Approximation Algorithm for Embedding Metrics into a Two-dimensional Space

Mihai Bădoiu*

Abstract

In this paper, we present a polynomial-time approximation algorithm for computing an embedding of an arbitrary metric into a two-dimensional space. The algorithm finds an embedding whose additive distortion is at most $c\epsilon^*$, where ϵ^* is the smallest additive distortion possible and c is an absolute constant. To our knowledge, this is the first result of this type, i.e., it gives an algorithm that finds (approximately) optimal embedding of a given distance matrix into a fixed d-dimensional space, where $d > 1$ is low, under *any* standard definition of embedding (see Related Work).

1 Introduction

Embedding of arbitrary distance matrices into geometric spaces (most notably, two-dimensional plane) is a fundamental problem occurring in many applications. In the context of data visualization, this approach allows the user to observe the structure of the data set and discover its interesting properties. In computational chemistry, this approach is used to recreate the geometric structure of the data from the distance information. Other application areas are discussed in [MDS].

The methods for computing such embeddings have their roots in work going back to the first half of the 20th century, and modern roots in work of Shepard [She62a, She62b], Kruskal [Kru64a, Kru64b], and others. This area is usually called *Multi-dimensional Scaling* (MDS) and is a subject of extensive research [MDS]. However, despite significant practical interests, very few theoretical results exist in this area. The most commonly used algorithms are heuristic (e.g., gradient-based method, simulated annealing, etc) and are often not satisfactory in terms of the running time and/or quality of the embeddings. The only theoretical result in this area [HIL98, Iva00] is a constant factor approximation algorithm for minimum distortion embedding into a *line*.

In this paper, we make progress towards an algorithmic theory of multi-dimensional scaling. In particular, we present a polynomial-time algorithm, which approximates a given distance matrix $D[\cdot,\cdot]$ by a matrix of distances induced by a set of points in a two-dimensional plane under l_1 norm. Specifically, consider $\epsilon = \min_f\{\max_{p,q}|D[p,q] - ||f(p)-f(q)||_1|\}$. Our algorithm computes f such that $\max_{p,q}|D[p,q] - ||f(p) - f(q)||_1| \le c\epsilon$. The constant c

*MIT Laboratory for Computer Science; 545 Technology Square, NE43-371; Cambridge, Massachusetts 02139-3594; mihai@theory.lcs.mit.edu

guaranteed by our algorithm is equal to 30. However, it is likely that it can be made smaller by a more careful analysis.

To our knowledge, this is the first algorithm that finds an (approximately) optimal embedding of a given distance matrix into a fixed d-dimensional space, where $d > 1$ is low, under *any* standard definition of embedding (see Related Work).

1.1 Related work

The two typical ways of measuring the distortion of embedding is to consider *multiplicative* error or *additive* error. The multiplicative distortion measures the ratio of the new distance over the old distance (we need to assume that the mapping is non-contractive). In the additive case, we take the absolute value of the difference between the new distance and the old distance. The overall distortion is usually computed as the maximum of distortion over all the pairs of points. However, different notions of overall distortion are not uncommon; for example, one can take the sum of distortions or the sum of squares of distortions over all the pairs of points.

As mentioned earlier, the only theoretical result on embedding metrics into a low-dimensional space has been obtained by [HIL98, Iva00]; they considered the maximum additive distortion. If one allows the host metric to have higher (say, logarithmic in n) dimensions, Bourgain showed in [Bou85] how to embed an arbitrary metric into $l_2^{O(\log^2 n)}$ with maximum multiplicative distortion $O(\log n)$. It should be noted, however, that for the applications mentioned above, the most interesting spaces are two- or three-dimensional.

In general, one can choose non-geometric metric spaces to serve as the host space. For example, in computational biology, approximating a matrix of distances between different genetic sequences by an ultrametric or a tree metric allows one to retrace the evolution path that led to formation of the genetic sequences. Motivated by these applications M. Farach-Colton and S. Kannan show how to find an *ultrametric* T with minimum possible maximum additive distortion [FCK96]. There is also an approximation algorithm for the case of tree metrics [ABFC$^+$96].

