

Approximating the Earth Mover’s Distance between sets of geometric objects*

Marc van Kreveld Frank Staals Amir Vaxman Jordi L. Vermeulen

Abstract

Given two distributions P and S of equal total mass, the Earth Mover’s Distance measures the cost of transforming one distribution into the other, where the cost of moving a unit of mass is equal to the distance over which it is moved.

We give approximation algorithms for the Earth Mover’s Distance between various sets of geometric objects. We give a $(1 + \varepsilon)$ -approximation when P is a set of weighted points and S is a set of line segments, triangles or d -dimensional simplices. When P and S are both sets of line segments, sets of triangles or sets of simplices, we give a $(1 + \varepsilon)$ -approximation with a small additive term. All algorithms run in time polynomial in the size of P and S , and actually calculate the transport plan (that is, a specification of how to move the mass), rather than just the cost. To our knowledge, these are the first combinatorial algorithms with a provable approximation ratio for the Earth Mover’s Distance when the objects are continuous rather than discrete points.

1 Introduction

The Earth Mover’s Distance (EMD) is a metric that is widely used in fields such as image retrieval [17], shape matching [7, 13, 20] and mesh reconstruction [4]. It models two sets P and S as distributions of mass, and takes their distance $\mathcal{D}(P, S)$ to be the minimum cost of transforming one distribution into the other, where cost is measured by the amount of mass moved multiplied by the distance over which it is moved. More formally,

$$\mathcal{D}(P, S) = \inf_{\mu \in M} \int_P \int_S d(p, s) \cdot \mu(p, s) \, dp \, ds$$

where M is the set of all mappings of mass between P and S and $d(\cdot, \cdot)$ is any metric. In the case where P and S are sets of (weighted) points, we can rewrite this as

$$\mathcal{D}(P, S) = \min_{\mu \in M} \sum_{p \in P} \sum_{s \in S} d(p, s) \cdot \mu(p, s)$$

For unweighted point sets, the solution can be obtained by solving an assignment problem; for weighted point sets, this is an instance of a minimum cost flow problem.

Recently, much attention has been devoted to computing the Earth Mover’s Distance when both P and S are sets of points [1, 5, 10, 18]. In this paper we expand on this by letting P and S be sets of points, line segments, triangles or d -dimensional simplices in \mathbb{R}^d . We describe a unified framework for calculating the EMD between points and segments, points and triangles, points and simplices, segments and segments, triangles and triangles and simplices and simplices. Our approach provides polynomial-time algorithms that give a

*This research was supported by NWO TOP grant no. 612.001.65.

$(1 + \varepsilon)$ -approximation to the Earth Mover’s Distance between P and S , for some arbitrarily small $\varepsilon > 0$. Moreover, our algorithms produce an assignment of mass that realises this cost. For triangles and simplices, the running time also depends on the largest edge length (note that we normalise the total area/volume of each set to one, so we cannot improve the running time by scaling the input). When neither set contains points, there is a small extra additive term in our approximation. For all our algorithms, our approach is to subdivide the elements of the input into sufficiently small pieces, and then approximate each piece by a point. The approximate optimal transport plan can then be obtained by solving a transport problem on these points. Our results are summarised in Table 1. Note that all our algorithms give the solution with high probability; this is simply a consequence of using Fox and Lu’s algorithm [5] to solve the optimal transport problem on points. Substituting a deterministic algorithm here would make our results deterministic as well.

To our knowledge, these are the first combinatorial algorithms with a provable approximation ratio for the Earth Mover’s Distance when the objects are continuous rather than discrete points. We give algorithms for moving mass from points to segments (Section 5), points to triangles (Section 6), points to simplices (Section 9.1), segments to segments (Section 7), triangles to triangles (Section 8) and simplices to simplices (Section 9.2).

2 Related work

The general problem of optimally moving a distribution of mass was first described by Monge in 1781 [16], and was reformulated by Kantorovich in 1942 [9]. It is known as the Earth Mover’s Distance due to the analogy of moving piles of dirt around; it is also known as the 1-Wasserstein distance, and is a special case of the more general optimal transport problem. For a full treatment of the problem’s history and connections to other areas of mathematics, the reader is referred to Villani’s book [21].

The Earth Mover’s Distance has been studied in many geometric contexts. Agarwal et al. [1] give both exact and approximation algorithms for the case where both sets are points under some L_p metric. When both sets are weighted points, Khesin et al. [10] give two $(1 + \varepsilon)$ -approximation algorithms running in $O(n\varepsilon^{-O(d)} \log(\Delta)^{O(d)} \log n)$ and $O(n\varepsilon^{-O(d)} \log(U^{O(d)} \log(n)^2))$ time, where d is the dimension, Δ the aspect ratio of the input, and U the total mass. However, their algorithm assumes that the point weights are integers, whereas our weights can be arbitrary real numbers, as they correspond to

Objects	Running time	Additive term
Points to segments	$O\left(\frac{nm}{\varepsilon^c} \text{polylog} \frac{nm}{\varepsilon}\right)$	-
Points to triangles	$O\left(\frac{nm}{\varepsilon^c} \text{polylog} \frac{nm\Delta}{\varepsilon}\right)$	-
Points to simplices	$O\left(6^d d^2 d^d m + \frac{105^d d^2 d^{d/2} nm}{\varepsilon^{d+c}} \text{polylog}\left(\frac{d nm \Delta}{\varepsilon^d}\right)\right)$	-
Segments to segments	$O\left(\frac{nm}{\varepsilon^c} \text{polylog} \frac{nm}{\varepsilon}\right)$	$O\left(\frac{\varepsilon}{nm}\right)$
Triangles to triangles	$O\left(\frac{nm\Delta(n+m)}{\varepsilon^c} \text{polylog} \frac{nm\Delta}{\varepsilon}\right)$	$O\left(\frac{\varepsilon}{\sqrt{nm}}\right)$
Simplices to simplices	$O\left(\frac{\sqrt{d}(nm)^{1/d} \Delta^d (n+m)}{\varepsilon^c} \text{polylog}\left(\frac{d(nm)^{1/d} \Delta^d}{\varepsilon}\right)\right)$	$O\left(\frac{\sqrt{d\varepsilon}}{(nm)^{1/d}}\right)$

Table 1: A summary of our results for different choices of sets P and S of sizes n and m . d is the dimension, Δ is the largest diameter of any element of the sets, and c is some constant.

lengths and areas. This result was improved by Fox and Lu [5], using a similar method to obtain, with high probability, a $(1 + \varepsilon)$ -approximation in $O(n\varepsilon^{-O(d)} \log^{O(d)} n)$ time. The EMD was also studied when the input sets may be transformed: Cabello et al. [2] give a $(2 + \varepsilon)$ -approximation to minimising the EMD between two point sets under rigid transformations.

For continuous distributions, rather than discrete point sets, many numerical algorithms are known (see e.g. De Goes et al. [3], Lavenant et al. [12], Mérigot [14, 15] and Solomon et al. [19]). For the case where one set contains weighted points and the other is a bounded set $C \subset \mathbb{R}^d$, Geiß et al. [6] give a geometric proof that there exists an additively weighted Voronoi diagram such that transporting mass from each point p to the part of C contained in its Voronoi cell is optimal. The weights of this Voronoi diagram can be determined numerically.

De Goes et al. [4] discuss a problem similar to our own, but in the context of the reconstruction and simplification of 2D shapes. Given a set of points, they want to reconstruct a simplicial complex of a given number of vertices that closely represents the shape of the point set. They start with computing the Delaunay triangulation of the point set, then iteratively collapse the edge that minimises the increase in the EMD between the point set and the triangulation. They use a variant of the EMD in which the cost is proportional to the square of the distance (2-Wasserstein distance). This allows them to calculate this variant of the EMD between a given set of points and a given edge of the triangulation exactly, as the squared distance can be decomposed into a normal and a tangential component. However, they determine the assignment of points to edges heuristically. In this work, we show how to obtain a $(1 + \varepsilon)$ -approximation to the true optimal solution.

3 Preliminaries

We are given a set of points $P = \{p_1, \dots, p_n\}$ in the plane with weights $\{w_1, \dots, w_n\}$ and a set of geometric objects $S = \{s_1, \dots, s_m\}$, with lengths, areas or volumes $\{l_1, \dots, l_m\}$. It is given that $\sum w_i = \sum l_j$. We assume the mass associated with an object s_i is distributed uniformly over the object, and that all objects have the same mass density. For convenience, and without loss of generality, we scale the input such that the total mass in either set is one. We want to compute a “transport plan” of mass from P to S that minimises the cost according to the Earth Mover’s Distance. We define for each pair $(p_i, s_j) \in P \times S$ a function $\mu_{i,j}(x, y)$, that describes the density of mass being moved from p_i to the point $(x, y) \in s_j$. All these functions together describe the function μ used in the definition of $\mathcal{D}(A, B)$. Such a set of functions needs to satisfy the following conditions to be a valid transport plan:

$$\begin{aligned} \forall i, j : \quad & 0 \leq \mu_{i,j}(x, y) \leq 1 \\ \forall i : \quad & \sum_{j=1}^m \int_{s_j} \mu_{i,j}(x, y) \, dt = w_i \\ \forall j, (x, y) \in s_j : \quad & \sum_{i=1}^n \mu_{i,j}(x, y) = 1 \end{aligned}$$

We can then define the cost of a given transport plan μ as

$$|\mu| = \sum_{i=1}^n \sum_{j=1}^m \int_{s_j} \mu_{i,j}(x, y) \cdot d(p_i, (x, y)) \, dt$$




Figure 1: The regions of points for one subsegment q_j , and the parts of the segment they assign their mass to.

where $d(\cdot, \cdot)$ is any metric. Our problem is to find a transport plan μ^* with minimal cost.

In the following section, we give an exact algorithm to calculate an optimal transport plan between a set of weighted points and line segments when $d(\cdot, \cdot)$ is the L_1 metric. However, the approach we use does not seem to generalise to Euclidean distances, objects with areas, or even two sets of segments. This motivated us to look towards approximation algorithms for more general versions of the EMD problem. In the rest of this paper, we describe approximation algorithms and only consider the case where $d(\cdot, \cdot)$ is the L_2 metric.

4 Points to segments under the L_1 metric

When S is a set of line segments and distances are measured by the L_1 metric, we can solve the problem exactly by a convex quadratic program. We first subdivide all segments on the x - and y -coordinates of the points; call the set of subdivided segments $Q = (q_1, \dots, q_k)$. Note that the horizontal and vertical strip induced by each segment is now empty of points, while the axis-aligned quadrants starting at each corner of its bounding box may contain points; see Figure 1 for an illustration. Let Q_1 be the set of segments in Q with slope between -1 and 1 , and let Q_2 be $Q \setminus Q_1$.

