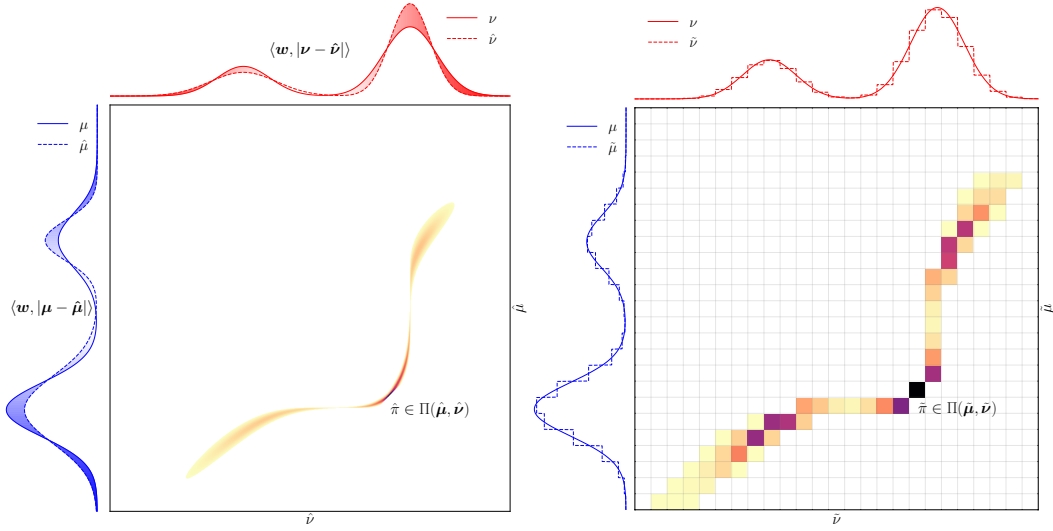


Figure 1: Optimal transport bounds visualization. Left: Illustration of optimal transport between discrete probability measures, approximating the marginal measures. Right: Illustration of optimal transport between quantized discrete probability measures, as approximation.

**Theorem 3.4.** *The optimal transport loss under the marginally weighted coarse cost  $L_{\tilde{C}}(\tilde{\mu}, \tilde{\nu})$  is an upper bound to the optimal transport loss  $L_C(\mu, \nu)$ , and similarly for the Wasserstein distance,*

$$L_C(\tilde{\mu}, \tilde{\nu})^{\frac{1}{p}} \geq \mathcal{W}_p(\mu, \nu). \quad (17)$$

The proof uses an auxiliary coupling transferring the mass between each pair of sub-regions  $(X_i, Y_j)$  using the trivial coupling, weighted by a coarse coupling, choosing the optimal coarse coupling. A detailed proof is provided in [Appendix C.2](#).

### 3.3 MIN-COST LOWER BOUND

For the same coarsening, we define the locally minimal cost  $C_{k\ell}^{\min} := \min_{x \in X_k, y \in Y_\ell} \rho(x, y)^p$ , and compute the optimal coupling for this coarse cost using a network simplex solver, yielding a lower bound  $\mathcal{W}_p^{\min} := L_{C^{\min}}(\tilde{\mu}, \tilde{\nu})^{\frac{1}{p}}$ . The following theorem is proved in [Appendix C.3](#).

**Theorem 3.5.** *The coarse optimal transport cost set by the locally minimal cost, is a lower bound of the optimal transport.  $L_{C^{\min}}(\tilde{\mu}, \tilde{\nu}) \leq L_C(\mu, \nu)$ .*

### 3.4 PRIMAL UPSCALING UPPER BOUND

In this approach we upscale an optimal coupling for the coarse cost  $\tilde{c}_{k\ell} = \rho(\tilde{x}_k, \tilde{y}_\ell)^p$  computed using a network simplex solver  $\tilde{\pi}^* = \arg \min_{\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})} \langle \tilde{\pi}, \tilde{C} \rangle$  back to the original problem size.

**Up-scaled coupling** A coupling matrix  $\pi$  of dimensions  $n^d \times n^d$  can equivalently be represented as a  $2d$ -tensor of the shape  $n \times n \times \dots \times n$ . We formally define operations for reshaping matrices into tensors and back. Let  $\text{reshape}$  be a cardinality-preserving transformation from  $\mathbf{A}$  to  $\mathbf{B}$  such that  $a_{ij} = b_{u_1, \dots, u_d, v_1, \dots, v_d}$  where  $(u_1, \dots, u_d)$  is a multi-index that corresponds to a matrix row by  $i = 1 + \sum_{k=1}^d u_k n^{k-1}$ . Similarly the multi-index  $(v_1, \dots, v_d)$  corresponds to the column index  $j$ . Reshaping the coarse optimal coupling  $\tilde{\pi}^*$  into a  $2d$ -tensor  $\tilde{\mathbf{P}}^*$ , we up-scale the optimal coupling using a normalized positive-valued kernel  $\mathbf{K}$ , a  $2d$ -tensor representing a hypercube of width  $\kappa$ .  $\hat{\mathbf{P}} := \tilde{\mathbf{P}}^* \otimes \mathbf{K}$ . By reshaping the up-scaled tensor  $\hat{\mathbf{P}}$  into an up-scaled matrix  $\hat{\pi}$ , we obtain the approximate up-scaled coupling. Using a uniform kernel  $\mathbf{K}$  performs nearest-neighbor interpolation.

**Lemma 3.6.** Using a normalized positive-valued kernel  $\mathbf{K}$ , satisfying  $\sum_{t \in [\kappa]^{2d}} \mathbf{K}_t = 1$ , ensures  $\hat{\pi}$  represents a coupling, thus satisfying  $\hat{\pi} \in \mathbb{R}_+^{N^d \times N^d}$  and  $\sum_{i=1}^{N^d} \sum_{j=1}^{N^d} \hat{\pi}_{ij} = 1$ , such that  $\hat{\pi} \in \Pi(\hat{\pi} \mathbf{1}_{N^d}, \hat{\pi}^\top \mathbf{1}_{N^d})$ .

**Iterative proportional fitting** The approximate up-scaled coupling is  $\xi$ -fitted into  $\hat{\pi}_\xi = \text{diag}(\mathbf{a}) \hat{\pi} \text{diag}(\mathbf{b})$  by iterative proportional fitting using Sinkhorn–Knopp algorithm, until the marginals  $\hat{\boldsymbol{\mu}}_\xi = \mathbf{a} \odot (\hat{\pi} \mathbf{b})$ ,  $\hat{\boldsymbol{\nu}}_\xi = \mathbf{b} \odot (\hat{\pi}^\top \mathbf{a})$  are converged to  $\|\hat{\boldsymbol{\mu}}_\xi - \boldsymbol{\mu}\|_1 + \|\boldsymbol{\nu} - \hat{\boldsymbol{\nu}}_\xi\|_1 < \xi$ , where  $\mathbf{a}, \mathbf{b} \in \mathbb{R}_+^{N^d}$  are the vector scale factors, yielding the approximation  $\widehat{\mathcal{W}}_p(\mu, \nu) := \langle \hat{\pi}_\xi, C \rangle^{\frac{1}{p}}$ .

**Upper bound** Using weighted total variation (Equation (12)) we define an upscaling upper bound

$$\overline{\mathcal{W}}_p(\mu, \nu) := \widehat{\mathcal{W}}_p(\mu, \nu) + \Delta_{\hat{\boldsymbol{\mu}}_\xi} + \Delta_{\hat{\boldsymbol{\nu}}_\xi}, \quad (18)$$

**Theorem 3.7** (Upscaling Upper Bound).  $\overline{\mathcal{W}}_p(\mu, \nu)$  is upper bound of the Wasserstein distance,

$$\mathcal{W}_p(\mu, \nu) \leq \overline{\mathcal{W}}_p(\mu, \nu). \quad (19)$$

*Proof of Theorem 3.7.* The up-scaled matrix  $\hat{\pi} \in \Pi(\hat{\pi} \mathbf{1}_{N^d}, \hat{\pi}^\top \mathbf{1}_{N^d})$  is normalized as a coupling, by Lemma 3.6, and  $\hat{\pi}_\xi$  retains this normalization, by Sinkhorn’s theorem.  $\overline{\mathcal{W}}_p(\mu, \nu)$  is an upper bound of the Wasserstein distance, by Lemma 3.2.  $\square$

*Remark 3.8.* Considering  $\mathcal{X} = [N]^d$  with  $L^2$  norm, the radius becomes  $r = \frac{1}{2} d^{\frac{1}{2}} N$  (see Proposition 3.3) such that the weighted total variation correction  $\Delta_{\hat{\boldsymbol{\mu}}} + \Delta_{\hat{\boldsymbol{\nu}}}$  is at most  $2^{1-\frac{2}{p}} d^{\frac{1}{2}} N \xi^{\frac{1}{p}}$ . Thus, the weighted total variation correction is negligible for  $\xi \ll N^{-p}$ , and can be ignored for many practical use cases.

### 3.5 DUAL UPSCALING LOWER BOUND

We construct a lower bound for the Wasserstein distance  $\mathcal{W}_p(\mu, \nu)$  by solving a down-scaled optimal transport problem using coarsened measures. The coarse optimal Kantorovich potentials are then up-scaled using a multi-linear interpolation and improved using a c-transform. Considering the same setting as in Section 3.4, we solve for the optimal potentials of the down-scaled discrete measures

$$(\tilde{\mathbf{f}}^*, \tilde{\mathbf{g}}^*) = \arg \max_{(\tilde{\mathbf{f}}, \tilde{\mathbf{g}}) \in \mathcal{R}(\tilde{C})} \langle \tilde{\mathbf{f}}, \tilde{\boldsymbol{\mu}} \rangle + \langle \tilde{\mathbf{g}}, \tilde{\boldsymbol{\nu}} \rangle \quad (20)$$

and evaluate the dual transport cost at the original scale by upscaling the optimal potentials. Upscaling can be performed by any multivariate interpolation method such as nearest-neighbor, spline, multi-linear and polynomial methods. Using an interpolation function  $R : \mathcal{X} \cup \tilde{\mathcal{X}} \rightarrow \mathbb{R}$ , the up-scaled potential  $\hat{\mathbf{f}}$  is defined by

$$\hat{\mathbf{f}} := \{R_{\tilde{\mathbf{f}}, \tilde{\mathcal{X}}}(x_i)\}_{i \in [N^d]}. \quad (21)$$

An important property of the dual formulation is that for every potential  $\mathbf{f} \in \mathbb{R}^n$  we can easily find a *tight* potential  $\mathbf{f}^c \in \mathbb{R}^m$  such that  $(\mathbf{f}, \mathbf{f}^c) \in \mathcal{R}(C)$ , by  $f_j^c := \min_i C_{ij} - f_i$ . This is known as a *c-transform*. It can be shown that repeating this process once more achieves a tight pair  $(\mathbf{f}^c, \mathbf{f}^{cc}) \in \mathcal{R}(C)$ , where  $f_i^{cc} := \min_j C_{ij} - f_j^c$ . Thus, a lower bound is guaranteed by using the c-transform to generate the potential pair from the up-scaled potential  $\mathbf{f} = \hat{\mathbf{f}}^{cc}$  and  $\mathbf{g} = \hat{\mathbf{f}}^c$ , which yields the upscaling lower bound

$$\underline{\mathcal{W}}_p(\mu, \nu) := (\langle \mathbf{f}, \boldsymbol{\mu} \rangle + \langle \mathbf{g}, \boldsymbol{\nu} \rangle)^{\frac{1}{p}}. \quad (22)$$

Finally, by the admissibility of a c-transformed pair  $(\mathbf{f}, \mathbf{g}) \in \mathcal{R}(C)$  we can write

**Proposition 3.9.** Considering an approximate potential  $\hat{\mathbf{f}} \in \mathbb{R}^{N^d}$ . For  $\mathbf{f} = \hat{\mathbf{f}}^{cc}$  and  $\mathbf{g} = \hat{\mathbf{f}}^c$

$$\langle \mathbf{f}, \boldsymbol{\mu} \rangle + \langle \mathbf{g}, \boldsymbol{\nu} \rangle \leq L_C(\mu, \nu). \quad (23)$$

## 4 COMPUTATIONAL COMPLEXITY ANALYSIS

The quantization-based bounds involve the following steps: computing the Wasserstein metric on the quantized measures, upscaling the dual potentials or couplings to the original scale, and the calculation of weighted total variation correction terms. The latter is calculated in linear time and space, thus negligible w.r.t. the other steps. In the following, we detail the computational gains provided by the proposed methods.