We can now label the regions of points for each segment: let X_{j1} and X_{j2} be the regions to the left of q_j , with X_{j1} being the region starting at the leftmost endpoint of q_j , and X_{j2} being the other. Similarly, let X_{j3} be the region starting at the rightmost endpoint of q_j , and let X_{j4} be the other region on the right. In case of a horizontal or vertical segment, X_{j2} and X_{j4} are simply merged into X_{j1} and X_{j3} , and it does not matter if X_{j1} is the top or bottom region.

For all points the distance to any point on the a segment q_j is the same as the distance via one of the corners of the axis-aligned bounding box of q_j . Therefore, for each region, we can separately consider the cost to reach the bounding box of q_j with a certain amount of mass, and the cost to spread that mass out over the segment. Furthermore, the order in which a segment with a given slope receives mass from the different regions in an optimal solution is fixed depending on its slope.

Let u_{ij} be variables containing the amount of mass moved from p_i to q_j , let d_{ij} be the precomputed distance from p_i to the bounding box of q_j , let W_j and H_j be the width and height of the bounding box of q_j , and let w_j and h_j be constants such that $w_j \cdot \ell$ is the




Figure 2: The calculations of the distances to the midpoints of the regions to which the mass from X_{j1} and X_{j2} will be assigned for a segment with slope between -1 and 1 .

difference in x -coordinate when moving a distance of ℓ along segment q_j , and $h_j \cdot \ell$ is the difference in y -coordinate. Writing $\sum_{i \in X_{j1}} u_{ij}$ as x_{j1} for convenience, for a given segment $q_j \in Q_1$ we can write the cost c_{jk} for moving the mass from the corner of the region X_{jk} to the segment as follows:

$$\begin{aligned}
c_{j1} &= \frac{1}{2}x_{j1}^2(w_j + h_j) \\
c_{j2} &= x_{j2}(w_j x_{j1} + w_j \frac{1}{2}x_{j2} + (H_j - h_j x_{j1} - h_j \frac{1}{2}x_{j2})) = x_{j2}((w_j - h_j)(x_{j1} + \frac{1}{2}x_{j2}) + H_j) \\
c_{j3} &= \frac{1}{2}x_{j3}^2(w_j + h_j) \\
c_{j4} &= x_{j4}(w_j x_{j3} + w_j \frac{1}{2}x_{j4} + (H_j - h_j x_{j3} - h_j \frac{1}{2}x_{j4})) = x_{j4}((w_j - h_j)(x_{j3} + \frac{1}{2}x_{j4}) + H_j)
\end{aligned}$$

Here we use the fact that under the L_1 distance, the cost of sending mass to some connected region of a segment is the same as the cost of sending everything to the midpoint of the connected region; see Figure 2 for an illustration of the calculations of the distances to these midpoints. Note that we omit the cost of sending the mass from the points to the corners of the bounding box; this will be accounted for later. Further note that $w_j - h_j$ is always positive here. Symmetrically, the costs $c'_k(q_j)$ for a given segment $q_j \in Q_2$ are as follows:

$$\begin{aligned}
c'_{j1} &= \frac{1}{2}x_{j1}^2(w_j + h_j) \\
c'_{j2} &= x_{j2}((h_j - w_j)(x_{j3} + \frac{1}{2}x_{j2}) + W_j) \\
c'_{j3} &= \frac{1}{2}x_{j3}^2(w_j + h_j) \\
c'_{j4} &= x_{j4}((h_j - w_j)(x_{j1} + \frac{1}{2}x_{j4}) + W_j)
\end{aligned}$$

Note that $h_j - w_j$ is always positive here. We can now formalise our problem as follows:

$$\begin{aligned}
\mu^* &= \arg \min_{\mathbf{u}} \left(\sum_j \sum_i d_{ij} \cdot u_{ij} \right) + \sum_k \left(\sum_{j \in Q_1} c_{jk} + \sum_{j \in Q_2} c'_{jk} \right) \\
\text{subject to } & u_{ij} \geq 0 && \forall i, j \\
& \sum_j u_{ij} = w_i && \forall i \\
& \sum_i u_{ij} = |q_j| && \forall j
\end{aligned}$$




Figure 3: An example subdivision of a set of segments. Small perpendicular line segments delimit the generated subsegments. A green circle denotes the distance of δ/nm from each point. Note that ε is set to a very large value here for the clarity of the resulting image.

Since all d_{ij} , c_{jk} and c'_{jk} are non-negative, this is a quadratic program with a convex objective function, and as such it can be solved in (weakly) polynomial time [11].

Theorem 1. *Let P be a set of n weighted points and S be a set of m line segments with equal total weight. It is possible to construct an exact optimal transport plan between P and S under the L_1 metric in weakly polynomial time.*

5 Points to segments

We now describe a polynomial-time algorithm that finds a transport plan with a cost that is at most $1 + \varepsilon$ times the cost of the optimal transport plan when one set contains points and the other contains line segments. The main idea is to reduce our instance to a transport problem on two weighted sets of points. Our strategy is as follows: we subdivide each segment such that for each subsegment s' the ratio of the distance to the closest and furthest point on s' for every $p_i \in P$ is at most $1 + \delta$ for some appropriate choice of $\delta \in O(\varepsilon)$. We then approximate a minimum cost flow problem on a bipartite graph between P and the subdivided segments, where the cost of any edge is equal to the shortest distance between a point and a subsegment. Finally, we use the solution to this flow problem to build a discrete transport plan. For an appropriate choice of δ , this gives a $(1 + \varepsilon)$ -approximation.

The naive approach to subdividing the segments would be to make all the pieces some equal, appropriately small length. However, we can reduce the number of subsegments required by subdividing the segments as follows*. We repeatedly perform the following procedure for each subsegment. If there exists a point in P such that the entire subsegment lies within distance δ/nm of that point, do nothing. Otherwise, if there is a point in P for which the ratio of the longest and shortest distance between that point and the current subsegment is more than $1 + \delta$, cut the subsegment in half. Call the resulting set of subsegments Q ; see Figure 3 for an example.

We now define a complete bipartite graph $G = (P \cup Q, P \times Q)$, with edges between each point-subsegment pair (note that this graph is used for analysis only; our algorithm does not construct it). The cost of each edge will simply be the shortest distance between the point and segment it connects. A solution to a flow problem in G can be transformed into a transport plan by assigning a piece of subsegment to a point with length equal to the amount of flow along the corresponding edge. We will show that the EMD between P and S is approximated by the cost of any transport plan derived from a minimum cost flow in G .

*This reduces the total number of subsegments required from $O(nm/\varepsilon^2)$ to $O(\frac{nm}{\varepsilon} \log \frac{1}{\varepsilon})$.

First note the following general lower bound on the cost of an optimal solution:

Lemma 2. *The Earth Mover's Distance $|\mu^*|$ between P and Q is bounded from below by the cost $|\mathcal{W}|$ of a minimum cost flow \mathcal{W} in G .*

Proof. Consider any transport plan μ^* that minimises the Earth Mover's Distance. If instead of spreading the mass moved from each point p_i equally over the whole segment, we move all the mass to the point on the segment closest to p_i , we obtain a plan λ with cost $|\lambda| \leq |\mu^*|$. Such a plan is a solution to a flow problem in G , as it moves all available mass. It follows that the cost $|\mathcal{W}|$ of a minimum cost flow \mathcal{W} in G satisfies $|\mathcal{W}| \leq |\mu^*|$. \square

We also note the following lower bound on the value of $|\mathcal{W}|$:

Lemma 3. $|\mathcal{W}| \geq \frac{\delta - 2\delta^2 - 2\delta^3}{nm}$.

Proof. For a given point-segment pair $(p, s) \in P \times S$, consider the segments in Q derived from s that have a point within distance δ/nm of p . By construction, such a segment has its furthest point at distance at most $(1 + \delta) \cdot \delta/nm = \delta/nm + \delta^2/nm$. Therefore, the total length of these segments is at most $2(\delta/nm + \delta^2/nm)$ for a given p and s . Over all point-segment pairs, this gives a total length of at most $2\delta + 2\delta^2$. This means the total length of segments in Q with distance at least δ/nm is at least $1 - 2\delta - 2\delta^2$. The cost of a minimum flow in \mathcal{W} is therefore at least $(1 - 2\delta - 2\delta^2) \cdot \delta/nm = (\delta - 2\delta^2 - 2\delta^3)/nm$. \square

We calculate a transport plan μ between P and Q as follows. First, we approximate each segment $q \in Q$ by a point somewhere on that segment with weight equal to the length of q ; call this set of points T . We obtain μ by calculating an optimal transport plan ν between P and T , and then spreading the mass sent to each point $t \in T$ evenly over the segment in Q that point was derived from. We now bound the cost of μ in terms of $|\mathcal{W}|$:

Lemma 4. $|\mathcal{W}| \leq |\mu| \leq (1 + \delta)^2 |\mathcal{W}| + \frac{4\delta^2}{nm} + \frac{2\delta^3}{nm}$.

Proof. We first bound the cost of ν . In \mathcal{W} , we measured all the distances to the closest point on each subsegment. Imagine that we picked all the points in T to be the furthest point on the subsegment. For the subsegments with furthest distance at least δ/nm , the ratio of these distances is at most $1 + \delta$ by construction. We can therefore bound all the parts of ν where the furthest distance is at least δ/nm by $(1 + \delta)|\mathcal{W}|$. The total mass being moved over distance at most δ/nm in ν is at most 2δ , giving a cost of at most $2\delta^2/nm$. The total cost of ν is therefore at most $(1 + \delta)|\mathcal{W}| + 2\delta^2/nm$.

Now consider the extra cost incurred when transforming ν into μ by spreading the mass out evenly over all the segments. We can use the same argument as before: for the parts of ν with distance at least δ/nm , the cost increases by a factor of at most $1 + \delta$, and the total cost of the part within distance δ/nm is at most $2\delta^2/nm$. We can therefore bound the cost of μ by $(1 + \delta)|\nu| + 2\delta^2/nm$.

We now obtain the upper bound stated in the lemma by plugging the bound on ν into the bound on μ . The lower bound follows directly from the fact that none of the distances in μ are smaller than the distances between the same objects in \mathcal{W} . \square

We now show that $|\mu|$ approximates $|\mu^*|$.

Theorem 5. $|\mu|$ is a $(1 + 17\delta)$ -approximation to the Earth Mover's Distance $|\mu^*|$ between P and S for $0 < \delta \leq \frac{1}{4}$.




Figure 4: An example of a single R_{ij} . Small perpendicular segments delimit the generated subsegments. The green circle denotes the distance of δ/nm from the projected point.