**Downscaled optimal transport** The solution to the Kantorovich problem of the scaled measures can be solved by dedicated linear programming methods, such as the network simplex used in [Bonnel et al. \(2011\)](#), with  $O(n^{3d} \log n)$  time complexity ([Ahuja et al., 1993](#)). By solving only for the optimal transport of the coarse measures, we produce a computational speedup of  $\Theta(\kappa^{3d})$  (up to log factors). The space complexity can also be significantly reduced, since one can avoid storing the full cost matrix of size  $N^d \times N^d$ , by using coarse cost matrices, e.g.  $\bar{C}$ ,  $C^{\min}$ , and  $\tilde{C}$  of size  $n^d \times n^d$ , realizing a memory gain of  $\Theta(\kappa^{2d})$ .

**Upscaled optimal coupling** The optimal coarse coupling  $\tilde{\pi}^*$  is a sparse matrix with at most  $2n^d - 1$  positive entries ([Peyré & Cuturi, 2019](#), Proposition 3.4). Thus, the up-scaled approximate coupling from [Theorem 3.7](#) conserves this sparsity with  $\#\text{supp } \hat{\pi} \leq \kappa^{2d}(2n^d - 1)$ , allowing to calculate the approximate optimal transport

$$\langle \hat{\pi}_\xi, C \rangle = \sum_{(i,j) \in \text{supp } \hat{\pi}} \overbrace{a_i \hat{\pi}_{ij} b_j}^{(\hat{\pi}_\xi)_{ij}} \rho(x_i, y_j)^p. \quad (24)$$

without impacting the total time and space complexity of the coarse optimal transport solution.

$$\frac{\#C}{\#\text{supp } \hat{\pi}} = \frac{N^{2d}}{\kappa^{2d}(2n^d - 1)} = \Theta((N/\kappa)^d). \quad (25)$$

**Lazy c-transform** To reduce the memory requirements of [Equation \(22\)](#), we evaluate the c-transform on-demand (i.e. "lazy") without storing the entire cost matrix  $C$ .

$$g_j \leftarrow \min_i \rho(x_i, y_j)^p - \hat{f}_i \quad (26)$$

$$f_i \leftarrow \min_j \rho(x_i, y_j)^p - g_j \quad (27)$$

Method	Time Complexity	Space Complexity
Exact Wasserstein computation	$\tilde{O}(N^{3d})$	$O(N^{2d})$
Entropic Regularization-Based Bounds ( <a href="#">Lin et al., 2022</a> )	$\tilde{O}(N^{2d}/\varepsilon^2)$	$O(N^{2d})$
Quantization-Based Bounds	$\tilde{O}((N/\kappa)^{3d})$	$O((N/\kappa)^{2d})$

Table 1: Complexity of different bounds in terms of the fine-scale cardinality  $\#\mathcal{X} = N^d = (n\kappa)^d$ . The first row corresponds to the methods in [Section 3.1](#) and the second to [Sections 3.2](#) to [3.5](#).

While the entropic regularization-based bounds enjoy better asymptotic time complexity than the proposed quantization-based bounds, in practice, to achieve comparable accuracy we are required to pick very small  $\varepsilon$  values, eliminating this effect, as shown in [Section 5](#).

## 5 EXPERIMENTS

The methods were implemented in Python and optimized for GPU acceleration using the JAX numerical computing library ([Bradbury et al., 2018](#)). For solving entropy regularized optimal transport we used `sinkhorn` solver from OTT-JAX ([Cuturi et al., 2022](#)) and for solving exact Kantorovich problems we used the network-simplex `emd` solver from Python Optimal Transport (POT) ([Flamary et al., 2021](#)). The benchmarks were run on a machine with an AMD EPYC 9654 CPU and an NVIDIA L40 GPU.

EMDB In the field of structural biology, approximations of the Wasserstein metric are increasingly being used on 2D projection images and 3D volumetric reconstructions of proteins and other macromolecules. Specific applications include molecular alignment, clustering and dimensionality reduction, with most methods substituting the Wasserstein metric with a crude approximation that is fast to compute (Rao et al., 2020; Riahi et al., 2023; Singer & Yang, 2024; Kileel et al., 2021). We evaluate our algorithms in a challenging 3D alignment setting, where we wish to compute the Wasserstein- $p$  metric  $p \in \{1, 2\}$  between rotated 3D density maps of the same molecule. The volumetric density maps are downloaded from the electron microscopy data bank (EMDB) (wwPDB Consortium, 2024) using the ASPIRE package (Wright et al., 2025). Figure 2 shows the computed bounds for rotations between  $0^\circ$  and  $180^\circ$  of the Plasmodium falciparum 80S ribosome 3D density map (Wong et al., 2014). Plots for other molecules are shown in the supplementary material. Dual upscaling (lower bound) and weighted-cost (upper bound) methods at  $\kappa = 2$  provide the best approximations with average relative bounding region of only  $31.37\% \pm 12.42\%$  ( $p = 1$ ) and  $28.06\% \pm 13.15\%$  ( $p = 2$ ), where for each angle the relative bounding region is  $\frac{upper-lower}{upper}$ . A summary of the computational speed-up is provided in Table 2.




Figure 2: Wasserstein distance bounds between rotated 3D density maps of the 80S ribosome. The left panel shows an isosurface plot of the 3D density map that we rotated around the  $z$ -axis. The other two panels compare the different algorithms for producing upper and lower bounds on the Wasserstein- $p$  metric. Due to the logarithmic scale of the  $y$ -axis, the zero-valued dual entropic regularization bounds do not appear in the plot. (center)  $p = 1$ ; (right)  $p = 2$ .

p	Upper Bounds						Lower Bounds					
	Weighted-Cost		Primal Upscaling		Entropic Regularization		Dual Upscaling		Min-Cost		Entropic Regularization	
	$\kappa_2$	$\kappa_4$	$\kappa_2$	$\kappa_4$	$\epsilon_1$	$\epsilon_4$	$\kappa_2$	$\kappa_4$	$\kappa_2$	$\kappa_4$	$\epsilon_1$	$\epsilon_4$
1	0.22%	<b>0.12%</b>	1.09%	27.23%	130.95%	144.97%	0.55%	0.33%	0.20%	<b>0.11%</b>	190.88%	197.35%
	$\pm 0.06\%$	$\pm 0.26\%$	$\pm 0.34\%$	$\pm 15.99\%$	$\pm 29.46\%$	$\pm 39.00\%$	$\pm 0.21\%$	$\pm 0.21\%$	$\pm 0.06\%$	$\pm 0.29\%$	$\pm 58.51\%$	$\pm 52.12\%$
2	0.22%	<b>0.06%</b>	0.30%	0.16%	164.64%	171.42%	0.19%	<b>0.04%</b>	0.21%	<b>0.04%</b>	195.49%	197.76%
	$\pm 0.04\%$	$\pm 0.02\%$	$\pm 0.07\%$	$\pm 0.08\%$	$\pm 42.18\%$	$\pm 41.89\%$	$\pm 0.05\%$	$\pm 0.02\%$	$\pm 0.05\%$	$\pm 0.02\%$	$\pm 42.43\%$	$\pm 44.34\%$

Table 2: Computation time relative to the exact Wasserstein distance computation for 3D Cryo-EM data at  $32 \times 32 \times 32$  resolution. Results show mean  $\pm$  standard deviation across rotations.

**DOTmark** The methods were evaluated on 2D images from the discrete optimal transport benchmark (Schrieber et al., 2017), using  $\rho = L^2$  the euclidean metric, at  $p = \{1, 2\}$ . To examine the effect of the scaling factor, the quantization-based methods were evaluated using  $\kappa = \{2, 4\}$ . To examine the effect of the entropic-regularization parameter, the regularization-based methods were evaluated using  $\epsilon = \{0.001N^p, 0.004N^p\}$  explicitly dependent on  $N^p$  term to avoid large  $\|C\|_\infty/\epsilon$  causing numerical instability (Altschuler et al., 2018), since  $\|C\|_\infty \propto N^p$  in our setting. For upper bounds, while at  $p = 1$  the entropic-regularization upper bound at  $\epsilon = 0.001N^p$  delivers the best

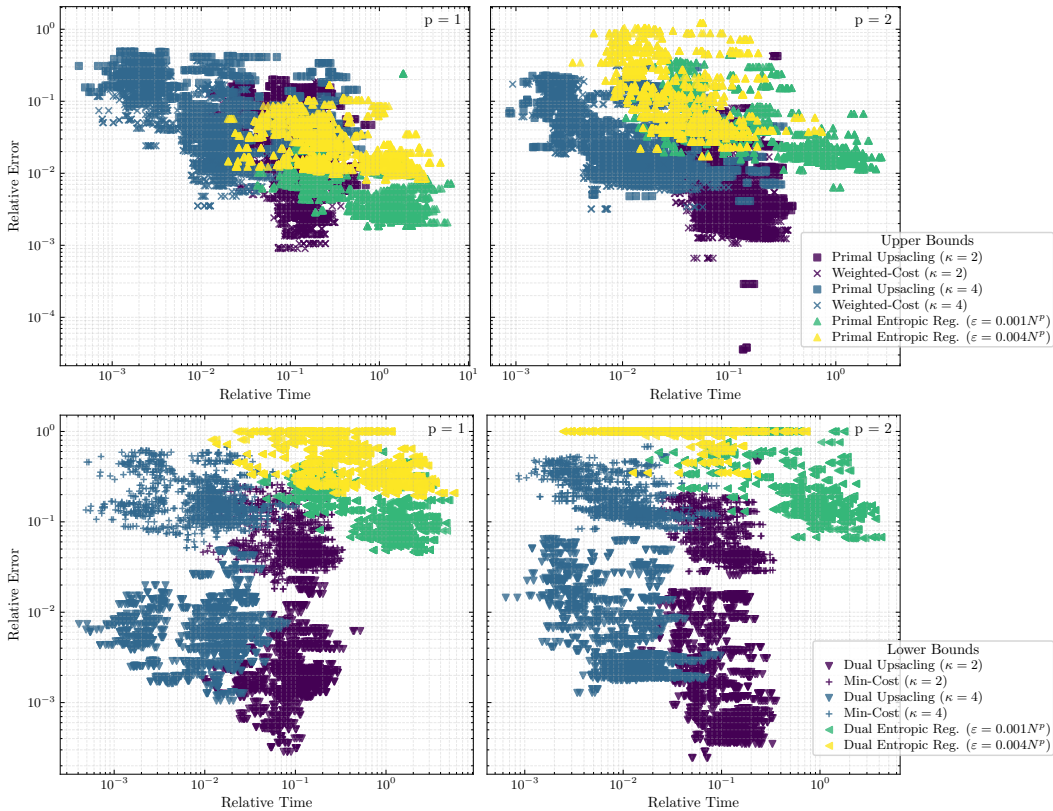


Figure 3: This figure shows the time/accuracy tradeoff of the different bounds on image pairs from the DOTMark dataset. Each point in the plots represents an image pair whose Wasserstein metric is estimated using a particular method. The  $x$  value of the point is the time taken to compute the estimate, relative to the runtime of the exact Wasserstein computation. The  $y$  value is the relative error of the bound. (top left) upper bounds for  $W_1$ . (top right) upper bounds for  $W_2$ . (bottom left) lower bounds for  $W_1$ . (lower right) lower bounds for  $W_2$ .

approximation, summarized in Table 4 in Appendix B, it does so with significant impact on the computation time, as seen in Figure 3. In contrast, for both upper and lower bounds, the quantization methods yield the best approximations at  $\kappa = 2$ , maintaining low computation time and ranking second in the relative time benchmark after the quantization methods scaled at  $\kappa = 4$ .

## 6 CONCLUSION AND DISCUSSION

In this paper, we proposed several methods for computing fast approximations that lower or upper-bound the Wasserstein metric between discrete distributions on a regular grid. Our experiments on 2D images and 3D volumetric data demonstrate significant improvements in computational efficiency and accuracy compared to bounds based on entropic OT. In future work, our approach could be refined and extended by exploring different methods for the downscaling and upscaling stages in addition to multi-scale approaches. The methods could also be extended to domains beyond regular grids such as point clouds in  $\mathbb{R}^n$  and graphs by picking coarsening and interpolation operations that are suitable for the domain.

**Code availability and reproducibility** Python implementations of all the methods in the paper, as well as code for reproducing the figures and tables, will be released as free software on Github.

486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539

---

## REFERENCES

- Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice-Hall, Inc., USA, February 1993. ISBN 978-0-13-617549-0.
- Jason Altschuler, Jonathan Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration, February 2018.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 214–223. PMLR, July 2017. URL <https://proceedings.mlr.press/v70/arjovsky17a.html>.
- Kubilay Atasu and Thomas Mittelholzer. Linear-Complexity Data-Parallel Earth Mover’s Distance Approximations. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 364–373. PMLR, May 2019. URL <https://proceedings.mlr.press/v97/atasul9a.html>.
- Arturs Backurs, Yihe Dong, Piotr Indyk, Ilya Razenshteyn, and Tal Wagner. Scalable nearest neighbor search for optimal transport. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *ICML’20*, pp. 497–506. JMLR.org, July 2020.
- Amartya Banerjee, Harlin Lee, Nir Sharon, and Caroline Moosmüller. Efficient Trajectory Inference in Wasserstein Space Using Consecutive Averaging. In *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, pp. 2260–2268. PMLR, April 2025. URL <https://proceedings.mlr.press/v258/banerjee25a.html>.
- Alberto Bartesaghi, Alan Merk, Mario J. Borgnia, Jacqueline L. S. Milne, and Sriram Subramaniam. Prefusion structure of trimeric HIV-1 envelope glycoprotein determined by cryo-electron microscopy. *Nature Structural & Molecular Biology*, 20(12):1352–1357, December 2013. doi:10.1038/nsmb.2711.
- Gaspard Beugnot, Aude Genevay, Kristjan Greenewald, and Justin Solomon. Improving approximate optimal transport distances using quantization. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, pp. 290–300. PMLR, December 2021. URL <https://proceedings.mlr.press/v161/beugnot21a.html>.
- Nicolas Bonneel, Michiel van de Panne, Sylvain Paris, and Wolfgang Heidrich. Displacement interpolation using Lagrangian mass transport. *ACM Transactions on Graphics*, 30(6):1–12, December 2011. doi:10.1145/2070781.2024192.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Yidong Chen, Chen Li, and Zhonghua Lu. Computing Wasserstein-p Distance Between Images With Linear Cost. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 519–528, 2022. URL [https://openaccess.thecvf.com/content/CVPR2022/html/Chen\\_Computing\\_Wasserstein-p\\_Distance\\_Between\\_Images\\_With\\_Linear\\_Cost\\_CVPR\\_2022\\_paper.html](https://openaccess.thecvf.com/content/CVPR2022/html/Chen_Computing_Wasserstein-p_Distance_Between_Images_With_Linear_Cost_CVPR_2022_paper.html).
- Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal Transport for Domain Adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1853–1865, September 2017. doi:10.1109/TPAMI.2016.2615921.
- Katy Craig and Haoqing Yu. Wavelet Wasserstein distances for  $0 < p \leq 1$ . *Sampling Theory, Signal Processing, and Data Analysis*, 23(2):19, August 2025. doi:10.1007/s43670-025-00113-4.
- Marco Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- Marco Cuturi, Laetitia Meng-Papaxanthos, Yingtao Tian, Charlotte Bunne, Geoff Davis, and Olivier Teboul. Optimal Transport Tools (OTT): A JAX Toolbox for all things Wasserstein, January 2022.

- 
- 540 Ishan Deshpande, Ziyu Zhang, and Alexander G. Schwing. Generative Model-  
541 ing Using the Sliced Wasserstein Distance. In *Proceedings of the IEEE Confer-*  
542 *ence on Computer Vision and Pattern Recognition*, pp. 3483–3491, 2018. URL  
543 [https://openaccess.thecvf.com/content\\_cvpr\\_2018/html/Deshpande\\_](https://openaccess.thecvf.com/content_cvpr_2018/html/Deshpande_Generative_Modeling_Using_CVPR_2018_paper.html)  
544 [Generative\\_Modeling\\_Using\\_CVPR\\_2018\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2018/html/Deshpande_Generative_Modeling_Using_CVPR_2018_paper.html).
- 545  
546 Jean Feydy, Pierre Roussillon, Alain Trouvé, and Pietro Gori. Fast and Scalable Optimal Transport  
547 for Brain Tractograms, July 2021.
- 548 Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas  
549 Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron,  
550 Nathalie T. H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet,  
551 Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and  
552 Titouan Vayer. POT: Python Optimal Transport. *Journal of Machine Learning Research*, 22(78):  
553 1–8, 2021. URL <http://jmlr.org/papers/v22/20-451.html>.
- 554  
555 Samuel Gerber and Mauro Maggioni. Multiscale strategies for computing optimal transport. *The*  
556 *Journal of Machine Learning Research*, 18(1):2440–2471, January 2017.
- 557  
558 Piotr Indyk and Nitin Thaper. Fast Image Retrieval via Embeddings. In *3rd International Workshop*  
559 *on Statistical and Computational Theories of Vision (at ICCV)*, 2003.
- 560  
561 Joe Kileel, Amit Moscovich, Nathan Zelesko, and Amit Singer. Manifold Learning with Arbitrary  
562 Norms. *Journal of Fourier Analysis and Applications*, 27(5), October 2021. doi:10.1007/s00041-  
563 021-09879-2.
- 563  
564 Paul Knopp and Richard Sinkhorn. Concerning nonnegative matrices  
565 and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21  
566 (2):343–348, January 1967. URL [https://projecteuclid.org/](https://projecteuclid.org/journals/pacific-journal-of-mathematics/volume-21/issue-2/Concerning-nonnegative-matrices-and-doubly-stochastic-matrices/pjm/1102992505.full)  
567 [journals/pacific-journal-of-mathematics/volume-21/issue-2/](https://projecteuclid.org/journals/pacific-journal-of-mathematics/volume-21/issue-2/Concerning-nonnegative-matrices-and-doubly-stochastic-matrices/pjm/1102992505.full)  
568 [Concerning-nonnegative-matrices-and-doubly-stochastic-matrices/](https://projecteuclid.org/journals/pacific-journal-of-mathematics/volume-21/issue-2/Concerning-nonnegative-matrices-and-doubly-stochastic-matrices/pjm/1102992505.full)  
569 [pjm/1102992505.full](https://projecteuclid.org/journals/pacific-journal-of-mathematics/volume-21/issue-2/Concerning-nonnegative-matrices-and-doubly-stochastic-matrices/pjm/1102992505.full).
- 570  
571 Soheil Kolouri, Kimia Nadjahi, Umut Simsekli, Roland Badeau, and Gustavo Ro-  
572 hde. Generalized Sliced Wasserstein Distances. In *Advances in Neural Infor-*  
573 *mation Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL  
574 [https://proceedings.neurips.cc/paper\\_files/paper/2019/hash/](https://proceedings.neurips.cc/paper_files/paper/2019/hash/f0935e4cd5920aa6c7c996a5ee53a70f-Abstract.html)  
575 [f0935e4cd5920aa6c7c996a5ee53a70f-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/2019/hash/f0935e4cd5920aa6c7c996a5ee53a70f-Abstract.html).
- 576  
577 Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. From Word Embeddings To Document  
578 Distances. In *Proceedings of the 32nd International Conference on Machine Learning*, pp. 957–966.  
579 PMLR, June 2015. URL <https://proceedings.mlr.press/v37/kusnerb15.html>.
- 580  
581 Tianyi Lin, Nhat Ho, and Michael I Jordan. On the Efficiency of Entropic Regularized Algorithms  
582 for Optimal Transport. *Journal of Machine Learning Research*, 23(137):1–42, 2022. URL  
583 <http://jmlr.org/papers/v23/20-277.html>.
- 584  
585 Elaine C. Meng, Thomas D. Goddard, Eric F. Pettersen, Greg S. Couch, Zach J. Pearson, John H.  
586 Morris, and Thomas E. Ferrin. UCSF ChimeraX: Tools for structure building and analysis. *Protein*  
587 *Science*, 32(11):e4792, 2023. doi:10.1002/pro.4792.
- 588  
589 Quentin Mérigot. A Multiscale Approach to Optimal Transport. *Computer Graphics Forum*, 30(5):  
590 1583–1592, August 2011. doi:10.1111/j.1467-8659.2011.02032.x.
- 591  
592 Eduardo Fernandes Montesuma, Fred Maurice Ngolè Mboula, and Antoine Souloumiac. Recent  
593 Advances in Optimal Transport for Machine Learning. *IEEE Transactions on Pattern Analysis and*  
*Machine Intelligence*, 47(2):1161–1180, February 2025. doi:10.1109/TPAMI.2024.3489030.
- 594  
595 Caroline Moosmüller and Alexander Cloninger. Linear optimal transport embedding: provable  
596 Wasserstein classification for certain rigid transformations and perturbations. *Information and*  
597 *Inference: A Journal of the IMA*, 12(1):363–389, March 2023. doi:10.1093/imaiai/iaac023.

- 
- 594 Thi Hoang Duong Nguyen, Wojciech P. Galej, Xiao-chen Bai, Chris Oubridge, Andrew J. Newman,  
595 Sjors H. W. Scheres, and Kiyoshi Nagai. Cryo-EM structure of the yeast U4/U6.U5 tri-snRNP at  
596 3.7 Å resolution. *Nature*, 530(7590):298–302, February 2016. doi:[10.1038/nature16940](https://doi.org/10.1038/nature16940).
- 597
- 598 G. I. Papayiannis, G. N. Domazakis, D. Drivaliaris, S. Koukoulas, A. E. Tsekrekos, and A. N.  
599 Yannacopoulos. On clustering uncertain and structured data with Wasserstein barycenters and a  
600 geodesic criterion for the number of clusters. *Journal of Statistical Computation and Simulation*,  
601 91(13):2569–2594, September 2021. doi:[10.1080/00949655.2021.1903463](https://doi.org/10.1080/00949655.2021.1903463).
- 602 Gabriel Peyré and Marco Cuturi. Computational Optimal Transport: With Applications to  
603 Data Science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.  
604 doi:[10.1561/22000000073](https://doi.org/10.1561/22000000073).
- 605
- 606 Rohan Rao, Amit Moscovich, and Amit Singer. Wasserstein K-Means for Clustering Tomographic  
607 Projections. In *Machine Learning for Structural Biology Workshop, Neural Information Processing  
608 Systems (NeurIPS)*, pp. 1–12, 2020. URL [https://www.mlsb.io/papers/MLSB2020\\_  
609 Wasserstein\\_K-Means\\_for\\_Clustering.pdf](https://www.mlsb.io/papers/MLSB2020_Wasserstein_K-Means_for_Clustering.pdf).
- 610 Aryan Tajmir Riahi, Geoffrey Woollard, Frédéric Poitevin, Anne Condon, and Khanh Dao Duc.  
611 AlignOT: An Optimal Transport Based Algorithm for Fast 3D Alignment With Applications to  
612 Cryogenic Electron Microscopy Density Maps. *IEEE/ACM Transactions on Computational Biology  
613 and Bioinformatics*, 20(6):3842–3850, November 2023. doi:[10.1109/TCBB.2023.3327633](https://doi.org/10.1109/TCBB.2023.3327633).
- 614 Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon,  
615 Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, Lia Lee, Jenny Chen, Justin Brumbaugh,  
616 Philippe Rigollet, Konrad Hochedlinger, Rudolf Jaenisch, Aviv Regev, and Eric S. Lander. Optimal-  
617 Transport Analysis of Single-Cell Gene Expression Identifies Developmental Trajectories in  
618 Reprogramming. *Cell*, 176(4):928–943.e22, February 2019. doi:[10.1016/j.cell.2019.01.006](https://doi.org/10.1016/j.cell.2019.01.006).
- 619 Jörn Schrieber, Dominic Schuhmacher, and Carsten Gottschlich. DOTmark – A Benchmark for  
620 Discrete Optimal Transport. *IEEE Access*, 5:271–282, 2017. doi:[10.1109/ACCESS.2016.2639065](https://doi.org/10.1109/ACCESS.2016.2639065).
- 621
- 622 Sameer Shirdhonkar and David W. Jacobs. Approximate earth mover’s distance in linear time.  
623 In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8, June 2008.  
624 doi:[10.1109/CVPR.2008.4587662](https://doi.org/10.1109/CVPR.2008.4587662).
- 625 Amit Singer and Ruiyi Yang. Alignment of density maps in Wasserstein distance. *Biological Imaging*,  
626 4:e5, January 2024. doi:[10.1017/S2633903X24000059](https://doi.org/10.1017/S2633903X24000059).
- 627
- 628 Justin Solomon, Fernando de Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen,  
629 Tao Du, and Leonidas Guibas. Convolutional Wasserstein Distances: Efficient Optimal Trans-  
630 portation on Geometric Domains. *ACM Transactions on Graphics*, 34(4):66:1–66:11, 2015.  
631 doi:[10.1145/2766963](https://doi.org/10.1145/2766963).
- 632 Soledad Stagnoli, Francesca Peccati, Sean R. Connell, Ane Martinez-Castillo, Diego Charro, Oscar  
633 Millet, Chiara Bruzzone, Asis Palazon, Ana Ardá, Jesús Jiménez-Barbero, June Ereño-Orbea,  
634 Nicola G. A. Abrescia, and Gonzalo Jiménez-Osés. Assessing the Mobility of Severe Acute  
635 Respiratory Syndrome Coronavirus-2 Spike Protein Glycans by Structural and Computational  
636 Methods. *Frontiers in Microbiology*, 13, April 2022. doi:[10.3389/fmicb.2022.870938](https://doi.org/10.3389/fmicb.2022.870938).
- 637 Cédric Villani. *Optimal Transport, old and new*, volume 338 of *Grundlehren der mathematischen  
638 Wissenschaften*. Springer, Berlin, Heidelberg, 2009. ISBN 978-3-540-71050-9. doi:[10.1007/978-  
639 3-540-71050-9](https://doi.org/10.1007/978-3-540-71050-9).
- 640
- 641 Wilson Wong, Xiao-chen Bai, Alan Brown, Israel S Fernandez, Eric Hanssen, Melanie Condon,  
642 Yan Hong Tan, Jake Baum, and Sjors HW Scheres. Cryo-EM structure of the Plasmodium  
643 falciparum 80S ribosome bound to the anti-protozoan drug emetine. *eLife*, 3:e03080, June 2014.  
644 doi:[10.7554/eLife.03080](https://doi.org/10.7554/eLife.03080).
- 645 Garrett Wright, Joakim Andén, Vineet Bansal, Junchao Xia, Chris Langfield, Josh Carmichael, Kris  
646 Sowattanangkul, Robbie Brook, Yunpeng Shi, Ayelet Heimowitz, Gabi Pragier, Itay Sason, Amit  
647 Moscovich, Yoel Shkolnisky, and Amit Singer. ComputationalCryoEM/ASPIRE-Python: v0.13.2.  
Zenodo, March 2025.