Proof. By Lemma 4, we know that

$$|\mu| \leq (1 + \delta)^2 |\mathcal{W}| + \frac{4\delta^2}{nm} + \frac{2\delta^3}{nm}$$

$|\mathcal{W}|$ is also a lower bound on $|\mu|$; the ratio between the upper and lower bound is

$$\frac{(1 + \delta)^2 |\mathcal{W}| + \frac{4\delta^2}{nm} + \frac{2\delta^3}{nm}}{|\mathcal{W}|}$$

This ratio is the largest for small values of $|\mathcal{W}|$, so we plug in the lower bound from Lemma 3:

$$\begin{aligned} & \frac{(1 + \delta)^2 \cdot \frac{\delta - 2\delta^2 - 2\delta^3}{nm} + \frac{4\delta^2}{nm} + \frac{2\delta^3}{nm}}{\frac{\delta - 2\delta^2 - 2\delta^3}{nm}} \\ &= \frac{1 + 4\delta - 3\delta^2 - 6\delta^3 - 2\delta^4}{1 - 2\delta - 2\delta^2} \\ &= 1 + \delta + \frac{5\delta + \delta^2 - 4\delta^3 - 2\delta^4}{1 - 2\delta - 2\delta^2} \\ &\leq 1 + \delta + \frac{6\delta}{1 - 2\delta - 2\delta^2} \\ &\leq 1 + 17\delta \end{aligned} \quad (\text{assuming } \delta \leq \frac{1}{4})$$

As $|\mathcal{W}|$ is also a lower bound for $|\mu^*|$ (Lemma 2), and μ can obviously not have lower cost than the optimal transport plan, this gives a $(1 + 17\delta)$ -approximation. \square

Setting $\delta = \varepsilon/17$ gives a $(1 + \varepsilon)$ -approximation. Note that the bound on δ is not restrictive: for any constant ε that would require a larger value of δ , we can simply use the value $1/4$ at the cost of a constant factor in the running time of our algorithm.

5.1 Running time analysis

We now analyse the number of subsegments in Q . We will count the number of subsegments in a different subdivision of S , and then show that Q has at most a constant factor more subsegments. The alternative subdivision of each s_j will be as follows: project each p_i onto the supporting line of s_j , call this point p_{ij} . We construct the one-dimensional Voronoi diagram of all p_{ij} along the supporting line of s_j ; let s_{ij} be the part of s_j inside the Voronoi cell of p_{ij} . From each p_{ij} , we subdivide s_{ij} into both directions. Up to a distance of δ/nm , we make subsegments of size δ^2/nm . Moving outward, we double the size of the subsegments whenever their ratio of distances to p_{ij} would still be below $1 + \delta$. Let R_{ij} be the resulting subdivision; see Figure 4 for an example.

Lemma 6. $R = \bigcup R_{ij}$ has $O\left(\frac{nm}{\delta} \log \frac{1}{\delta}\right)$ subsegments.

Proof. We define $\beta = \frac{\delta}{nm}$ and $\gamma = \frac{\delta^2}{nm}$. In the following, we only analyse the case where p_{ij} is on s_{ij} ; if it lies outside, the number of subsegments will be smaller, as the size of the subsegments increases with distance. The length covered as we add subsegments on s_{ij} can be written as

$$\beta + 2 \sum_{i=0}^k \alpha_i 2^i \gamma$$

where k is the number of times we double the size of the subsegments, and α_i is the number of subsegments with a size that has been doubled i times. The number of subsegments can then be calculated by finding the values of k and α_i . We start with α_0 , which can be found by considering the distance at which the next cell could be double the size:

$$\begin{aligned} \frac{\beta + \alpha_0 \gamma + 2\gamma}{\beta + \alpha_0 \gamma} &\leq 1 + \delta \\ \frac{2\gamma}{\beta + \alpha_0 \gamma} &\leq \delta \\ \alpha_0 &\geq \frac{2\gamma - \delta\beta}{\delta\gamma} \\ &= \frac{2}{\delta} - \frac{\beta}{\gamma} \\ &= \frac{1}{\delta} \end{aligned}$$

We want to take the values of α_i as small as possible, so we take $\alpha_0 = \frac{1}{\delta}$. Next, we can show by induction that all α_i are equal:

$$\text{IH: } \alpha_j = \alpha^* = \frac{1}{\delta} \text{ for } j < i.$$

$$\begin{aligned}
\frac{\beta + 2^{i+1}\gamma + \sum_{j=0}^i \alpha_j 2^j \gamma}{\beta + \sum_{j=0}^i \alpha_j 2^j \gamma} &\leq 1 + \delta \\
\frac{2^{i+1}\gamma}{\beta + \sum_{j=0}^i \alpha_j 2^j \gamma} &\leq \delta \\
2^{i+1}\gamma &\leq \delta\beta + \delta \sum_{j=0}^i \alpha_j 2^j \gamma \\
&\leq \delta\beta + \delta\alpha_i 2^i \gamma + \delta \sum_{j=0}^{i-1} \alpha_j 2^j \gamma \\
2^{i+1} &\leq 1 + \delta\alpha^*(2^i - 1) + \delta\alpha_i 2^i \\
\alpha_i &\geq \frac{2^{i+1}}{\delta 2^i} - \frac{1}{\delta 2^i} - \frac{\delta\alpha^*(2^i - 1)}{\delta 2^i} \\
&\geq \frac{2}{\delta} - \frac{1}{\delta 2^i} - \frac{2^i - 1}{\delta 2^i} \\
&\geq \frac{2}{\delta} - \frac{1}{\delta} \\
&= \frac{1}{\delta}
\end{aligned}$$

Knowing that all α_i are equal to $\frac{1}{\delta}$, we can determine the value of k :

$$\begin{aligned}
\beta + \sum_{i=0}^k \frac{1}{\delta} 2^i \gamma &\geq |s_{ij}| \\
\beta + \frac{1}{\delta} \gamma (2^{k+1} - 1) &\geq |s_{ij}| \\
\beta + \beta(2^{k+1} - 1) &\geq |s_{ij}| \\
2^{k+1} &\geq \frac{|s_{ij}|}{\beta} \\
k &\geq \log \frac{|s_{ij}|}{\beta} - 1
\end{aligned}$$

This gives a total number of subsegments of $O(\frac{1}{\delta} \log \frac{|s_{ij}|}{\beta})$ for each point-segment pair. The sum over all pairs is largest when all $|s_{ij}|$ are equal, i.e. $1/nm$. This gives us a total number of subsegments for all pairs of $O\left(\frac{nm}{\delta} \log \frac{1/nm}{\beta}\right) = O\left(\frac{nm}{\delta} \log \frac{1}{\delta}\right)$. \square

Lemma 7. Q has $O\left(\frac{nm}{\delta} \log \frac{1}{\delta}\right)$ subsegments.

Proof. Consider any subsegment $r \in R$. Any subsegment $q \in Q$ that overlaps with r has $|q| \geq |r|/4$: otherwise q was subdivided unnecessarily. As the subsegments in Q are disjoint, it follows that r can overlap with at most 5 subsegments in Q . As such, Q contains at most 5 times more subsegments than R , which, by Lemma 6, is $O\left(\frac{nm}{\delta} \log \frac{1}{\delta}\right)$. \square

Putting everything together, we obtain the following result:

Theorem 8. *Let P be a set of n weighted points and S be a set of m line segments with equal total weight, let $|\mu^*|$ be the cost of an optimal transport plan between them, and let δ be any constant > 0 . Given an algorithm that constructs a $(1 + \delta)$ -approximation between weighted sets of k points in $f_\delta(k)$ time, we can construct a transport plan between P and S with cost $\leq (1 + 25\delta)|\mu^*|$ in $O(f_\varepsilon(\frac{nm}{\varepsilon} \log(\frac{1}{\varepsilon})))$ time.*

Proof. In Theorem 5, we prove that an optimal transport plan ν between P and T approximates $|\mu^*|$. However, we may be able to compute a $(1 + \delta)$ -approximation to ν faster than we are able to compute it exactly. It remains to be shown that this approximation also suffices.

Plugging in a $(1 + \delta)$ -approximation to $|\nu|$, rather than the exact value, we obtain the ratio

$$\frac{(1 + \delta)^3 |\mathcal{W}| + \frac{4\delta^2}{nm} + \frac{4\delta^3}{nm} + \frac{2\delta^4}{nm}}{|\mathcal{W}|}$$

Following the same strategy as in the proof of Theorem 5, we derive the approximation ratio as follows:

$$\begin{aligned} & \frac{(1 + \delta)^3 \cdot (\delta - 2\delta^2 - 2\delta^3) + 4\delta^2 + 4\delta^3 + 2\delta^4}{\delta - 2\delta^2 - 2\delta^3} \\ = & \frac{(1 + 3\delta + 3\delta^2 + \delta^3)(1 - 2\delta - 2\delta^2) + 4\delta + 4\delta^2 + 2\delta^3}{1 - 2\delta - 2\delta^2} \\ = & \frac{1 + 5\delta - \delta^2 - 9\delta^3 - 8\delta^4 - 2\delta^5}{1 - 2\delta - 2\delta^2} \\ = & 1 + \frac{7\delta + \delta^2 - 9\delta^3 - 8\delta^4 - 2\delta^5}{1 - 2\delta - 2\delta^2} \\ = & 1 + \delta + \frac{6\delta + 3\delta^2 - 7\delta^3 - 8\delta^4 - 2\delta^5}{1 - 2\delta - 2\delta^2} \\ \leq & 1 + \delta + \frac{9\delta}{1 - 2\delta - 2\delta^2} \\ \leq & 1 + 25\delta \quad \text{(assuming } \delta \leq \frac{1}{4}) \end{aligned}$$

As such, using an approximation of ν still gives us an approximation of μ^* , albeit with a somewhat worse dependency on δ . \square

To our knowledge, the current fastest algorithm to calculate a $(1 + \delta)$ -approximation to ν is that by Fox and Lu [5], which runs in $O(N\delta^{-O(1)} \text{polylog } N)$ time. Setting $\delta = \varepsilon/25$, this gives the following corollary to the previous theorem:

Corollary 9. *For any constant $\varepsilon > 0$, a transport plan between P and S with cost $\leq (1 + \varepsilon)|\mu^*|$ can be constructed in $O(\frac{nm}{\varepsilon^c} \text{polylog}(\frac{nm}{\varepsilon}))$ time with high probability.*

6 Points to triangles

We consider the case where P is a set of weighted points with total weight one and S is a set of m triangles with total area one. We denote the longest edge of any triangle by Δ . Our strategy is similar to before: we subdivide the triangles such that for each subregion, the ratio between its shortest and longest distance to each point is at most $1 + \delta$ for some appropriate choice of $\delta \in O(\varepsilon)$. We then show that a solution based on an optimal transport plan between P and some points inside the subregions approximates an optimal solution.




Figure 5: An example subdivision of a set of triangles. Each cell records the total area of triangles it intersects. Cells that are part of Q are shown in black; empty cells are shown in grey. A green disk denotes the distance of δ/\sqrt{nm} from each point. Note that ε is set to a very large value here for the clarity of the resulting image.