---

648 wwPDB Consortium. EMDB—the Electron Microscopy Data Bank. *Nucleic Acids Research*, 52  
649 (D1):D456–D465, January 2024. doi:[10.1093/nar/gkad1019](https://doi.org/10.1093/nar/gkad1019).  
650

651 Wenzhe Yang, Sheng Wang, Yuan Sun, and Zhiyong Peng. Fast dataset search with earth mover’s  
652 distance. *Proc. VLDB Endow.*, 15(11):2517–2529, July 2022. doi:[10.14778/3551793.3551811](https://doi.org/10.14778/3551793.3551811).  
653  
654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701

---

702 APPENDIX A ALGORITHMS

703  
704  
705 Lower-bound based on entropic regularization described in Section 3.1. This method simply runs the  
706 iterative Sinkhorn algorithm and then returns the unregularized cost term of the dual solution to the  
707 entropic regularization problem.  
708

---

709  
710  
711 **Algorithm 1** Regularization-Based Lower Bound

---

712 **Require:**  $\mu, \nu \in \Sigma_{N^d}$  on  $\mathcal{X} = [N]^d$ ,  $p \geq 1, \varepsilon > 0, \xi > 0$  and metric  $\rho$ .

713  $C \leftarrow \{\rho(x_i, y_j)^p\}_{ij}$

714  $K \leftarrow \exp\left(-\frac{C}{\varepsilon}\right)$

715 Initialize  $\mathbf{f} \leftarrow \mathbf{0}_n, \mathbf{g} \leftarrow \mathbf{0}_m$

716 Initialize  $\mathbf{b} \leftarrow \exp\left(\frac{\mathbf{g}}{\varepsilon}\right)$

717 **repeat**

718  $\mathbf{a} \leftarrow \mu \odot (K\mathbf{b})$

719  $\mathbf{f} \leftarrow \varepsilon \log \mathbf{a}$

720  $\mathbf{b} \leftarrow \nu \odot (K^\top \mathbf{a})$

721  $\mathbf{g} \leftarrow \varepsilon \log \mathbf{b}$

722 **until**  $\|\mathbf{a} \odot K\mathbf{b} - \mu\|_1 + \|\nu - \mathbf{b} \odot K^\top \mathbf{a}\|_1 < \xi$

723 **return**  $(\langle \mathbf{f}, \mu \rangle + \langle \mathbf{g}, \nu \rangle)^{\frac{1}{p}}$

---

724  
725  
726  
727  
728 Algorithm for entropic regularization-based upper bound described in Section 3.1, using the un-  
729 regularized term of the primal solution to the entropic regularization problem, with weighted total  
730 variation marginal corrections.  
731

---

732  
733  
734 **Algorithm 2** Regularization-Based Upper Bound

---

735 **Require:**  $\mu, \nu \in \Sigma_{N^d}$  on  $\mathcal{X} = [N]^d$ ,  $p \geq 1, \varepsilon > 0, \xi > 0$  and metric  $\rho$ .

736  $C \leftarrow \{\rho(x_i, y_j)^p\}_{ij}$

737  $K \leftarrow \exp\left(-\frac{C}{\varepsilon}\right)$

738 Initialize  $\mathbf{b} \leftarrow \mathbf{1}_{N^d}$

739 **repeat**

740  $\mathbf{a} \leftarrow \mu \odot K\mathbf{b}$

741  $\mathbf{b} \leftarrow \nu \odot K^\top \mathbf{a}$

742  $\hat{\mu} \leftarrow \mathbf{a} \odot (K\mathbf{b})$

743  $\hat{\nu} \leftarrow \mathbf{b} \odot (K^\top \mathbf{a})$

744 **until**  $\|\hat{\mu} - \mu\|_1 + \|\nu - \hat{\nu}\|_1 < \xi$

745  $\hat{\pi}_\varepsilon \leftarrow \text{diag}(\mathbf{a})K \text{diag}(\mathbf{b})$

746  $\bar{x} \leftarrow \text{mean}(\mathcal{X})$

747  $\mathbf{w} \leftarrow \{\rho(\bar{x}, x_i)^p\}_i$

748  $\Delta_{\hat{\mu}} \leftarrow 2^{1-\frac{1}{p}} \langle \mathbf{w}, |\hat{\mu} - \mu| \rangle^{\frac{1}{p}}$

749  $\Delta_{\hat{\nu}} \leftarrow 2^{1-\frac{1}{p}} \langle \mathbf{w}, |\nu - \hat{\nu}| \rangle^{\frac{1}{p}}$

750 **return**  $\langle \hat{\pi}_\varepsilon, C \rangle^{\frac{1}{p}} + \Delta_{\hat{\mu}} + \Delta_{\hat{\nu}}$

---

751  
752  
753  
754  
755 Algorithm for quantization-based upper bound described in Section 3.2, using coarse cost weighted  
by the trivial coupling.

756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809

---

**Algorithm 3** Weighted-Cost Upper Bound

---

**Require:**  $\mu, \nu \in \Sigma_{N^d}$  on  $\mathcal{X} = [N]^d$ ,  $p \geq 1$ ,  $\kappa \in \mathbb{N}$  and metric  $\rho$ .  
 $\tilde{\mu} \leftarrow \text{SumPool}(\mu; \kappa)$   
 $\tilde{\nu} \leftarrow \text{SumPool}(\nu; \kappa)$   
 $\bar{C}_{k\ell} \leftarrow \left\{ \frac{1}{\tilde{\mu}_k \tilde{\nu}_\ell} \sum_{\substack{x \in X_k \\ y \in Y_\ell}} \rho(x, y)^p \mu(x) \nu(y) \right\}_{k\ell}$   
Solve  $L_{\bar{C}} \leftarrow \min_{\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})} \langle \tilde{\pi}, \bar{C} \rangle$   
**return**  $L_{\bar{C}}^{\frac{1}{p}}$

---

Algorithm for bi-level quantization-based upper bound described in Section 3.4, using nearest-neighbor upscaling of the optimal coarse coupling, iterative proportional fitting of the marginals (i.e. Sinkhorn iterations), with weighted total variation marginal corrections.

---

**Algorithm 4** Upscaling Upper Bound

---

**Require:**  $\mu, \nu \in \Sigma_{N^d}$  on  $\mathcal{X} = [N]^d$ ,  $p \geq 1$ ,  $\kappa \in \mathbb{N}$ ,  $\xi > 0$  and metric  $\rho$ .  
 $\mathcal{X} \leftarrow \text{AvgPool}(\mathcal{X}; \kappa)$   
 $\tilde{\mu} \leftarrow \text{SumPool}(\mu; \kappa)$   
 $\tilde{\nu} \leftarrow \text{SumPool}(\nu; \kappa)$   
 $\tilde{C} \leftarrow \{\rho(\tilde{x}_k, \tilde{x}_\ell)^p\}_{k\ell}$   
Solve  $\tilde{\pi}^* \leftarrow \arg \min_{\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})} \langle \tilde{\pi}, \tilde{C} \rangle$   
▷ **Up-scaled coupling**  
 $\tilde{\mathbf{P}}^* \leftarrow \text{reshape}(\tilde{\pi}^*; n)$  ▷ Reshape as tensor  
 $\mathbf{K} \leftarrow \{\kappa^{-2d}\}_{t \in [n]^{2d}}$   
 $\hat{\mathbf{P}} \leftarrow \tilde{\mathbf{P}}^* \otimes \mathbf{K}$  ▷ Upscaling  
 $\hat{\pi} \leftarrow \text{reshape}^{-1}(\hat{\mathbf{P}}; N)$   
▷ **Iterative proportional fitting**  
Initialize  $\mathbf{b} \leftarrow \mathbf{1}_{N^d}$   
**repeat**  
 $\mathbf{a} \leftarrow \mu \otimes \hat{\pi} \mathbf{b}$   
 $\mathbf{b} \leftarrow \nu \otimes \hat{\pi}^\top \mathbf{a}$   
 $\hat{\mu} \leftarrow \mathbf{a} \odot (\hat{\pi} \mathbf{b})$   
 $\hat{\nu} \leftarrow \mathbf{b} \odot (\hat{\pi}^\top \mathbf{a})$   
**until**  $\|\hat{\mu} - \mu\|_1 + \|\hat{\nu} - \nu\|_1 < \xi$   
▷ **Upper bound**  
 $\widehat{\mathcal{W}}_p \leftarrow \left( \sum_{(i,j) \in \text{supp}(\hat{\pi})} a_i \hat{\pi}_{ij} b_j \rho(x_i, x_j)^p \right)^{\frac{1}{p}}$   
 $\bar{x} \leftarrow \text{mean}(\mathcal{X})$   
 $\mathbf{w} \leftarrow \{\rho(\bar{x}, x_i)^p\}_i$   
 $\Delta_{\hat{\mu}} \leftarrow 2^{1-\frac{1}{p}} \langle \mathbf{w}, |\hat{\mu} - \mu| \rangle^{\frac{1}{p}}$   
 $\Delta_{\hat{\nu}} \leftarrow 2^{1-\frac{1}{p}} \langle \mathbf{w}, |\nu - \hat{\nu}| \rangle^{\frac{1}{p}}$   
**return**  $\widehat{\mathcal{W}}_p + \Delta_{\hat{\mu}} + \Delta_{\hat{\nu}}$

---

Algorithm for bi-level quantization-based lower bound described in Section 3.5, using interpolation for upscaling the optimal coarse dual potentials, and c-transform to achieve optimized admissible dual potentials pair.

---

**Algorithm 5** Upscaling Lower Bound

---

**Require:**  $\mu, \nu \in \Sigma_{N^d}$  on  $\mathcal{X} = [N]^d$ ,  $p \geq 1$ ,  $\kappa \in \mathbb{N}$  and metric  $\rho$ .