We first overlay a uniform grid onto our triangles with grid cells of size $\Delta \times \Delta$. We can identify the cells of this grid that contain a triangle in $O(m \log m)$ time using point-location in a compressed quadtree where the smallest cell size is $\Delta \times \Delta$ [8]. As each triangle can intersect at most four cells, the total size of this set of cells is $O(m)$. We now recursively subdivide each cell as follows: if there is a point in P such that the whole cell is within distance δ/\sqrt{nm} of it, we stop; otherwise, if for any point the ratio of distances to the furthest and closest point in this cell is more than $1 + \delta$, we subdivide this cell into four cells of one quarter the area. If the ratio holds for all points, we stop. Call the resulting set of cells Q ; see Figure 5 for an example.

During each subdivision, we keep track of the total area of triangles contained inside that cell. We can then once again build a complete bipartite graph $G = (P \cup Q, P \times Q)$, with the capacity of each vertex set to the weight of the corresponding point or the total area of triangles contained in the corresponding cell, and the weight of each edge equal to the shortest distance between the point and the cell it connects. The cost of a minimum cost flow \mathcal{W} is now once again a lower bound to the EMD, exactly as in Lemma 2. In an analogous way to Lemma 3, we obtain a lower bound on the cost of \mathcal{W} :

Lemma 10.
$$|\mathcal{W}| \geq \frac{\delta}{\sqrt{nm}} - \frac{\pi\delta^3 + 2\pi\delta^4 + \pi\delta^5}{\sqrt{nm}}.$$

Proof. For a given point-triangle pair $(p, s) \in P \times S$, consider the cells in Q intersecting s that have a point within distance δ/\sqrt{nm} of p . By construction, such a cell has its furthest point at distance at most $(1 + \delta) \cdot \delta/\sqrt{nm} = \delta/\sqrt{nm} + \delta^2/\sqrt{nm}$. Therefore, the total area of these cells is at most $\pi(\delta/\sqrt{nm} + \delta^2/\sqrt{nm})^2 = \pi(\delta^2 + 2\delta^3 + \delta^4)/nm$. Over all point-triangle pairs, this gives a total area of at most $\pi(\delta^2 + 2\delta^3 + \delta^4)$. This leaves $1 - \pi(\delta^2 + 2\delta^3 + \delta^4)$ with distance at least δ/\sqrt{nm} in \mathcal{W} . The cost is therefore at least $(1 - \pi(\delta^2 + 2\delta^3 + \delta^4)) \cdot \delta/\sqrt{nm} = \delta/\sqrt{nm} - \pi(\delta^3 + 2\delta^4 + \delta^5)/\sqrt{nm}$. \square

We now once again approximate $|\mathcal{W}|$ by reducing the flow problem to a transportation problem between two sets of weighted points. Again, we pick any point in each cell $q \in Q$ and give it a weight equal to the area of triangles contained in q ; call this set of points T .

Lemma 11.
$$|\mathcal{W}| \leq |\mu| \leq (1 + \delta)^2 |\mathcal{W}| + \frac{2\pi\delta^3}{\sqrt{nm}} + \frac{\pi\delta^4}{\sqrt{nm}}$$

Proof. Let ν be an optimal transport plan between P and T , and let $|\nu|$ be its cost. We can upper bound $|\nu|$ by measuring all distances to the furthest point in each cell. We

constructed Q such that the ratio of the closest and furthest distance between any point-cell pair is $1 + \delta$ when the furthest distance is at least δ/\sqrt{nm} . We can therefore bound all parts of ν where the distance is at least δ/\sqrt{nm} by $(1 + \delta)|\mathcal{W}|$. The total mass being moved over a distance at most δ/\sqrt{nm} in ν is at most $\pi\delta^2$, giving a cost of $\pi\delta^3/\sqrt{nm}$. The total cost when measuring to the furthest point is therefore $(1 + \delta)|\mathcal{W}| + \pi\delta^3/\sqrt{nm}$.

We now turn ν into a transport plan μ between P and Q by spreading the mass sent to each point $t \in T$ out evenly over the parts of the triangles in the cell in Q that t was derived from. By construction, for cells with a distance of at least δ/\sqrt{nm} , this increases the cost by at most a factor $1 + \delta$. We can therefore bound the cost of this part of μ by $(1 + \delta)|\nu|$. The remaining part has a total mass of at most $\pi\delta^2$, giving a cost of $\pi\delta^3/\sqrt{nm}$. The total cost of μ is then bound by $(1 + \delta)|nu| + \pi\delta^3/\sqrt{nm}$.

Plugging in the bound on $|\nu|$ obtained above, we obtain an upper bound of $(1 + \delta)^2|\mathcal{W}| + 2\pi\delta^3/\sqrt{nm} + \pi\delta^4/\sqrt{nm}$. The lower bound follows directly from the fact that none of the distance in μ are smaller than the distances between the same objects in \mathcal{W} . \square

Putting this all together, we can show that $|\mu|$ approximates $|\mu^*|$.

Theorem 12. $|\mu|$ is a $(1 + 9\delta)$ -approximation to the Earth Mover's Distance $|\mu^*|$ between P and S for $0 < \delta \leq \frac{1}{2\pi}$.

Proof. By Lemma 11 have that

$$|\mu| \leq (1 + \delta)^2|\mathcal{W}| + \frac{2\pi\delta^3}{\sqrt{nm}} + \frac{\pi\delta^4}{\sqrt{nm}}$$

$|\mathcal{W}|$ is also a lower bound on $|\mu|$; the ratio between the upper and lower bound is

$$\frac{(1 + \delta)^2|\mathcal{W}| + \frac{2\pi\delta^3}{\sqrt{nm}} + \frac{\pi\delta^4}{\sqrt{nm}}}{|\mathcal{W}|}$$

This ratio is largest for small values of $|\mathcal{W}|$, so we plug in the lower bound from

Lemma 10:

$$\begin{aligned}
& \frac{(1 + \delta)^2 |\mathcal{W}| + \frac{2\pi\delta^3}{\sqrt{nm}} + \frac{\pi\delta^4}{\sqrt{nm}}}{|\mathcal{W}|} \\
\leq & \frac{(1 + \delta)^2 \left(\frac{\delta}{\sqrt{nm}} - \frac{\pi\delta^3 + 2\pi\delta^4 + \pi\delta^5}{\sqrt{nm}} \right) + \frac{2\pi\delta^3}{\sqrt{nm}} + \frac{\pi\delta^4}{\sqrt{nm}}}{\frac{\delta}{\sqrt{nm}} - \frac{\pi\delta^3 + 2\pi\delta^4 + \pi\delta^5}{\sqrt{nm}}} \\
\leq & \frac{(1 + 2\delta + \delta^2)(1 - \pi\delta - 2\pi\delta^2 - \pi\delta^3) + 2\pi\delta^2 + \pi\delta^3}{1 - \pi\delta - 2\pi\delta^2 - \pi\delta^3} \\
= & \frac{1 + 2\delta + \delta^2 - 2\pi\delta^2 - 5\pi\delta^3 - 4\pi\delta^4 - \pi\delta^5}{1 - \pi\delta - 2\pi\delta^2 - \pi\delta^3} \\
= & 1 + \frac{2\delta + \delta^2 - \pi\delta - 4\pi\delta^3 - 4\pi\delta^4 - \pi\delta^5}{1 - \pi\delta - 2\pi\delta^2 - \pi\delta^3} \\
= & 1 + \delta + \frac{\delta + \delta^2 - \pi\delta - 2\pi\delta^3 - 3\pi\delta^4 - \pi\delta^5}{1 - \pi\delta - 2\pi\delta^2 - \pi\delta^3} \\
< & 1 + \delta + \frac{\delta^2}{1 - \pi\delta - 2\pi\delta^2 - \pi\delta^3} \\
\leq & 1 + \delta + \frac{\delta^2}{1 - \frac{1}{2} - \frac{1}{\pi} - \frac{1}{2\pi^2}} \quad (\text{assuming } \delta \leq \frac{1}{2\pi}) \\
< & 1 + \delta + 8\delta^2 \\
< & 1 + 9\delta
\end{aligned}$$

As $|\mathcal{W}|$ is also a lower bound for $|\mu|$ (Lemma 2), and μ can obviously not have lower cost than the optimal transport plan, this gives a $(1 + 9\delta)$ -approximation. \square

Setting $\delta = \varepsilon/9$ gives a $(1 + \varepsilon)$ -approximation.

6.1 Running time analysis

Our analysis will be the same as in Section 5.1; we just need to determine the size of Q . We will once again make an alternative subdivision of each $s_j \in S$, count the number of cells in that subdivision, and then argue that $|Q|$ differs by at most a constant factor. Our alternative subdivision is a direct adaptation of the one used in Section 5.1 to two dimensions: for each point p_i and triangle s_j , we fill a square with side length $2\delta/\sqrt{nm}$ centred on p_i with cells of size δ^2/\sqrt{nm} . From there, we add rings of cells of side length δ^2/\sqrt{nm} around the square, until the next full ring could have cells double the size without violating the ratio of $1 + \delta$ between the shortest and longest distance to p_i for any cell in the ring. We repeat this process until we have covered a square of size $\Delta \times \Delta$. Let R_{ij} be the resulting set of cells; see Figure 6 for an example. The proof is similar to Lemma 6.

Lemma 13. $R = \bigcup R_{ij}$ has $O\left(\frac{nm}{\delta^2} \log \frac{nm\Delta}{\delta}\right)$ cells.

Proof. We define $\beta = \frac{\delta}{\sqrt{nm}}$ and $\gamma = \frac{\delta^2}{\sqrt{nm}}$. In the following, we only analyse the case where p_{ij} is inside s_{ij} ; if it lies outside, the number of cells will be smaller, as the size of the cells increases with distance. We also analyse the number of cells in one quadrant only; the total number is simply four times as many. See Figure 6 for an illustration of a quadrant.




Figure 6: On the left, part of one quadrant of the construction of R_{ij} . There are α_i layers of cells of size $2^i\gamma$ before the size is doubled. On the right, an illustration of the argument that a cell of Q (blue) has at least one quarter the edge length of a cell of R that it intersects (red).




Figure 7: A cell of size δx has the correct ratio after one subdivision.

The number of cells created as we add rings of cells on s_{ij} can then be written as

$$\frac{\beta^2}{\gamma^2} + \sum_{i=0}^k 2\alpha_i \cdot \frac{\beta + \sum_{j=0}^{i-1} \alpha_j 2^j \gamma}{2^i \gamma} + \alpha_i^2$$

where k is the number of times we double the size of the cells, and α_i is the number of rings containing cells of a size that has been doubled i times. The number of cells can then be calculated by finding the values of k and α_i . We take the values of α_i to be the same as in Lemma 6 (i.e. $1/\delta$): along a horizontal or vertical line through p_i these values are exact, and cells not on this line can be made to have the correct ratio through one extra subdivision.