$\tilde{\mathcal{X}} \leftarrow \text{AvgPool}(\mathcal{X}; \kappa)$

$\tilde{\mu} \leftarrow \text{SumPool}(\mu; \kappa)$

$\tilde{\nu} \leftarrow \text{SumPool}(\nu; \kappa)$

$\tilde{C} \leftarrow \{\rho(\tilde{x}_k, \tilde{x}_\ell)^p\}_{k\ell}$

Solve  $(\tilde{\mathbf{f}}^*, \tilde{\mathbf{g}}^*) \leftarrow \arg \max_{(\tilde{\mathbf{f}}, \tilde{\mathbf{g}}) \in \mathcal{R}(\tilde{C})} \langle \tilde{\mathbf{f}}, \tilde{\mu} \rangle + \langle \tilde{\mathbf{g}}, \tilde{\nu} \rangle$

$\hat{\mathbf{f}} \leftarrow \{R_{\tilde{\mathbf{f}}, \tilde{\mathcal{X}}}(x_i)\}_{i \in [N^d]}$

$\mathbf{g} \leftarrow \left\{ \min_i \rho(x_i, x_j)^p - \hat{f}_i \right\}_j$

$\mathbf{f} \leftarrow \left\{ \min_j \rho(x_i, x_j)^p - g_j \right\}_i$

**return**  $(\langle \mathbf{f}, \mu \rangle + \langle \mathbf{g}, \nu \rangle)^{\frac{1}{p}}$

---

## APPENDIX B ADDITIONAL EXPERIMENTS

### APPENDIX B.1 DOTMARK

In this section we present additional figures and results evaluating the proposed Wasserstein bounds on the discrete optimal transport benchmark (DOTMark) (Schrieber et al., 2017) presented in the main text. The computational speed up of the proposed methods compared to the exact OT solver are summarized in Table 3. The results show that the quantization methods achieve the most significant speed ups. Notably, the dual upscaling method at  $\kappa = 4$  are calculated in 0.2-2.2% of the time, while making almost no sacrifice in accuracy. Maintaining no more than 2.4% average error.

The relative accuracy of the proposed methods exponentially improves for large values of the exact Wasserstein distance as evident in Figures 4 and 5. Negative-valued lower bounds are trivially clipped to 0, when evaluate in the benchmark.

### APPENDIX B.2 EMDB

The Electron Microscopy Data Bank (EMDB) (wwPDB Consortium, 2024) is a repository of volumetric density maps that contains many interesting molecules that were reconstructed from cryogenic electron microscopy (cryo-EM) experiments. These reconstructions are estimates of the 3D electric potential at every point in the molecule. For our 3D experiments, we downloaded and processed four maps of famous molecules, detailed in Table 5 using the ASPIRE package (Wright et al., 2025). In Figure 2 you can see 3D renderings of these molecules that we generated using UCSF ChimeraX (Meng et al., 2023). The volumetric maps were downloaded from EMDB, masked inside a spherical region of radius 128 pixels, rotated around the Z axis in increments of 20 degrees and downsampled to  $16 \times 16 \times 16$  voxels. The computational speed up is summarized in Table 6, showing that even at  $\kappa = 2$  the quantization methods provide substantial speedups. Figure 7 shows the Wasserstein metrics and bounds between the 3D density map of the molecule in its base orientation and its rotations around the Z axis. The exact Wasserstein metric is shown as the thick black line with upper and lower bounds next to it using the various markers.

The quantization-based methods dominate in accuracy for the lower bounds of both  $p \in \{1, 2\}$ , whereas for the upper bounds of the Wasserstein-1 metric, the upper bound based on entropic regularization with  $\varepsilon = 0.001N^p$  achieves the best accuracy, although at a significant computational cost. The triangular symmetry of EMDB-14621 (SARS-CoV-2 spike protein) and EMDB-2484 (Trimeric HIV-1 envelope glycoprotein) seen in Figure 7 are easily detectable as the dips at  $120^\circ$  rotation angle.

864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917

		Upper Bounds						Lower Bounds					
		Weighted-Cost		Primal Upscaling		Entropic Regularization		Dual Upscaling		Min-Cost		Entropic Regularization	
		$\kappa_2$	$\kappa_4$	$\kappa_2$	$\kappa_4$	$\varepsilon_1$	$\varepsilon_4$	$\kappa_2$	$\kappa_4$	$\kappa_2$	$\kappa_4$	$\varepsilon_1$	$\varepsilon_4$
Class	p												
Classic Images	1	4.4%	<b>0.3%</b>	5.0%	<b>0.3%</b>	19.8%	16.0%	4.9%	<b>0.2%</b>	4.7%	0.3%	17.6%	18.7%
		$\pm 2.1\%$	$\pm 0.2\%$	$\pm 2.6\%$	$\pm 0.4\%$	$\pm 10.1\%$	$\pm 11.2\%$	$\pm 2.6\%$	$\pm 0.1\%$	$\pm 2.3\%$	$\pm 0.1\%$	$\pm 7.8\%$	$\pm 13.3\%$
Microscopy	2	6.1%	0.4%	6.1%	<b>0.3%</b>	7.9%	1.3%	6.1%	<b>0.3%</b>	6.1%	<b>0.3%</b>	12.3%	1.3%
		$\pm 3.1\%$	$\pm 0.1\%$	$\pm 3.0\%$	$\pm 0.1\%$	$\pm 6.6\%$	$\pm 0.5\%$	$\pm 3.0\%$	$\pm 0.1\%$	$\pm 2.9\%$	$\pm 0.1\%$	$\pm 10.4\%$	$\pm 0.6\%$
Shapes	1	16.7%	<b>2.3%</b>	20.7%	6.6%	128.0%	90.7%	14.4%	2.2%	15.7%	<b>1.9%</b>	120.8%	86.9%
		$\pm 5.6\%$	$\pm 1.4\%$	$\pm 9.6\%$	$\pm 10.6\%$	$\pm 68.3\%$	$\pm 79.2\%$	$\pm 5.0\%$	$\pm 1.2\%$	$\pm 4.8\%$	$\pm 1.0\%$	$\pm 57.2\%$	$\pm 81.3\%$
Shapes	2	15.4%	<b>1.9%</b>	17.4%	3.0%	75.6%	7.7%	15.6%	2.0%	14.9%	<b>1.7%</b>	99.9%	8.1%
		$\pm 5.2\%$	$\pm 1.2\%$	$\pm 5.9\%$	$\pm 3.4\%$	$\pm 51.3\%$	$\pm 4.9\%$	$\pm 5.1\%$	$\pm 1.3\%$	$\pm 5.0\%$	$\pm 1.0\%$	$\pm 72.8\%$	$\pm 5.5\%$
Shapes	1	10.2%	<b>1.9%</b>	15.1%	10.8%	172.9%	68.8%	7.8%	1.5%	8.5%	<b>1.4%</b>	163.6%	77.0%
		$\pm 3.6\%$	$\pm 1.7\%$	$\pm 8.9\%$	$\pm 11.9\%$	$\pm 99.0\%$	$\pm 72.7\%$	$\pm 2.0\%$	$\pm 0.7\%$	$\pm 2.4\%$	$\pm 1.2\%$	$\pm 88.8\%$	$\pm 77.6\%$
Shapes	2	7.9%	<b>1.5%</b>	9.6%	4.2%	39.5%	6.0%	7.4%	1.2%	7.3%	<b>1.0%</b>	56.2%	5.7%
		$\pm 3.8\%$	$\pm 1.8\%$	$\pm 5.6\%$	$\pm 6.4\%$	$\pm 45.0\%$	$\pm 11.6\%$	$\pm 2.6\%$	$\pm 0.6\%$	$\pm 3.0\%$	$\pm 1.4\%$	$\pm 69.1\%$	$\pm 11.8\%$

Table 3: Computational time comparison of different methods showing the mean  $\pm$  standard deviation of the relative computation time compared to exact OT solver. Each method is evaluated at different fidelity level  $\kappa_2 = 2$  and  $\kappa_4 = 4$  and different values of  $\varepsilon_1 = 1 \cdot 10^{-3} N^p$  and  $\varepsilon_4 = 4 \cdot 10^{-3} N^p$ .

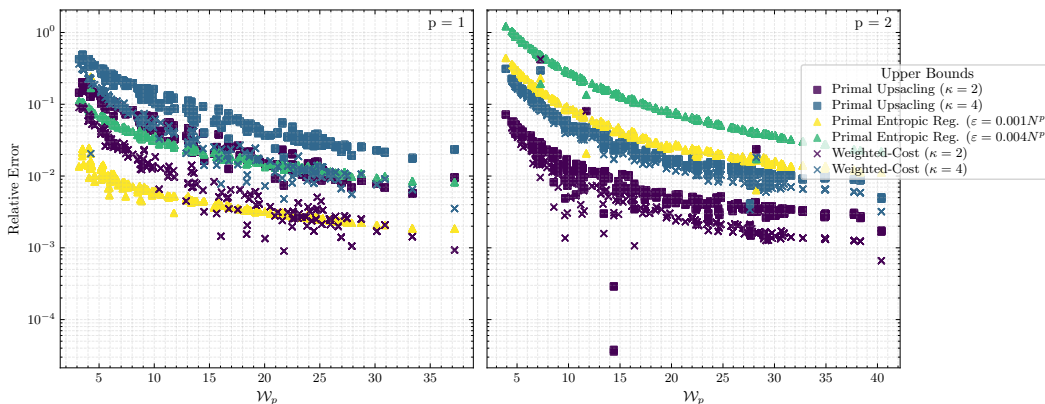


Figure 4: Accuracy of Wasserstein Upper Bounds

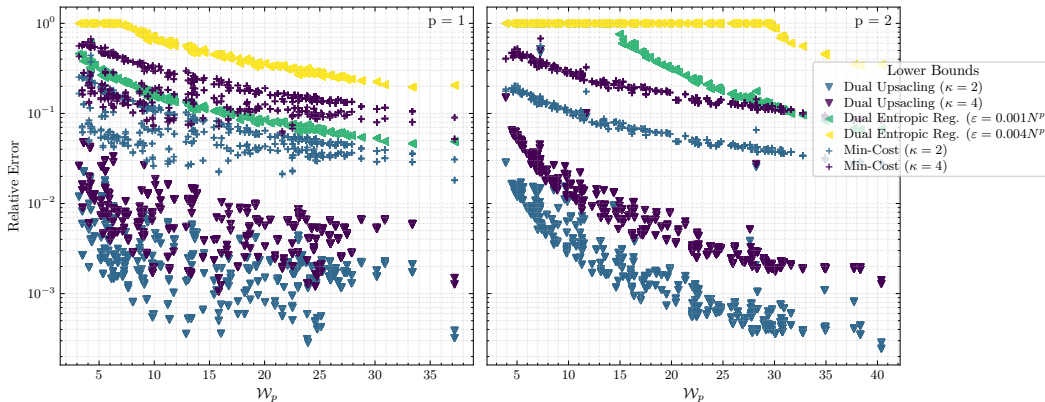


Figure 5: Accuracy of Wasserstein Lower Bounds. Negative-valued bounds are clipped to 0, evaluating as 100% relative error.

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971




Figure 6: Isosurfaces of 3D molecular densities from the Electron Microscopy Data Bank (EMDB). In our experiments, the molecules on the left are rotated around the Z axis, which corresponds to the depth direction here. The middle and right columns show the same molecules rotated by 90 degrees around the X and Y axis (respectively).