Let (x, y) be the vector from p_i to the closest point on the cell. Assume w.l.o.g. that $0 \leq y \leq x$; the other cases are symmetrical. By construction of our subdivision, we know that the cell has size at most δx . We will now show that by dividing the cell on extra time (i.e. to a size of $\delta x/2$), the furthest point will have the desired ratio irrespective of the value of y . See Figure 7 for an illustration.

$$\begin{aligned}
\frac{\sqrt{(x + \frac{\delta x}{2})^2 + (y + \frac{\delta x}{2})^2}}{\sqrt{x^2 + y^2}} &\leq 1 + \delta \\
\frac{(x + \frac{\delta x}{2})^2 + (y + \frac{\delta x}{2})^2}{x^2 + y^2} &\leq (1 + \delta)^2 \\
\frac{x^2 + y^2 + \delta x^2 + \delta xy + \frac{\delta^2 x^2}{2}}{x^2 + y^2} &\leq 1 + 2\delta + \delta^2 \\
\frac{\delta x^2 + \delta xy + \frac{\delta^2 x^2}{2}}{x^2 + y^2} &\leq 2\delta + \delta^2 \\
\delta xy &\leq \delta x^2 + \frac{\delta^2 x^2}{2} + 2\delta y^2 + \delta^2 y^2
\end{aligned}$$

As $\delta xy \leq \delta x^2$, and the other terms on the right-hand side are positive, the inequality holds. As such, the construction described can be turned into one where all cells have the desired ratio with one extra subdivision.

Plugging the values of α_i , β , γ into our initial formula, we can obtain the number of cells as a function of k :

$$\begin{aligned}
&\left(\frac{\frac{\delta}{\sqrt{nm}}}{\frac{\delta^2}{\sqrt{nm}}}\right)^2 + \sum_{i=0}^k \frac{2}{\delta} \cdot \frac{\frac{\delta}{\sqrt{nm}} + \sum_{j=0}^{i-1} \frac{1}{\delta} 2^j \frac{\delta^2}{\sqrt{nm}}}{2^i \frac{\delta^2}{\sqrt{nm}}} + \frac{1}{\delta^2} \\
&= \frac{1}{\delta^2} + \frac{k}{\delta^2} + \frac{2}{\delta} \sum_{i=0}^k \frac{\frac{\delta}{\sqrt{nm}} + \frac{\delta}{\sqrt{nm}} \sum_{j=0}^{i-1} 2^j}{2^i \frac{\delta^2}{\sqrt{nm}}} \\
&= \frac{1}{\delta^2} + \frac{k}{\delta^2} + \frac{2}{\delta} \sum_{i=0}^k \frac{\frac{\delta}{\sqrt{nm}} + \frac{\delta}{\sqrt{nm}} (2^i - 1)}{2^i \frac{\delta^2}{\sqrt{nm}}} \\
&= \frac{1}{\delta^2} + \frac{k}{\delta^2} + \frac{2}{\delta} \sum_{i=0}^k \frac{2^i \frac{\delta}{\sqrt{nm}}}{2^i \frac{\delta^2}{\sqrt{nm}}} \\
&= \frac{1}{\delta^2} + \frac{k}{\delta^2} + \frac{2}{\delta} \sum_{i=0}^k \frac{1}{\delta} \\
&= \frac{1}{\delta^2} + \frac{k}{\delta^2} + \frac{2k}{\delta^2} \\
&\in O\left(\frac{k}{\delta^2}\right)
\end{aligned}$$

We can directly calculate the value of k by considering the number of doublings needed

to cover a horizontal line segment of length Δ starting at p_i :

$$\begin{aligned} \frac{\delta}{\sqrt{nm}} + \sum_{i=0}^k \frac{1}{\delta} \cdot 2^i \cdot \frac{\delta^2}{\sqrt{nm}} &= \Delta \\ \frac{\delta}{\sqrt{nm}} + \frac{\delta}{\sqrt{nm}} \sum_{i=0}^k 2^i &= \Delta \\ 1 + \sum_{i=0}^k 2^i &= \frac{\sqrt{nm}\Delta}{\delta} \\ 2^{k+1} &= \frac{\sqrt{nm}\Delta}{\delta} \\ k &\in O\left(\log \frac{nm\Delta}{\delta}\right) \end{aligned}$$

This gives a total number of cells of $O\left(\frac{1}{\delta^2} \log \frac{nm\Delta}{\delta}\right)$ per point-triangle pair. Over all pairs, we obtain a total number of cells of $O\left(\frac{nm}{\delta^2} \log \frac{nm\Delta}{\delta}\right)$. \square

Lemma 14. *Q has $O\left(\frac{nm}{\delta^2} \log \frac{nm\Delta}{\delta}\right)$ cells.*

Proof. Consider any cell $r \in R$. Any cell $q \in Q$ that overlaps with r has $|q| \geq |r|/16$: otherwise q was subdivided unnecessarily; see Figure 6. As the cells in Q are disjoint, it follows that r can overlap with at most 25 cells in Q . As such, Q contains at most 25 times more cells than R , which, by Lemma 13, is $O\left(\frac{nm}{\delta^2} \log \frac{nm\Delta}{\delta}\right)$. \square

This leads to the following result:

Theorem 15. *Let P be a set of n weighted points and S be a set of m triangles with equal total weight, let Δ be the longest edge length in S after normalising its total area to one, let $|\mu^*|$ be the cost of an optimal transport plan between P and S , and let δ be any constant > 0 . Given an algorithm that constructs a $(1 + \delta)$ -approximation between weighted sets of k points in $f_\delta(k)$ time, we can construct a transport plan between P and S with cost $\leq (1 + 9\delta)|\mu^*|$ can be constructed in $O\left(f_\delta\left(\frac{nm}{\delta^2} \text{polylog}\left(\frac{nm\Delta}{\delta}\right)\right)\right)$ time.*

We can again calculate a $(1 + \delta)$ -approximation to ν in $O(N\delta^{-O(1)} \text{polylog } N)$ time using the algorithm by Fox and Lu [5], giving the following corollary to the previous theorem:

Corollary 16. *For any constant $\varepsilon > 0$, a transport plan between P and S with cost $\leq (1 + \varepsilon)|\mu^*|$ can be constructed in $O\left(\frac{nm}{\varepsilon^c} \text{polylog}\left(\frac{nm\Delta}{\varepsilon}\right)\right)$ time with high probability.*

7 Segments to segments

In the previous section, we considered the case when one of our two input sets contains points. We now describe an algorithm to compute the EMD between two sets of line segments. Here, we cannot directly apply our general approach of subdividing: the optimal transport plan may have a cost arbitrarily close to zero. As such, if we disregard everything within some radius of one of the sets, there may be nothing left. We solve this by introducing an additive term into the approximation. The cost of a plan generated by our algorithm is




Figure 8: Two points, with their mass assignment in an optimal solution shown in purple. We greedily match a point x on u to $\tau(x)$ on v , and obtain the assignment of mass shown in green.

$(1 + \varepsilon)|\mu^*| + A$, for some value A depending on ε . This allows us to greedily match parts of the input within a small distance of each other, and then solve the remainder with our previous approach.

Let $P = \{p_1, \dots, p_n\}$ and $S = \{s_1, \dots, s_m\}$ be sets of line segments with equal total length. Our algorithm is then as follows. First, we greedily match equal-length pieces of P and S that are within distance δ/nm of each other, until no such pieces remain; we describe this process in more detail later. Let P' and S' be the remaining parts of P and S , respectively. We subdivide P' and S' as before: for every $p \in P'$, if there is an $s \in S'$ such that the ratio between the closest and furthest distance is more than $1 + \delta$, cut p in half; after processing P' , do the same for S' . Call the resulting sets Q and R . We then choose a point on each $q \in Q$ and $r \in R$, with a weight equal to the length of the subsegment, and solve an optimal transport problem between these two point sets. Our final transport plan is then obtained by spreading the mass moved between any two points evenly over the segments they were chosen on.

We first prove that greedily matching parts of the input within distance δ/nm increases the cost of an optimal solution by at most an additive term. The proof for the approximation algorithm then follows the same structure as in the previous sections. Let μ_M be a transport plan between the parts of the input that were greedily matched, in which the longest distance is at most δ/nm , and let μ_G^* be an optimal transport plan for the remainder of the input.

Lemma 17. *Let u and v be two subsegments with length l of P and S , respectively. If all mass from u can be transported to v with distance at most κ , then a transport plan between $P \setminus \{u\}$ and $S \setminus \{v\}$ has cost at most $|\mu^*| + l\kappa$.*

Proof. We will construct a transport plan in which u and v are removed, having cost at most $|\mu^*| + l\kappa$. The cost of an optimal solution on the remainder is then not higher.

Let $\mu_s(x, t) : \mathbb{R}^2 \times [0, 1] \rightarrow \mathbb{R}^2$ be a function describing, for a point $x \in s$, where its mass comes from or goes to (recall that each point sends or receives mass density one), let $d_s(x, t) : \mathbb{R}^2 \times [0, 1] \rightarrow \mathbb{R}$ be defined as $d(x, \mu_s(x, t))$, and let $\tau : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a mapping of points on u to points on v such that for all $x \in u$, $d(x, \tau(x)) \leq \kappa$. The cost of the part of μ^* involving segments u and v can then be written as

$$c^*(u) = \int_{x \in u} \int d_u(x, t) \, dt \, dx$$

$$c^*(v) = \int_{x \in u} \int d_v(\tau(x), t) \, dt \, dx$$

We modify μ^* by removing u and v , and moving all mass that each point $\tau(x)$ receives in μ^* to where x moved it in μ^* ; see Figure 8. We can distribute this mass in any way we

like, as the total incoming and outgoing mass is one by definition. This gives a transport plan μ with cost

$$\begin{aligned} |\mu| &= |\mu^*| - \left(\int_{x \in u} \int d_u(x, t) \, dt \, dx \right) - \left(\int_{x \in u} \int d_v(\tau(x), t) \, dt \, dx \right) \\ &\quad + \left(\int_{x \in u} \int d(\mu_u(x, t), \mu_v(\tau(x), t)) \, dt \, dx \right) \end{aligned}$$

By the triangle inequality, $d(\mu_u(x, t), \mu_v(\tau(x), t)) \leq d_v(\tau(x), t) + \kappa + d_u(x, t)$. It follows that

$$\begin{aligned} |\mu| &\leq |\mu^*| - \left(\int_{x \in u} \int d_u(x, t) \, dt \, dx \right) - \left(\int_{x \in u} \int d_v(\tau(x), t) \, dt \, dx \right) \\ &\quad + \left(\int_{x \in u} \int d_v(\tau(x), t) + \kappa + d_u(x, t) \, dt \, dx \right) \\ &= |\mu^*| + l\kappa \end{aligned} \quad \square$$

We can now bound the costs of μ_G^* and μ_M .

Lemma 18. $|\mu_G^*| \leq |\mu^*| + \frac{\delta}{nm}$.

Proof. Any subsegments of P and S with length l that are greedily matched increase the cost of an optimal solution in the remaining part by at most $\delta l/nm$ (Lemma 17). The total length that can be greedily matched is at most one, so the total extra cost is at most δ/nm . \square

Lemma 19. $|\mu_M| \leq \frac{\delta}{nm}$.

Proof. By construction, the distance over which any mass is transported in μ_M is at most δ/nm . The total mass transported is at most one, giving the bound. \square

For each segment $p \in P$, we can straightforwardly compute a maximal subset that can be transported over distance at most δ/nm . Consider each segment $s \in S$: the supporting lines of p and s can intersect inside p or s , outside both, or not at all. If they don't intersect (i.e. are parallel), computation of the parts that can be transported within the required distance is trivial. If they intersect outside both, we can find the points on p and s furthest from the intersection point that are within the required distance, then find the largest distance we can move towards the intersection point while staying within the required distance. If they intersect inside one or both of the segments, we split the segments at the intersection point and handle both sides using the case for intersections outside the segments.

Let P' and S' be the parts of P and S that remain after the greedy matching, with $|P'| = |S'| = \ell$. We subdivide P' and S' into Q and R as described above. As before, we define a complete bipartite graph $G = (Q \cup R, Q \times R)$, where the weight of each edge is equal to the shortest distance between the two subsegments it connects, and the capacity of each vertex is equal to the length of the subsegment it represents. Let \mathcal{W} be a minimum cost flow in G ; we observe the following lower bound on its cost:

Lemma 20. $|\mathcal{W}| \geq \frac{\delta \ell}{nm}$.

Proof. By construction, the distances in G are at least δ/nm . As the total mass moved is ℓ , we obtain the bound stated in the lemma. \square

Lemma 2 also still applies to the part of the input that remains after greedy matching. We now approximate $|\mathcal{W}|$ by reducing the flow problem to a transportation problem between two weighted point sets. We pick any point on each $q \in Q$ and $r \in R$, and give them weights equal to $|q|$ and $|r|$. Call these sets of points U and V . We can now bound the cost of μ_G in terms of μ_G^* using the flow problem.

Lemma 21. $|\mu_G^*| \leq |\mu_G| \leq (1 + \delta)^2 |\mu_G^*|$.

Proof. Let ν be an optimal transport plan between U and V , and let $|\nu|$ be its cost. We can upper bound $|\nu|$ by measuring all distances to the furthest points inside the segments. By construction of Q and R , the ratio of longest to shortest distance is at most $1 + \delta$. The cost $|\nu|$ of ν can therefore not be more than $(1 + \delta)|\mathcal{W}|$.

We can turn μ into a valid transport plan μ_G between Q and R by spreading the mass moved to each point in U and V evenly over the segments in Q and R that they were derived from. Again, by construction, the distances increase by a factor of at most $1 + \delta$, giving $\mu_G \leq (1 + \delta)|\nu|$.

Plugging in the bound on $|\nu|$ obtained above, we obtain an upper bound of $(1 + \delta)^2 |\mathcal{W}|$. As $|\mathcal{W}| \leq |\mu_G^*|$, we obtain that $|\mu_G| \leq (1 + \delta)^2 |\mu_G^*|$. The lower bound follows directly from the fact that μ_G^* is optimal, and therefore cannot have a cost higher than that of μ_G . \square

We can then show that, for a transport plan $\mu = \mu_G + \mu_M$, $|\mu|$ approximates $|\mu^*|$:

Theorem 22. $|\mu| \leq (1 + 3\delta)|\mu^*| + \frac{5\delta}{nm}$.

Proof. By Lemma 21, we know that $|\mu_G| \leq (1 + \delta)^2 |\mu_G^*|$. As $\delta \leq 1$, $(1 + \delta)^2 \leq 1 + 3\delta$. By Lemma 18, we have that $|\mu_G^*| \leq |\mu^*| + (1 - \ell)\delta/nm$. Combining the two results, we get that

$$\begin{aligned} |\mu_G| &\leq (1 + 3\delta)|\mu_G^*| \\ &\leq (1 + 3\delta) \left(|\mu^*| + (1 - \ell) \frac{\delta}{nm} \right) \\ &\leq (1 + 3\delta)|\mu^*| + (1 - \ell) \frac{\delta + 3\delta^2}{nm} \\ &\leq (1 + 3\delta)|\mu^*| + \frac{4\delta}{nm} \end{aligned}$$

By Lemma 19, $|\mu_M| \leq \delta/nm$. As $|\mu| = |\mu_G| + |\mu_M|$, we obtain the bound stated in the lemma. \square

Setting $\delta = \varepsilon/3$ gives a $(1 + \varepsilon)$ -approximation with an additive term of $5\varepsilon/3nm$

7.1 Running time analysis

During the greedy matching, each $p \in P$ may have been cut into m pieces, and each $s \in S$ into n pieces. As such, P' and S' (the parts remaining after greedy matching) both contain $O(nm)$ subsegments. In the worst case, P' and S' are close to each other everywhere, causing them to be subdivided into the smallest possible subsegments. As the minimum distance is δ/nm , and the ratio of the longest and shortest distance between any two subsegments is $1 + \delta$, the smallest possible subsegment has size $\Theta(\frac{\delta^2}{nm})$. Each subsegment of P' and S' may give rise to one extra subsegment in Q and R , as the length may not be exactly divisible by δ/nm . This gives sets Q and R a size of $O(\frac{nm}{\delta^2} + nm) = O(\frac{nm}{\varepsilon^2})$, leading to the following result:

Theorem 23. *Let P and S be sets of n and m line segments in the plane, both having equal total length, let $|\mu^*|$ be the cost of an optimal transport plan between them, and let δ be any constant > 0 . Given an algorithm that constructs a $(1 + \delta)$ -approximation between weighted sets of k points in $f_\delta(k)$ time, we can construct a transport plan between P and S with cost $\leq (1 + c'\delta)|\mu^*| + \frac{5\varepsilon}{3nm}$ for some constant c' in $O(f_\delta(\frac{nm}{\delta^2}))$ time with high probability.*

We can again calculate a $(1 + \delta)$ -approximation to ν in $O(N\delta^{-O(1)} \text{polylog } N)$ time using the algorithm by Fox and Lu [5], giving the following corollary to the previous theorem:

Corollary 24. *For any constant $\varepsilon > 0$, a transport plan between P and S with cost $\leq (1 + \varepsilon)|\mu^*|$ can be constructed in $O(\frac{nm}{\varepsilon^c} \text{polylog}(\frac{nm}{\varepsilon}))$ time with high probability.*

8 Triangles to triangles

We consider the case where P and S are both sets of triangles with total area one and longest edge length Δ . The algorithm is completely analogous to the one for transport between sets of segments: we greedily match parts of the input within a certain distance, subdivide the remainder and approximate the optimal transport plan by reduction to a minimum cost flow. As the setup and proofs are exactly the same as in the previous section (just substitute the integrals over segments with integrals over area), this is omitted. All we need is an algorithm that can greedily match parts of the input within a given distance.

We do this greedy matching as follows. We can first remove the parts where P and S overlap: they have cost zero. We then overlay a grid with cells of size $\delta/(2\sqrt{nm})$ onto our input, and keep only the cells that contain an edge or are adjacent to one that does (the other cells already have the desired clearance from cells containing triangles from the other set). Inside each cell, we record the total area of triangles from P and S that lie inside it separately. For parts of P and S that lie inside the same cell, we match as much as possible, resulting in a grid where each cell only contains parts of P or S . We then match as much of each cell as possible to each of its eight neighbours. The maximum distance over which we have greedily matched weight is $\sqrt{2}\delta/\sqrt{nm}$, and the remaining parts of P and S have a minimum distance of δ/\sqrt{nm} to each other. Applying the analogous version of the segment-to-segment algorithm then gives us a $(1 + \varepsilon)$ -approximation with an additive term of $O(\varepsilon/\sqrt{nm})$.

8.1 Running time analysis

The number of cells examined during the greedy matching is $O(\frac{nm\Delta}{\delta})$ per triangle, so $O(\frac{nm\Delta(n+m)}{\delta})$ in total. The part of the input remaining after greedy matching can be close to each other everywhere, causing it to be subdivided into cells of size $\Theta(\frac{\delta^2}{\sqrt{nm}})$. There may be $O(\frac{nm\Delta}{\delta^2})$ cells that intersect the boundaries of a triangle, or $O(\frac{nm\Delta(n+m)}{\delta^2})$ in total; the other cells are interior to some triangle, and as the total area is one, there can be at most $O(\frac{\sqrt{nm}}{\delta^2})$ of them. The total number of cells is therefore at most $O(\frac{nm\Delta(n+m)}{\delta^2})$. This gives the following result:

Theorem 25. *Let P be a set of n and S a set of m triangles in the plane, both having equal total area and longest edge length at most Δ after normalising their total areas to one, let $|\mu^*|$ be the cost of an optimal transport plan, and let δ be any constant > 0 . Given an algorithm that constructs a $(1 + \delta)$ -approximation between weighted sets of k points in $f_\delta(k)$*

time, we can construct a transport plan between P and S with cost $\leq (1 + c'\delta)|\mu^*| + O(\frac{\delta}{\sqrt{nm}})$ for some constant c' can be constructed in $O\left(f_\delta\left(\frac{nm\Delta(n+m)}{\delta^2}\right)\right)$ time.

We can again calculate a $(1 + \delta)$ -approximation to ν in $O(N\delta^{-O(1)})$ polylog N time using the algorithm by Fox and Lu [5], giving the following corollary to the previous theorem:

Corollary 26. *For any constant $\varepsilon > 0$, a transport plan between P and S with cost $\leq (1 + \varepsilon)|\mu^*| + O(\frac{\varepsilon}{\sqrt{nm}})$ can be constructed in $O\left(\frac{nm\Delta(n+m)}{\varepsilon^c}$ polylog $(\frac{nm\Delta}{\varepsilon})\right)$ time with high probability.*

9 Higher dimensions

In this section we show how our approach can be extended to work in d -dimensional space. We discuss the case of transporting mass from points to d -dimensional simplices, and from one set of simplices to another.