972  
 973  
 974  
 975  
 976  
 977  
 978  
 979  
 980  
 981  
 982  
 983  
 984  
 985  
 986  
 987  
 988  
 989  
 990  
 991  
 992  
 993  
 994  
 995  
 996  
 997  
 998  
 999  
 1000  
 1001  
 1002  
 1003  
 1004  
 1005  
 1006  
 1007  
 1008  
 1009  
 1010  
 1011  
 1012  
 1013  
 1014  
 1015  
 1016  
 1017  
 1018  
 1019  
 1020  
 1021  
 1022  
 1023  
 1024  
 1025

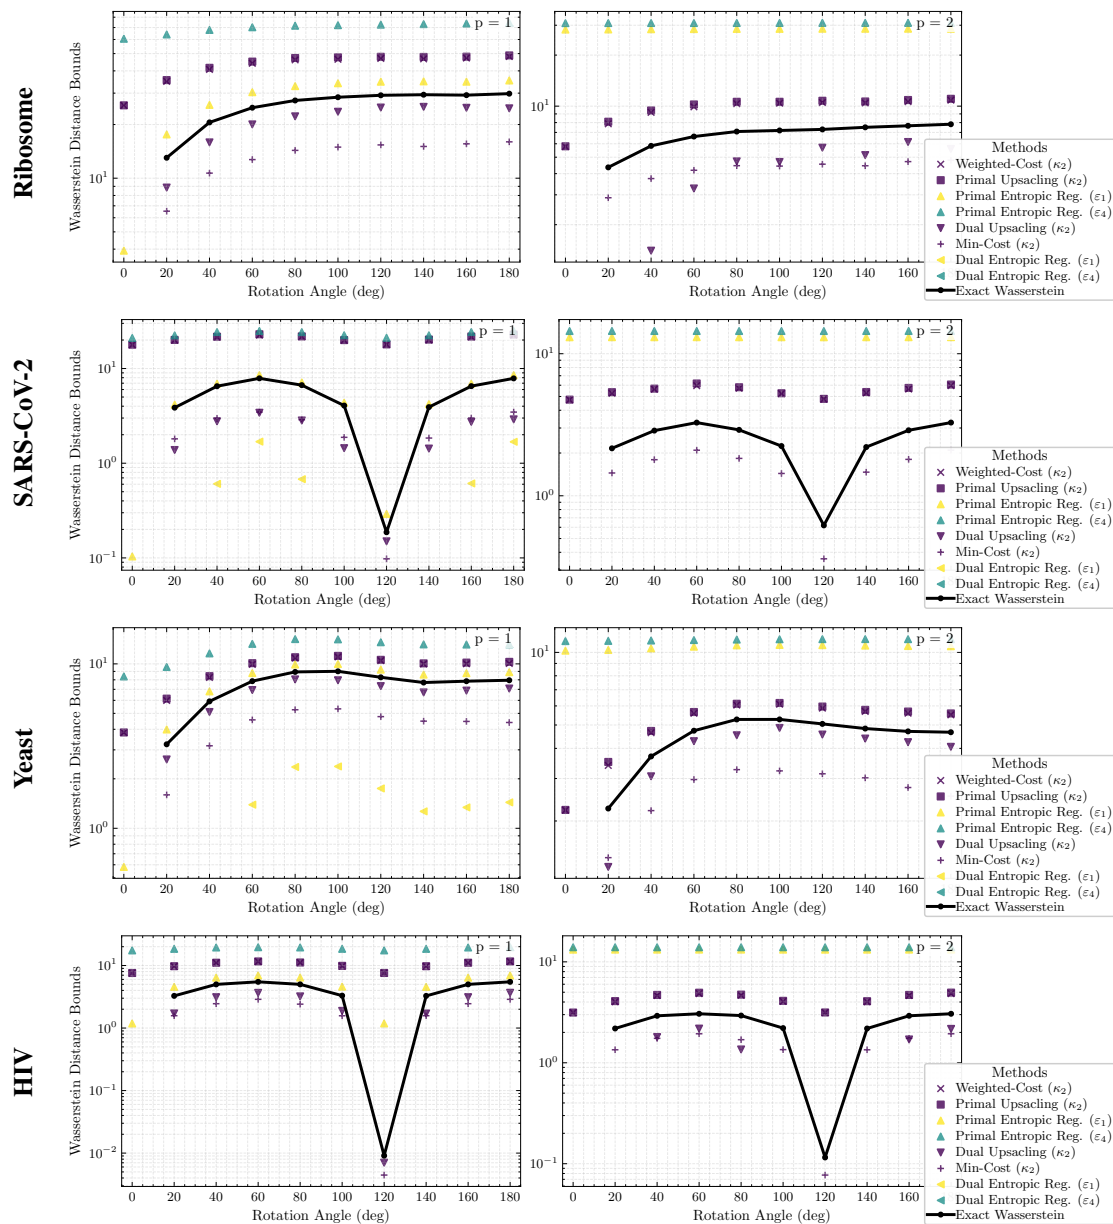


Figure 7: Wasserstein- $p$  metric and bounds between rotated 3D density maps of the molecules described in Table 5. From top to bottom: **Ribosome**, **SARS-CoV-2**, **Yeast**, **HIV**. The thick black line is the exact Wasserstein metric between the molecule and its rotated self, as a function of the rotation angle. The various upper and lower bounds are shown as different color markers. Note the drop around 120 degrees for the SARS-CoV-2 and HIV-1 spikes due to their 3-fold symmetry.

1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039  
1040  
1041  
1042  
1043  
1044  
1045  
1046  
1047  
1048  
1049  
1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057  
1058  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079

		Upper Bounds						Lower Bounds					
		Weighted-Cost		Primal Upscaling		Entropic Regularization		Dual Upscaling		Min-Cost		Entropic Regularization	
		$\kappa_2$	$\kappa_4$	$\kappa_2$	$\kappa_4$	$\varepsilon_1$	$\varepsilon_4$	$\kappa_2$	$\kappa_4$	$\kappa_2$	$\kappa_4$	$\varepsilon_1$	$\varepsilon_4$
Class	p												
Classic Images	1	3.1%	11.0%	9.6%	23.0%	<b>0.9%</b>	5.2%	<b>0.3%</b>	0.7%	10.0%	27.0%	24.0%	88.0%
		$\pm 2.0%$	$\pm 7.1%$	$\pm 4.1%$	$\pm 9.9%$	$\pm 0.5%$	$\pm 2.1%$	$\pm 0.2%$	$\pm 0.4%$	$\pm 6.4%$	$\pm 16.0%$	$\pm 8.3%$	$\pm 15.0%$
Microscopy	2	<b>1.6%</b>	7.9%	2.2%	8.8%	14.0%	44.0%	<b>0.7%</b>	2.4%	13.0%	33.0%	98.0%	100.0%
		$\pm 1.2%$	$\pm 5.2%$	$\pm 1.4%$	$\pm 5.5%$	$\pm 8.8%$	$\pm 25.0%$	$\pm 0.5%$	$\pm 1.6%$	$\pm 3.7%$	$\pm 8.8%$	$\pm 8.4%$	$\pm 0.0%$
Shapes	1	0.9%	3.4%	2.4%	6.5%	<b>0.4%</b>	2.0%	<b>0.4%</b>	0.9%	6.2%	17.0%	9.6%	38.0%
		$\pm 1.7%$	$\pm 5.9%$	$\pm 3.2%$	$\pm 8.4%$	$\pm 0.3%$	$\pm 2.0%$	$\pm 0.5%$	$\pm 0.9%$	$\pm 4.2%$	$\pm 10.0%$	$\pm 7.1%$	$\pm 22.0%$
Shapes	2	<b>0.5%</b>	2.2%	0.7%	2.7%	3.8%	11.0%	<b>0.2%</b>	0.7%	5.5%	16.0%	30.0%	90.0%
		$\pm 0.7%$	$\pm 3.5%$	$\pm 1.0%$	$\pm 3.9%$	$\pm 5.9%$	$\pm 18.0%$	$\pm 0.4%$	$\pm 1.2%$	$\pm 3.4%$	$\pm 8.4%$	$\pm 32.0%$	$\pm 19.0%$
Shapes	1	1.1%	3.6%	3.2%	7.8%	<b>0.7%</b>	2.6%	<b>0.5%</b>	1.0%	7.3%	20.0%	13.0%	51.0%
		$\pm 2.2%$	$\pm 4.8%$	$\pm 2.6%$	$\pm 6.3%$	$\pm 2.5%$	$\pm 2.2%$	$\pm 2.8%$	$\pm 3.0%$	$\pm 5.2%$	$\pm 11.0%$	$\pm 8.1%$	$\pm 22.0%$
Shapes	2	<b>1.2%</b>	3.2%	1.4%	3.6%	5.1%	15.0%	<b>0.9%</b>	1.7%	7.7%	20.0%	52.0%	99.0%
		$\pm 4.5%$	$\pm 5.0%$	$\pm 4.5%$	$\pm 5.0%$	$\pm 6.2%$	$\pm 17.0%$	$\pm 5.1%$	$\pm 5.6%$	$\pm 6.1%$	$\pm 9.2%$	$\pm 34.0%$	$\pm 9.6%$

Table 4: Accuracy comparison of different methods showing the mean  $\pm$  standard deviation of the relative error computed across all the pairwise distances in the DOTmark class at  $128 \times 128$  resolution. Each method is evaluated at different fidelity level  $\kappa_2 = 2$  and  $\kappa_4 = 4$  and different values of  $\varepsilon_1 = 1 \cdot 10^{-3} N^p$  and  $\varepsilon_4 = 4 \cdot 10^{-3} N^p$ .

Table 5: Selected cryo-EM structures from the Electron Microscopy Data Bank (EMDB).

Name	EMDB ID	Description
Ribosome	EMD-2660	Ribosome of the Plasmodium falciparum parasite which causes malaria in humans (Wong et al., 2014)
SARS-CoV-2	EMD-14621	SARS-CoV-2 spike protein (Stagnoli et al., 2022)
Yeast	EMD-8012	Yeast spliceosome (Nguyen et al., 2016)
HIV	EMD-2484	HIV-1 trimeric spike pre-fusion (Bartesaghi et al., 2013)

1080  
1081  
1082  
1083  
1084  
1085  
1086  
1087  
1088  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

Table 6: Relative computation time compared to exact OT solver of  $16 \times 16 \times 16$  downscaled EMDB density maps. Results show mean and standard deviation across rotation angles. Lower is better.

Molecule	$p$	Upper Bounds				Lower Bounds			
		Weighted-Cost	Primal Upscaling	Entropic Regularization		Dual Upscaling	Min-Cost	Entropic Regularization	
		$\kappa_2$	$\kappa_2$	$\epsilon_1$	$\epsilon_4$	$\kappa_2$	$\kappa_2$	$\epsilon_1$	$\epsilon_4$
Ribosome	1	<b>3.7%</b> $\pm 3.2\%$	175.6% $\pm 56.2\%$	156.2% $\pm 111.3\%$	155.8% $\pm 111.1\%$	74.5% $\pm 21.8\%$	<b>3.5%</b> $\pm 3.0\%$	154.1% $\pm 109.8\%$	154.1% $\pm 109.7\%$
	2	<b>12.4%</b> $\pm 3.7\%$	178.1% $\pm 226.5\%$	569.8% $\pm 252.8\%$	570.2% $\pm 252.3\%$	<b>12.2%</b> $\pm 3.6\%$	12.5% $\pm 5.0\%$	563.2% $\pm 249.8\%$	563.6% $\pm 249.6\%$
SARS-CoV-2	1	<b>11.4%</b> $\pm 26.7\%$	164.6% $\pm 96.4\%$	69.1% $\pm 4.7\%$	62.7% $\pm 15.1\%$	69.9% $\pm 35.7\%$	<b>2.2%</b> $\pm 0.3\%$	61.4% $\pm 14.0\%$	61.2% $\pm 14.8\%$
	2	<b>19.7%</b> $\pm 3.5\%$	350.5% $\pm 421.2\%$	850.9% $\pm 895.4\%$	658.8% $\pm 303.0\%$	77.9% $\pm 184.2\%$	<b>19.0%</b> $\pm 5.7\%$	642.5% $\pm 295.5\%$	641.6% $\pm 294.0\%$
Yeast	1	4.8% $\pm 4.9\%$	295.4% $\pm 379.0\%$	20.8% $\pm 31.2\%$	<b>3.1%</b> $\pm 3.5\%$	127.5% $\pm 153.4\%$	3.3% $\pm 3.3\%$	5.6% $\pm 4.6\%$	<b>0.9%</b> $\pm 0.9\%$
	2	<b>12.7%</b> $\pm 0.7\%$	172.2% $\pm 69.3\%$	919.4% $\pm 497.3\%$	828.9% $\pm 215.6\%$	13.3% $\pm 1.0\%$	<b>11.6%</b> $\pm 0.7\%$	820.4% $\pm 213.8\%$	819.7% $\pm 212.5\%$
HIV	1	<b>10.4%</b> $\pm 23.6\%$	243.1% $\pm 215.3\%$	163.1% $\pm 127.3\%$	163.0% $\pm 127.1\%$	99.6% $\pm 89.8\%$	<b>2.9%</b> $\pm 2.7\%$	161.1% $\pm 125.8\%$	161.1% $\pm 125.8\%$
	2	<b>40.7%</b> $\pm 97.7\%$	198.2% $\pm 255.3\%$	603.9% $\pm 252.2\%$	603.9% $\pm 252.0\%$	45.6% $\pm 114.3\%$	<b>40.4%</b> $\pm 99.0\%$	597.0% $\pm 249.3\%$	597.0% $\pm 249.3\%$

1134 APPENDIX C PROOFS OF THE MAIN CLAIMS

1135  
1136 APPENDIX C.1 WEIGHTED TOTAL VARIATION CORRECTION TERMS

1137  
1138 *Proof of Lemma 3.2.* Using the triangle inequality for the Wasserstein metric, we can write

1139 
$$\mathcal{W}_p(\mu, \nu) \leq \mathcal{W}_p(\mu, \hat{\mu}) + \mathcal{W}_p(\hat{\mu}, \hat{\nu}) + \mathcal{W}_p(\hat{\nu}, \nu) \tag{28}$$

1140  
1141 controlling for each element separately, we have

1142 
$$\mathcal{W}_p(\mu, \hat{\mu}) \leq \mathcal{TV}_p^w(\hat{\mu}, \mu) \quad \text{and} \quad \mathcal{W}_p(\hat{\nu}, \nu) \leq \mathcal{TV}_p^w(\nu, \hat{\nu}) \tag{29}$$

1143  
1144 by the property of weighted total variation, and

1145  
1146 
$$\mathcal{W}_p(\hat{\mu}, \hat{\nu}) = \left( \min_{\pi \in \Pi(\hat{\mu}, \hat{\nu})} \langle \pi, C \rangle \right)^{\frac{1}{p}} \leq \langle \hat{\pi}, C \rangle^{\frac{1}{p}} \tag{30}$$

1147  
1148 by evaluating the transport cost using a coupling in the problem's original space. □

1149  
1150  
1151 *Proof of Proposition 3.3.* Consider the definition of weighted total variation,

1152 
$$\mathcal{TV}_p(\hat{\mu}, \mu; w) + \mathcal{TV}_p(\nu, \hat{\nu}; w) \tag{31}$$

1153 
$$= 2^{1-\frac{1}{p}} \langle w, |\hat{\mu} - \mu| \rangle^{\frac{1}{p}} + 2^{1-\frac{1}{p}} \langle w, |\nu - \hat{\nu}| \rangle^{\frac{1}{p}}$$

1154 
$$= 2^{1-\frac{1}{p}} \left( \left( \sum \rho(\bar{x}, x)^p |\hat{\mu}_x - \mu_x| \right)^{\frac{1}{p}} + \left( \sum \rho(\bar{x}, x)^p |\nu_x - \hat{\nu}_x| \right)^{\frac{1}{p}} \right)$$

1155 
$$\leq 2^{1-\frac{1}{p}} \left( r \|\hat{\mu} - \mu\|_1^{\frac{1}{p}} + r \|\nu - \hat{\nu}\|_1^{\frac{1}{p}} \right) \tag{bounding radius}$$

1156 
$$\leq 2^{2-\frac{2}{p}} (\|\hat{\mu} - \mu\|_1 + \|\nu - \hat{\nu}\|_1)^{\frac{1}{p}} r \tag{Jensen's inequality}$$

1157 
$$< 2^{2-\frac{2}{p}} \xi^{\frac{1}{p}} r \tag{convergence criteria}$$

1162  
1163  
1164 APPENDIX C.2 WEIGHTED-COST UPPER BOUND

1165  
1166 First, let us consider the following lemma discussing a coupling constructed ad hoc using coarsened

1167 measures.  
1168 **Lemma Appendix C.1.** Let  $\mu, \nu$  measures with set of admissible couplings  $\Pi(\mu, \nu)$ , the trivial




Figure 8: Diagram showing the relationship between measures and their Wasserstein distances, and the quantities bounding each element.

1188  
1189  
1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199  
1200  
1201  
1202  
1203  
1204  
1205  
1206  
1207  
1208  
1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220  
1221  
1222  
1223  
1224  
1225  
1226  
1227  
1228  
1229  
1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1240  
1241

$$\pi_{\tilde{\pi}}(x, y) := \frac{\tilde{\pi}_{k_{\mathbb{X}}(x)\ell_{\mathbb{Y}}(y)}}{\mu(X_{k_{\mathbb{X}}(x)})\nu(Y_{\ell_{\mathbb{Y}}(y)})} \pi_{\otimes}(x, y) \quad (32)$$

is an admissible coupling  $\pi_{\tilde{\pi}} \in \Pi(\boldsymbol{\mu}, \boldsymbol{\nu})$ , where the coarsening inverse index functions  $k_{\mathbb{X}}(x)$ ,  $\ell_{\mathbb{Y}}(y)$  are defined as  $k_{\mathbb{X}}(x) := \{k : x \in X_k\}$ ,  $\ell_{\mathbb{Y}}(y) := \{\ell : y \in Y_{\ell}\}$ .

*Proof of Lemma Appendix C.1.* Following Definition 1.1 (Coupling) from (Villani, 2009), one can show  $\pi_{\tilde{\pi}}(x, y)$  Equation (32) is admissible. For  $\varphi, \psi$  be any integrable measurable functions on  $\mathcal{X}, \mathcal{Y}$  respectively, than  $\pi_{\tilde{\pi}}(x, y)$  admits

$$\begin{aligned} & \int_{\mathcal{X} \times \mathcal{Y}} (\varphi(x) + \psi(y)) d\pi_{\tilde{\pi}}(x, y) \quad (33) \\ &= \sum_{k, \ell} \int_{X_k \times Y_{\ell}} (\varphi(x) + \psi(y)) d\pi_{\tilde{\pi}}(x, y) \\ &= \sum_{k, \ell} \int_{X_k \times Y_{\ell}} (\varphi(x) + \psi(y)) \frac{\mathbf{\Pi}_{k\ell}}{\mu(X_k)\nu(Y_{\ell})} d\mu(x)d\nu(y) \quad \text{plug-in coupling's definition} \\ &= \sum_{k, \ell} \left( \frac{\tilde{\pi}_{k\ell}}{\mu(X_k)\nu(Y_{\ell})} \int_{X_k \times Y_{\ell}} \varphi(x) d\mu(x)d\nu(y) + \frac{\tilde{\pi}_{k\ell}}{\mu(X_k)\nu(Y_{\ell})} \int_{X_k \times Y_{\ell}} \psi(y) d\mu(x)d\nu(y) \right) \\ &= \sum_{k, \ell} \frac{\tilde{\pi}_{k\ell}}{\mu(X_k)} \int_{X_k} \varphi(x) d\mu(x) + \sum_{k, \ell} \frac{\tilde{\pi}_{k\ell}}{\nu(Y_{\ell})} \int_{Y_{\ell}} \psi(y) d\nu(y) \quad \text{sum over marginals} \\ &= \sum_k \int_{X_k} \varphi(x) d\mu(x) + \sum_{\ell} \int_{Y_{\ell}} \psi(y) d\nu(y) \\ &= \int_{\mathcal{X}} \varphi(x) d\mu(x) + \int_{\mathcal{Y}} \psi(y) d\nu(y) \end{aligned}$$

□

Next, we consider the transport cost of such a coupling.

**Lemma Appendix C.2.** *The transport loss assigned by the cost  $c(x, y)$  and a coupling  $\pi_{\tilde{\pi}}$  identifies with the coarse transport loss assigned by marginally weighted cost  $\bar{C}$  Equation (15) and the coarse coupling  $\tilde{\pi}$ ,*

$$\langle \pi_{\tilde{\pi}}, C \rangle = \langle \tilde{\pi}, \bar{C} \rangle \quad (34)$$

*Proof.*

$$\begin{aligned} \langle \pi_{\tilde{\pi}}, C \rangle &= \sum_{i, j} c(x_i, y_j) \pi_{\tilde{\pi}}(x_i, y_j) = \sum_{k, \ell} \sum_{\substack{x \in X_k \\ y \in Y_{\ell}}} c(x, y) \pi_{\tilde{\pi}}(x, y) \quad (35) \\ &= \sum_{k, \ell} \sum_{\substack{x \in X_k \\ y \in Y_{\ell}}} c(x, y) \frac{\tilde{\pi}_{k\ell}}{\mu(X_k)\nu(Y_{\ell})} \pi_{\otimes}(x, y) \\ &= \sum_{k, \ell} \frac{1}{\mu(X_k)\nu(Y_{\ell})} \sum_{\substack{x \in X_k \\ y \in Y_{\ell}}} c(x, y) \mu(x)\nu(y) \tilde{\pi}_{k\ell} \\ &= \sum_{k, \ell} \bar{C}_{k\ell} \tilde{\pi}_{k\ell} = \langle \tilde{\pi}, \bar{C} \rangle \end{aligned}$$

□

Finally, we can write

1242 *Proof of Theorem 3.4.* Based on admissibility of  $\pi_{\tilde{\pi}}$  shown in Lemma [Appendix C.1](#) the transport  
 1243 cost

$$1244 \langle \pi_{\tilde{\pi}}, C \rangle \geq L_C(\mu, \nu), \forall \tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu}). \quad (36)$$

1246 In particular, for  $\tilde{\pi}^* = \arg \min_{\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})} \langle \tilde{\pi}, \bar{C} \rangle$ ,

$$1247 \langle \tilde{\pi}^*, \bar{C} \rangle = \langle \pi_{\tilde{\pi}^*}, C \rangle \geq L_C(\mu, \nu) \quad (37)$$

1249 by the identity shown in Lemma [Appendix C.2](#). □

### 1251 APPENDIX C.3 MIN-COST LOWER BOUND

1253 *Proof of Theorem 3.5.* Consider

$$1254 \pi^* = \arg \min_{\pi \in \Pi(\mu, \nu)} \langle \pi, C \rangle \quad (38)$$

1257 and coarsening of the optimal coupling

$$1258 \hat{\pi}_{k\ell}^* := \sum_{\substack{x \in X_k \\ y \in Y_\ell}} \pi^*(x, y) \quad (39)$$

1261 such that,

$$1262 \begin{aligned} 1263 L_C(\mu, \nu) &= \langle \pi^*, C \rangle & (40) \\ 1264 &= \sum_{\substack{x \in \mathcal{X} \\ y \in \mathcal{Y}}} \rho(x, y)^p \pi^*(x, y) \\ 1265 &= \sum_{k, \ell} \sum_{\substack{x \in X_k \\ y \in Y_\ell}} \rho(x, y)^p \pi^*(x, y) \\ 1266 &\geq \sum_{k, \ell} C_{k\ell}^{\min} \sum_{\substack{x \in X_k \\ y \in Y_\ell}} \pi^*(x, y) \\ 1267 &= \langle \hat{\pi}^*, C^{\min} \rangle \\ 1268 &\geq \min_{\tilde{\pi} \in \Pi(\tilde{\mu}, \tilde{\nu})} \langle \tilde{\pi}, C^{\min} \rangle \\ 1269 &= L_{C^{\min}}(\tilde{\mu}, \tilde{\nu}). \end{aligned}$$

1278 □

### 1280 APPENDIX C.4 UP-SCALED COUPLING

1281 *Proof of Lemma 3.6.* Let  $\tilde{\pi}^* \in \Pi(\tilde{\mu}, \tilde{\nu})$  be the coarse optimal coupling and  $\mathbf{K}$  be the positive valued  
 1282 normalized kernel tensor satisfying  $\sum_{t \in [\kappa]^{2d}} \mathbf{K}_t = 1$ . Recall that  $\hat{\mathbf{P}} = \tilde{\mathbf{P}}^* \otimes \mathbf{K}$  and  $\hat{\pi}$  is obtained by  
 1283 reshaping  $\hat{\mathbf{P}}$ .

1285 First, we show that the sum of all elements equals 1:

$$1286 \begin{aligned} 1287 \sum_{i=1}^{N^d} \sum_{j=1}^{N^d} \hat{\pi}_{ij} &= \sum_{t \in [\kappa]^{2d}} \sum_{k=1}^{n^d} \sum_{\ell=1}^{n^d} \tilde{\pi}_{k\ell}^* \mathbf{K}_t & (41) \\ 1288 &= \sum_{k=1}^{n^d} \sum_{\ell=1}^{n^d} \tilde{\pi}_{k\ell}^* \sum_{t \in [\kappa]^{2d}} \mathbf{K}_t \\ 1289 &= \sum_{k=1}^{n^d} \sum_{\ell=1}^{n^d} \tilde{\pi}_{k\ell}^* \cdot 1 = 1 \end{aligned}$$

where the last equality follows from  $\tilde{\pi}^*$  being a coupling.

Second, we show that  $\hat{\pi}$  is non-negative. Since  $\tilde{\pi}^*$  is a coupling, it is non-negative, and  $\mathbf{K}$  is a positive-valued kernel, their tensor product  $\hat{\mathbf{P}}$  and its reshaped form  $\hat{\pi}$  are non-negative.