9.1 Points to simplices

The approach described here is a direct extension of the one detailed in Section 6. Let P be a set of n weighted points in d dimensions with total mass one, and let S be a set of m d -dimensional simplices with total volume one and longest edge length Δ . We start by overlaying an infinite grid of size Δ and identifying the cells intersected by any simplex in $O(dm \log(m) + 6^d d^2 d^d m)$ time using compressed quadtrees [8]. We then repeatedly subdivide each cell until the ratio between the shortest and longest distance is at most $1 + \delta$ for all points in P , or until it is wholly within $\delta/(nm)^{1/d}$ of any point in P . Call the resulting set of cells Q . We can again show that picking one point in each cell of Q with weight equal to the total volume of simplices contained in it, and then solving a transport problem between P and the resulting set of points, approximates the transport problem between P and S .

As the structure of the proof is very similar to that contained in Section 6, we omit some of the intermediate lemmas here. We start with the lower bound on the cost of a minimum cost flow \mathcal{W} in the bipartite graph $G = (P \cup Q, P \times Q)$:

Lemma 27. $|\mathcal{W}| \geq \frac{\delta}{(nm)^{1/d}} - \frac{\delta(2(\delta + \delta^2))^d}{(nm)^{1/d}}$.

Proof. For a given point-simplex pair $(p, s) \in P \times S$, consider the cells in Q intersecting s that have a point within distance $\delta/(nm)^{1/d}$ of p . Such a cell has its furthest point at most at distance $(\delta + \delta^2)/(nm)^{1/d}$. The total volume of these cells is then $(2(\delta + \delta^2))^d/nm$ (the volume of a d -dimensional hypercube with radius $(\delta + \delta^2)/(nm)^{1/d}$, which contains the hypersphere with the same radius). Over all point-simplex pairs, this gives a volume of at most $(2(\delta + \delta^2))^d$, leaving $1 - (2(\delta + \delta^2))^d$ with distance at least $\delta/(nm)^{1/d}$ in \mathcal{W} . The cost is therefore at least $(1 - (2(\delta + \delta^2))^d) \cdot \delta/(nm)^{1/d} = \delta/(nm)^{1/d} - \delta(2(\delta + \delta^2))^d/(nm)^{1/d}$. \square

We again approximate $|\mathcal{W}|$ by reducing the flow problem to a transportation problem between two sets of weighted points. We do this by picking a point in each cell $q \in Q$ and giving it a weight equal to the total volume of simplices contained in q . Call this set of points T :

Lemma 28. *The cost $|\nu|$ of an optimal transport plan ν between P and T satisfies $|\mathcal{W}| \leq |\nu| \leq (1 + \delta)|\mathcal{W}| + \frac{\delta(2(\delta + \delta^2))^d}{(nm)^{1/d}}$.*

Proof. The lower bound follows directly from the fact that none of the distances in ν are smaller than the distances between the same objects in \mathcal{W} (recall that the distances in \mathcal{W} are measured to the closest point on the cell). We can upper bound $|\nu|$ by measuring all distances to the furthest point in each cell. By construction, those distances are at most $1 + \delta$ times the distance to the closest point when the furthest distance is at least $\delta/(nm)^{1/d}$. We can therefore bound all parts of ν where the distance is at least $\delta/(nm)^{1/d}$ by $(1 + \delta)|\mathcal{W}|$. The total mass being moved over distance at most $\delta/(nm)^{1/d}$ in ν is at most $(2(\delta + \delta^2))^d$, giving a cost of $\delta(2(\delta + \delta^2))^d/(nm)^{1/d}$. The total cost when measuring to the furthest point is therefore $(1 + \delta)|\mathcal{W}| + \delta(2(\delta + \delta^2))^d/(nm)^{1/d}$. \square

As before, we turn ν into a valid transport plan μ between P and Q by spreading the mass moved to each point in T out evenly over the parts of simplices contained in the cell of Q the point was derived from. By the same argument used in Lemma 28, we obtain the following bound on the cost of μ :

Lemma 29. $|\mu| \leq (1 + \delta)|\nu| + \frac{\delta(2(\delta + \delta^2))^d}{(nm)^{1/d}}$.

Putting this all together, we can show that $|\mu|$ approximates the cost of the optimal transport plan $|\mu^*|$:

Theorem 30. $|\mu|$ is a $(1 + 21\delta)$ -approximation to the Earth Mover's Distance $|\mu^*|$ between P and S for $0 < \delta \leq \frac{1}{5}$.

Proof. We follow the same structure as Theorem 12, obtaining the following ratio, into which we plug the lower bound from Lemma 27:

$$\begin{aligned}
& \frac{(1 + \delta)^2|\mathcal{W}| + \frac{2\delta(2(\delta + \delta^2))^d}{(nm)^{1/d}} + \frac{\delta^2(2(\delta + \delta^2))^d}{(nm)^{1/d}}}{|\mathcal{W}|} \\
& \leq \frac{(1 + 2\delta + \delta^2) \left(\frac{\delta}{(nm)^{1/d}} - \frac{\delta(2(\delta + \delta^2))^d}{(nm)^{1/d}} \right) + \frac{2\delta(2(\delta + \delta^2))^d}{(nm)^{1/d}} + \frac{\delta^2(2(\delta + \delta^2))^d}{(nm)^{1/d}}}{\frac{\delta}{(nm)^{1/d}} - \frac{\delta(2(\delta + \delta^2))^d}{(nm)^{1/d}}} \\
& = \frac{1 + 2\delta + \delta^2 + (2(\delta + \delta^2))^d - \delta(2(\delta + \delta^2))^d - \delta^2(2(\delta + \delta^2))^d}{1 - (2(\delta + \delta^2))^d} \\
& = 1 + \delta + \frac{\delta + \delta^2 + 2(2(\delta + \delta^2))^d - \delta^2(2(\delta + \delta^2))^d}{1 - (2(\delta + \delta^2))^d} \\
& \leq 1 + \delta + \frac{\delta + \delta^2 + 2(2(\delta + \delta^2))((2(\delta + \delta^2)))^{d-1}}{1 - (2(\delta + \delta^2))^d} \\
& \leq 1 + \delta + \frac{\delta + \delta^2 + 8\delta}{1 - \frac{1}{2^d}} \quad (\text{Assuming } (\delta + \delta^2) \leq \frac{1}{4}) \\
& \leq 1 + \delta + \frac{9\delta + \delta^2}{\frac{1}{2}} \\
& < 1 + 21\delta
\end{aligned}$$

As $|\mathcal{W}|$ is also a lower bound for $|\mu|$ (Lemma 2), and μ can obviously not have lower cost than the optimal transport plan, this gives a $(1 + 21\delta)$ -approximation. \square

Note that the constant in our approximation is slightly worse than the one obtained in Theorem 12; this is because we approximate the volume of a hypersphere by the volume of its bounding cube, whereas before we could calculate the area of the disk exactly.

9.1.1 Running time analysis

We construct a structure similar to Section 6.1, then argue that the number of cells in our actual subdivision is similar. Our construction is the direct generalisation of the one described before: we build a layered structure of cells of increasing sizes. Let R_{ij} be the set of cells generated by point p_i and simplex s_j ; we now analyse how many cells are created.

Lemma 31. $R = \bigcup R_{ij}$ has $O\left(\frac{dd^{d/2}nm}{\delta^d} \log \frac{(nm)^{1/d}\Delta}{\delta}\right)$ cells.

Proof. We follow the structure of the proof of Lemma 14, again counting the number of cells in one “quadrant”, and then multiplying by the number of quadrants (2^d). The number of cells in a quadrant is

$$\left(\frac{\beta}{\gamma}\right)^d + \sum_{i=0}^k d\alpha_i \cdot \left(\frac{\beta + \sum_{j=0}^{i-1} \alpha_j 2^j \gamma}{2^i \gamma}\right)^{d-1} + \alpha_i^d$$

where α_i is the number of layers of cells that have doubled in size i times, β is the distance inside of which we use the smallest cell size ($\delta/(nm)^{1/d}$), γ is the smallest cell size ($\delta^2/(nm)^{1/d}$), and k is the number of times we need to double the cell size.

The value of each α_i is $1/\delta$: this value is exact along any axis, and we can show that all cells can be made to have the correct ratio with a given number of extra subdivisions. For a given cell $r \in R_{ij}$, let \mathbf{v} be the vector from p_i to the closest point on r , and let γ' be the edge length of cell q . W.l.o.g. assume that $v_0 = \max v_i$; by construction this gives us that $(v_0 + \gamma')/v_0 \leq 1 + \delta$. Let \mathbf{u} be the vector from the closest point on q to the furthest point; we want to find a value x such that $\frac{|\mathbf{v} + \mathbf{u}/x|}{|\mathbf{v}|} \leq 1 + \delta$. Through the triangle inequality, we can upper bound the distance to the furthest point on q as $|\mathbf{v}| + |\mathbf{u}/x|$. We can now calculate the required value of x :

$$\begin{aligned} \frac{|\mathbf{v}| + \frac{|\mathbf{u}|}{x}}{|\mathbf{v}|} &\leq 1 + \delta \\ \frac{|\mathbf{u}|}{x|\mathbf{v}|} &\leq \delta \\ \frac{\gamma' \sqrt{d}}{x|\mathbf{v}|} &\leq \delta \\ \frac{\gamma' \sqrt{d}}{xv_0} &\leq \delta \\ \frac{\sqrt{d}}{x} \delta &\leq \delta \\ x &\geq \sqrt{d} \end{aligned}$$

So all cells have the correct ratio if their edge length is reduced by a factor of at least \sqrt{d} , which means each cell needs to be replaced by at most $O(d^{d/2})$ cells.

This gives us the following derivation for the number of cells:

$$\begin{aligned}
& \frac{1}{\delta^d} + \sum_{i=0}^k \frac{d}{\delta} \cdot \left(\frac{\frac{\delta}{(nm)^{1/d}} + \sum_{j=0}^{i-1} 2^j \frac{\delta}{(nm)^{1/d}}}{2^i \frac{\delta^2}{(nm)^{1/d}}} \right)^{d-1} + \frac{1}{\delta^d} \\
&= \frac{1}{\delta^d} + \frac{k}{\delta^d} + \frac{d}{\delta} \sum_{i=0}^k \left(\frac{\frac{\delta}{(nm)^{1/d}} + \frac{\delta}{(nm)^{1/d}} (2^i - 1)}{2^i \frac{\delta^2}{(nm)^{1/d}}} \right)^{d-1} \\
&= \frac{1}{\delta^d} + \frac{k}{\delta^d} + \frac{d}{\delta} \sum_{i=0}^k \left(\frac{\frac{\delta}{(nm)^{1/d}}}{\frac{\delta^2}{(nm)^{1/d}}} \right)^{d-1} \\
&= \frac{1}{\delta^d} + \frac{k}{\delta^d} + \frac{d}{\delta} \cdot \frac{k}{\delta^{d-1}} \\
&= \frac{1}{\delta^d} + \frac{k}{\delta^d} + \frac{kd}{\delta^d} \\
&\in O\left(\frac{kd}{\delta^d}\right)
\end{aligned}$$

The value of k is derived in the same way as before, giving $k \in O(\log((nm)^{1/d}\Delta/\delta))$. This gives $O(\frac{d}{\delta^d} \log \frac{(nm)^{1/d}\Delta}{\delta})$ cells per point-simplex pair, where each cell needs to be divided into $O(d^{d/2})$ smaller cells, for a total of $O(\frac{dd^{d/2}nm}{\delta^d} \log \frac{(nm)^{1/d}\Delta}{\delta})$ cells. \square

Lemma 32. Q has $O\left(\frac{5^d dd^{d/2}nm}{\delta^d} \log \frac{(nm)^{1/d}\Delta}{\delta}\right)$ cells.