Thus,  $\hat{\pi}$  satisfies all the properties of a coupling measure.  $\square$

## APPENDIX D APPROXIMATION ERROR ANALYSIS

In this section, we derive the theoretical tightness rates for the proposed upper and lower bounds. Let us define covering Radius ( $r_\kappa$ ) and cell Diameter ( $\delta_\kappa$ ), the maximum distance from any point in the cell to its center and the maximum distance between any two points within the cell, respectively, for a coarse cell  $X_k$  with center  $\tilde{x}_k$ :

$$r_\kappa = \sup_{x \in X_k} \|x - \tilde{x}_k\|_2 = \frac{1}{2}\kappa\sqrt{d} \quad \delta_\kappa = \sup_{x, x' \in X_k} \|x - x'\|_2 = \kappa\sqrt{d} \quad (42)$$

We assume the underlying ground metric  $\rho(x, y)$  is  $L_\rho$ -Lipschitz with respect to the Euclidean norm. Furthermore, for the transport cost  $c(x, y) = \rho(x, y)^p$  on bounded domain  $\mathcal{X} \times \mathcal{X}$  with diameter  $D$ , for  $p \geq 1$  the cost  $c(x, y)$  is locally Lipschitz with constant  $L_c = pD^{p-1}L_\rho$ .

### APPENDIX D.1 WEIGHTED-COST UPPER BOUND

This bound approximates the transport solving the optimal transport using the expected cost between pair of coarse cells over the trivial coupling.

**Lemma Appendix D.1** (Cost Variation). *Let  $c(x, y)$  be  $L_c$ -Lipschitz. For any two coarse cells  $X_k$  and  $Y_\ell$ , and any specific points  $x \in X_k, y \in Y_\ell$ , the deviation of the weighted average cost  $\bar{C}_{k\ell}$  from the specific cost  $c(x, y)$  is bounded by the sum of the cell diameters:*

$$|\bar{C}_{k\ell} - c(x, y)| \leq 2L_c\delta_\kappa = 2L_c\kappa\sqrt{d}. \quad (43)$$

*Proof.* By the triangle inequality and Lipschitz continuity, for any  $u \in X_k, v \in Y_\ell$ :

$$|c(u, v) - c(x, y)| \leq L_c(\rho(u, x) + \rho(v, y)) \leq L_c(\delta_\kappa + \delta_\kappa).$$

Since  $\bar{C}_{k\ell} = \mathbb{E}_{u \sim \mu|_{X_k}, v \sim \nu|_{Y_\ell}}[c(u, v)]$ , the expectation is bounded by the maximum deviation.  $\square$

**Theorem Appendix D.2.** *The approximation error of the Weighted-Cost Upper Bound satisfies:*

$$\bar{\mathcal{W}}_p^\otimes(\mu, \nu) - \mathcal{W}_p(\mu, \nu) \leq 2D^{p-1}L_\rho \cdot \kappa\sqrt{d}.$$

*Proof.* Let  $\pi^*$  be the optimal coupling for the fine-scale problem. We construct a coarse coupling  $\tilde{\pi}^*$  by aggregating  $\pi^*$  such that  $\tilde{\pi}_{k\ell}^* = \sum_{x \in X_k, y \in Y_\ell} \pi^*(x, y)$ . The bound value  $\bar{L} = L_C(\tilde{\mu}, \tilde{\nu})$  minimizes the weighted-cost over coarse couplings, thus  $\bar{L} \leq \langle \tilde{\pi}^*, \bar{C} \rangle$ . The gap between  $\bar{L}$  and  $L_C(\mu, \nu)$  is bounded by:

$$\text{Gap} \leq \sum_{k, \ell} \tilde{\pi}_{k\ell}^* \bar{C}_{k\ell} - \sum_{x, y} \pi^*(x, y) c(x, y) = \sum_{k, \ell} \sum_{x \in X_k, y \in Y_\ell} \pi^*(x, y) (\bar{C}_{k\ell} - c(x, y)).$$

Applying Lemma Appendix D.1, the term in parentheses is bounded by  $2L_c\kappa\sqrt{d}$ . Summing over the probability mass yields the result, assuming  $\bar{L} \geq L_C \geq 1$ :

$$\bar{\mathcal{W}}_p^\otimes - \mathcal{W}_p = \bar{L}_C^{1/p} - L_C^{1/p} \quad (44)$$

$$\leq \frac{1}{p}(\bar{L}_C - L_C) \quad (45)$$

$$\leq \frac{1}{p}(L_c\kappa\sqrt{d}). \quad (46)$$

$\square$

1350 APPENDIX D.2 MIN-COST LOWER BOUND

1351  
1352 **Theorem Appendix D.3.** *The approximation error of the Min-Cost Lower Bound satisfies:*

1353 
$$\mathcal{W}_p(\mu, \nu) - \underline{\mathcal{W}}_p^{\min}(\mu, \nu) \leq 2D^{p-1}L_\rho \cdot \kappa\sqrt{d}.$$

1354  
1355  
1356 *Proof.* Using the same aggregated coupling  $\tilde{\pi}^*$  as above, the lower bound value is at most  $\langle \tilde{\pi}^*, C^{\min} \rangle$ .  
1357 The error is bounded by the maximum deviation between the actual fine cost  $c(x, y)$  and the optimistic  
1358 coarse cost  $C_{k\ell}^{\min} = \min_{u \in X_k, v \in Y_\ell} c(u, v)$ . Let  $(u^*, v^*)$  be the minimizers for the cell pair. For any  
1359  $x \in X_k, y \in Y_\ell$ :

1360 
$$c(x, y) - c(u^*, v^*) \leq L_c(\rho(x, u^*) + \rho(y, v^*)) \leq 2L_c\delta_\kappa.$$

1361 This holds for all transported mass, yielding a total error of  $2L_c\kappa\sqrt{d}$ . □

1362  
1363 APPENDIX D.3 PRIMAL UPSCALING UPPER BOUND

1364  
1365 **Theorem Appendix D.4.** *Neglecting the total variation correction terms (assuming  $\xi \rightarrow 0$ ), the*  
1366 *error of the Primal Upscaling Upper Bound satisfies:*

1367 
$$\overline{\mathcal{W}}_p(\mu, \nu) - \mathcal{W}_p(\mu, \nu) \leq L_\rho \cdot \kappa\sqrt{d}.$$

1368  
1369  
1370 *Proof.* The primal upscaling method constructs a solution based on the optimal coarse plan. Its cost  
1371 is equivalent to  $\mathcal{W}_p(\tilde{\mu}, \tilde{\nu})$  where  $\tilde{\mu}, \tilde{\nu}$  are the coarsened measures supported on cell centers. By the  
1372 triangle inequality of the Wasserstein metric:

1373 
$$\mathcal{W}_p(\mu, \nu) \leq \mathcal{W}_p(\tilde{\mu}, \tilde{\nu}) + \mathcal{W}_p(\mu, \tilde{\mu}) + \mathcal{W}_p(\nu, \tilde{\nu}).$$

1374  
1375 The term  $\mathcal{W}_p(\mu, \tilde{\mu})$  represents the quantization error of moving mass from the fine grid  $x \in X_k$  to  
1376 the center  $\tilde{x}_k$ . The displacement is bounded by the covering radius  $r_\kappa$ .

1377 
$$\mathcal{W}_p(\mu, \tilde{\mu}) \leq L_\rho r_\kappa = L_\rho \frac{\kappa}{2} \sqrt{d}.$$

1378  
1379 Summing the quantization errors for  $\mu$  and  $\nu$  yields  $L_\rho\kappa\sqrt{d}$ . □

1380  
1381 APPENDIX D.4 DUAL UPSCALING LOWER BOUND

1382  
1383 **Theorem Appendix D.5.** *Let the optimal dual potential  $f$  be  $L_c$ -Lipschitz. Assuming the optimal*  
1384 *transport cost is bounded away from zero such that  $L_C(\mu, \nu) \geq 1$ , the approximation error of the*  
1385 *Wasserstein distance satisfies:*

1386 
$$\mathcal{W}_p(\mu, \nu) - \underline{\mathcal{W}}_p(\mu, \nu) \leq 2D^{p-1}L_\rho \cdot \kappa\sqrt{d}.$$

1387  
1388  
1389 *Proof.* Let  $L_C(\mu, \nu) = \mathcal{W}_p^p(\mu, \nu)$  be the exact optimal transport cost, and  $\underline{L}_C = \underline{\mathcal{W}}_p^p(\mu, \nu)$  be the  
1390 approximate lower bound cost.

1391 First, we bound the error in the cost domain. The error is driven by the interpolation of the dual  
1392 potentials. Let  $f^*$  be the optimal potential on the fine grid.  $f^*$  is  $c$ -concave and thus inherits the  
1393 Lipschitz constant  $L_c$  from the cost function (Villani, 2009). The method approximates  $f^*$  using  $\hat{f}$ , a  
1394 multilinear interpolation of the coarse grid values. For an  $L_c$ -Lipschitz function, the sup-norm error  
1395 of linear interpolation on a cell of diameter  $\delta_\kappa$  is bounded by the variation of the function over the  
1396 cell:

1397 
$$\|f^*(x) - \hat{f}(x)\|_\infty \leq L_c \cdot \delta_\kappa.$$

1398 Since the dual objective is linear in the potentials for each measure, the gap is bounded by  $L_C - \underline{L}_C \leq$   
1399  $2L_c\delta_\kappa = 2L_c\kappa\sqrt{d}$ . □

1400  
1401  
1402  
1403

1404 APPENDIX E TABLE OF NOTATIONS

1405

1406

1407

Table 7: Table of Notations

1408

1409

Notation	Category	Description
<i>Sets and Spaces</i>		
$\mathbb{R}_+$	Set	Non-negative real numbers
$[n]$	Set	Set of integers $\{1, \dots, n\}$
$\Sigma_N$	Space	Probability simplex $\{(p_1, \dots, p_N) \in \mathbb{R}_+^N : \sum_i p_i = 1\}$
$\mathcal{X}, \mathcal{Y}$	Set	Point sets where measures are defined
$\mathbb{X}$	Set	Set of non-overlapping hypercubes covering the grid
$\tilde{\mathcal{X}}, \tilde{\mathcal{Y}}$	Set	Coarse grids (set of hypercube centers)
$X_k, Y_\ell$	Set	Individual hypercubes in the partition
<i>Measures and Vectors</i>		
$\mathbf{0}_n, \mathbf{1}_n$	Vector	All-zeros and all-ones vectors in $\mathbb{R}^n$ , respectively
$\mu, \nu$	Measure	Discrete probability measures
$\boldsymbol{\mu}, \boldsymbol{\nu}$	Vector	Vector representations of measures $\mu, \nu$
$\tilde{\mu}, \tilde{\nu}$	Measure	Coarsened measures
$\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{\nu}}$	Vector	Vector representations of coarsened measures
$\mathbf{f}, \mathbf{g}$	Vector	Kantorovich potentials
$\mathbf{a}, \mathbf{b}$	Vector	Sinkhorn scaling vectors
<i>Matrices and Tensors</i>		
$C$	Matrix	Ground-cost matrix
$\tilde{C}$	Matrix	Coarse cost matrix (center-based)
$\bar{C}$	Matrix	Coarse cost matrix (average-based)
$\pi$	Matrix	Transport plan (coupling matrix)
$\pi^*$	Matrix	Optimal transport coupling
$\tilde{\pi}$	Matrix	Coarse coupling
$\mathbf{K}$	Tensor	Normalized kernel tensor
<i>Functions and Operations</i>		
$\rho$	Function	Distance metric
$\mathcal{W}_p$	Function	Wasserstein- $p$ metric
$L_C$	Function	Optimal transport cost for ground-cost $C$
$\mathcal{TV}_p^w$	Function	Weighted total variation
$\otimes$	Operation	Tensor product
$\odot, \oslash$	Operation	Pointwise multiplication, division
$\langle \cdot, \cdot \rangle$	Operation	Standard vector/matrix inner product
<i>Parameters and Constants</i>		
$p$	Scalar	Order of Wasserstein metric ( $p \geq 1$ )
$\kappa$	Scalar	Scale factor for coarsening
$\xi$	Scalar	Convergence threshold for fitting
$N$	Scalar	Side length of regular grid
$d$	Scalar	Dimension of the space
$n$	Scalar	Side length of coarse grid ( $n = N/\kappa$ )
$\#S$	Scalar	Cardinality (size) of set $S$
$r$	Scalar	Radius of space $\mathcal{X}$
$\Delta_{\tilde{\mu}}, \Delta_{\tilde{\nu}}$	Scalar	Marginal weighed total variation corrections
<i>Code</i>		
AvgPool, SumPool	Function	Average/sum pooling layer with identical kernel size and stride
mean	Function	Mean of a set of points
reshape	Function	Cardinality preserving tensor shape transformation

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457