Proof. Consider any cell $r \in R$. As before, any cell $q \in Q$ that overlaps with r has $|q| \geq |r|/4^d$: otherwise q was subdivided unnecessarily. As the cells in Q are disjoint, it follows that r can overlap with at most 5^d cells in Q . As such, Q contains at most 5^d times more cells than R , which, by Lemma 31, is $O\left(\frac{dd^{d/2}nm}{\delta^d} \log \frac{(nm)^{1/d}\Delta}{\delta}\right)$. \square

Combined with the time required to build the quadtree that we use to find the starting cells of our subdivision, this gives the following result:

Theorem 33. Let P be a set of n weighted points and S be a set of m simplices in \mathbb{R}^d with equal total weight, let Δ be the longest edge length in S after normalising its total volume to one, let $|\mu^*|$ be the cost of an optimal transport plan between P and S , and let δ be any constant > 0 . Given an algorithm that constructs a $(1 + \delta)$ -approximation between weighted sets of k points in $f_\delta(k)$ time, we can construct a transport plan between P and S with cost $\leq (1 + 21\delta)|\mu^*|$ in $O\left(6^d d^2 d^d m + f_\delta\left(\frac{5^d dd^{d/2}nm}{\delta^d} \log \frac{(nm)^{1/d}\Delta}{\delta}\right)\right)$ time.

Note that if we set $d = 2$, this is the same running time as the one obtained in Theorem 15. We can again calculate a $(1 + \delta)$ -approximation to ν in $O(N\delta^{-O(1)} \text{polylog } N)$ time using the algorithm by Fox and Lu [5]. Setting $\delta = \varepsilon/21$, this gives a total running time of $O\left(6^d d^2 d^d m + \frac{105^d dd^{d/2}nm}{\varepsilon^{d+c}} \text{polylog}\left(\frac{5^d dnm\Delta}{\varepsilon^d}\right)\right) \in O\left(6^d d^2 d^d m + \frac{105^d d^2 d^{d/2}nm}{\varepsilon^{d+c}} \text{polylog}\left(\frac{dnm\Delta}{\varepsilon^d}\right)\right)$.

Corollary 34. For any constant $\varepsilon > 0$, a transport plan between P and S with cost $\leq (1 + \varepsilon)|\mu^*|$ can be constructed in $O\left(6^d d^2 d^d m + \frac{105^d d^2 d^{d/2}nm}{\varepsilon^{d+c}} \text{polylog}\left(\frac{dnm\Delta}{\varepsilon^d}\right)\right)$ time with high probability, where c is some constant.

9.2 Simplices to simplices

The approach from Sections 7 and 8 can also be extended to work on d -dimensional simplices in d dimensions. We take the same approach of overlaying a grid with cells of size $\delta/(4(nm)^{1/d})$ onto the input and greedily matching the parts of P and S that are close together. The maximum distance over which we greedily match weight is then $\sqrt{d}\delta/(2(nm)^{1/d})$, and the remaining parts of P and S have minimum distance $\delta/(2(nm)^{1/d})$ to each other. We then approximate the transport plan between the remaining cells with a minimum cost flow. The same analysis still works, and we obtain a $(1 + \varepsilon)$ -approximation with an additive term of $O(\sqrt{d}\varepsilon/(nm)^{1/d})$.

9.2.1 Running time analysis

The number of cells examined during the greedy matching is $O(\frac{(nm)^{1/d}\Delta^d}{\delta})$ per simplex, so $O(\frac{(nm)^{1/d}\Delta^d(n+m)}{\delta})$ in total (note that we simplify the analysis by simply considering the volume of a d -dimensional cube of side length Δ). The part of the input remaining after greedy matching can be close to each other everywhere, causing it to be subdivided into cells of size $\Theta(\frac{\delta^2}{\sqrt{d}(nm)^{1/d}})$. The total number number of these cells is then $O(\frac{\sqrt{d}(nm)^{1/d}\Delta^d(n+m)}{\delta^2})$. This gives the following result:

Theorem 35. *Let P be a set of n and S a set of m d -dimensional simplices in \mathbb{R}^d , both having equal total volume and longest edge length at most Δ after normalising their total volumes to one, let $|\mu^*|$ be the cost of an optimal transport plan between them, and let δ be any constant > 0 . Given an algorithm that constructs a $(1 + \delta)$ -approximation between weighted sets of k points in $f_\delta(k)$ time, we can construct a transport plan between P and S with cost $\leq (1 + c'\delta)|\mu^*| + O(\frac{\sqrt{d}\delta}{(nm)^{1/d}})$ for some constant c' can be constructed in $O\left(f_\delta\left(\frac{\sqrt{d}(nm)^{1/d}\Delta^d(n+m)}{\delta^2}\right)\right)$ time.*

We can again calculate a $(1 + \delta)$ -approximation to ν in $O(N\delta^{-O(1)} \text{polylog } N)$ time using the algorithm by Fox and Lu [5], giving the following corollary to the previous theorem:

Corollary 36. *For any constant $\varepsilon > 0$, a transport plan between P and S with cost $\leq (1 + \varepsilon)|\mu^*| + O(\frac{\sqrt{d}\varepsilon}{(nm)^{1/d}})$ can be constructed in $O\left(\frac{\sqrt{d}(nm)^{1/d}\Delta^d(n+m)}{\varepsilon^c} \text{polylog}\left(\frac{d(nm)^{1/d}\Delta^d}{\varepsilon}\right)\right)$ time with high probability.*

10 Conclusion

We have provided approximation algorithms to the Earth Mover's Distance between sets of points, line segments, triangles and d -dimensional simplices. These are the first combinatorial algorithms with a provable approximation ratio for this problem when the objects are continuous rather than discrete points.

Here we described the case where the total mass is spread uniformly over the available length or area. However, our approach also works when this is not the case. If the ratio of densities is bounded, the same running times hold; otherwise, this ratio will show up in the running times the same way that the longest edge length does for cases involving triangles.

We note that for points and line segments (in any dimension), the approximation scheme is free from undesired parameters, whereas for points and triangles (or simplices), the maximum edge length Δ appears in the running time, and when neither set is a set of

points, an additive term appears in the approximation. The most interesting open question is whether either of these two artifacts can be avoided.

References

- [1] P. K. Agarwal, K. Fox, D. Panigrahi, K. R. Varadarajan, and A. Xiao. Faster algorithms for the geometric transportation problem. In *Proc. 33rd Int. Symp. on Comp. Geom.*, pages 7:1–7:16, 2017.
- [2] S. Cabello, P. Giannopoulos, C. Knauer, and G. Rote. Matching point sets with respect to the Earth Mover’s Distance. *Comp. Geom.*, 39(2):118–133, 2008.
- [3] F. de Goes, K. Breeden, V. Ostromoukhov, and M. Desbrun. Blue noise through optimal transport. *ACM Trans. on Graph.*, 31(6):171:1–171:11, 2012.
- [4] F. de Goes, D. Cohen-Steiner, P. Alliez, and M. Desbrun. An optimal transport approach to robust reconstruction and simplification of 2D shapes. *Comp. Graph. Forum*, 30(5):1593–1602, 2011.
- [5] K. Fox and J. Lu. A near-linear time approximation scheme for geometric transportation with arbitrary supplies and spread. In *Proc. 36th Int. Symp. on Comp. Geom.*, 2020. To appear.
- [6] D. Geiß, R. Klein, R. Penninger, and G. Rote. Optimally solving a transportation problem using Voronoi diagrams. *Comp. Geom.*, 46(8):1009–1016, 2013.
- [7] K. Grauman and T. Darrell. Fast contour matching using approximate Earth Mover’s Distance. In *Proc. IEEE Conf. Comp. Vision and Pattern Recognition*, pages 220–227, 2004.
- [8] I. van der Hoog, E. Khrantcova, and M. Löffler. Dynamic smooth compressed quadtrees. In *Proc. 34th Int. Symp. on Comp. Geom.*, pages 45:1–45:15, 2018.
- [9] L. V. Kantorovich. On the translocation of masses. *Dok. Akad. Nauk SSSR*, 37:199–201, 1942.
- [10] A. B. Khesin, A. Nikolov, and D. Paramonov. Preconditioning for the geometric transportation problem. In *Proc. 35th Int. Symp. on Comp. Geom.*, pages 15:1–15:14, 2019.
- [11] M. K. Kozlov, S. P. Tarasov, and L. G. Khachiyan. The polynomial solvability of convex quadratic programming. *USSR Comp. Math. and Math. Phys.*, 20(5):223–228, 1980.
- [12] H. Lavenant, S. Claiici, E. Chien, and J. Solomon. Dynamical optimal transport on discrete surfaces. In *SIGGRAPH Asia Technical Papers*, pages 250:1–250:16, 2018.
- [13] F. Mémoli. Spectral Gromov-Wasserstein distances for shape matching. In *Proc. 12th IEEE Int. Conf. on Comp. Vision Workshops*, pages 256–263, 2009.
- [14] Q. Mérigot. A multiscale approach to optimal transport. *Comp. Graph. Forum*, 30(5):1583–1592, 2011.
- [15] Q. Mérigot, J. Meyron, and B. Thibert. An algorithm for optimal transport between a simplex soup and a point cloud. *SIAM J. on Imaging Sci.*, 11(2):1363–1389, 2018.

- [16] G. Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences de Paris*, 1781.
- [17] Y. Rubner, C. Tomasi, and L. J. Guibas. The Earth Mover's Distance as a metric for image retrieval. *Int. J. of Comp. Vision*, 40(2):99–121, 2000.
- [18] R. Sharathkumar and P. K. Agarwal. A near-linear time ε -approximation algorithm for geometric bipartite matching. In *Proc. ACM Symp. on Theory of Comp.*, pages 385–394, 2012.
- [19] 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. *ACM Trans. on Graph.*, 34(4):66:1–66:11, 2015.
- [20] Z. Su, Y. Wang, R. Shi, W. Zeng, J. Sun, F. Luo, and X. Gu. Optimal mass transport for shape matching and comparison. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 37(11):2246–2259, 2015.
- [21] C. Villani. *Optimal Transport: Old and New*. Springer Verlag, Heidelberg, 2008.