In addition, certain combinatorial results for low-distortion embeddings are known. For example, any metric can be embedded into d-dimensional Euclidean space with multiplicative distortion $O(n^{\frac{2}{d}} \log^{3/2} n)$ and no better than $\Omega(n^{1/\lfloor (d+1)/2 \rfloor})$ [Mat90]. However, in the applications, the most interesting metrics are those which can be embedded with much lower distortion.

Overview of the paper. In Section 3, we give an overview of the algorithm. In Section 4, we show how to solve the problem for a special case when we know the exact values of the x coordinates of the points and the value ϵ^* of the smallest error possible. In the appendix, in Section A, we show how to reduce our problem to the special case.

2 Preliminaries

Assume we are given a set P of n points and an $n \times n$ symmetric, positive and all-zero on the diagonal distance matrix D, which also satisfies the triangle inequality. The goal is to find an embedding $f : P \to \Re^2$ of the points into the plane, which minimizes the difference between the distances given by D and the distances given by the embedding. The distances in the plane are computed using the l_1 norm (or l_∞, which is isomorphic to l_1 in two dimensions).

Let ϵ^* be the optimal distortion. We guess ϵ^* by doing a binary search and we can assume we know its value. Given ϵ, let $f : P \to \Re^2$ be an embedding such that $\forall p, q \in P$,

$$|D[p,q] - ||f(p) - f(q)||_\infty| \leq \epsilon$$

Such an embedding exists for every $\epsilon \geq \epsilon^*$. To shorten the notation, we denote the x coordinate of $f(p)$ by p_1 and the y coordinate by p_2. Also, we write $||p - q||_\infty$ instead of $||f(p) - f(q)||_\infty$.

3 Overview of the algorithm

The algorithm works in two parts. In the first part, we approximate the x-coordinates of the embedding within $O(\epsilon^*)$. In the second part, assuming we know the approximate values of the x coordinates, we find the y values approximately.

The solution for the first part is easy in the case of the l_2 norm: we guess the diameter p, q, guess their placement, rotate the plane such that p and q are horizontal. Then we know that all v belong to the intersection of $\mathrm{Ball}(p, D[p,v] + \epsilon) - \mathrm{Ball}(p, D[p,v] - \epsilon)$ with $\mathrm{Ball}(q, D[q,v] + \epsilon) - \mathrm{Ball}(q, D[q,v] - \epsilon)$. This intersection gives little freedom to the x-coordinate of v, and we can guess it within $c\epsilon$ for a constant c. Unfortunately, the l_1 norm requires more elaborate techniques along these lines.

To do the second part we first find certain combinatorial structure of the point-set and then solve the problem using linear programming. Here we use properties of the l_∞ norm in a crucial way. We do not yet know how to do the second part in the case of the l_2 norm, but we believe a very similar method should work. In particular, we can roughly prove every lemma for l_2 except (4.4).

4 A special case

In this section, we are going to solve a special case in which we know the exact x coordinates of the points and ϵ for which there exists an embedding with distortion of at most ϵ. More exactly, we will compute a 5-approximation solution, i.e., we compute f such that $\forall p, q \in P$,

$$|D[p,q] - ||f(p) - f(q)||_\infty| \leq 5\epsilon$$

In the following sections we are going to reduce the main problem to this special case.

DEFINITION 4.1. *We connect two points p, q with an edge if $D[p,q] > |p_1 - q_1| + 3\epsilon$. We call such an edge a "strong" edge. We connect two points p, q with a "weak" edge if there is no strong edge between p and q and $D[p,q] > |p_1 - q_1| + \epsilon$.*

Let E be the set of the strong edges,

$$E = \{(p,q) | D[p,q] > |p_1 - q_1| + 3\epsilon\}$$

Intuitively, the strong edges are the edges we care about. Our goal is eventually to reduce the problem to linear programming. If there is no strong edge between two points p and q, by adding the constraint $-D[p,q] - \epsilon \leq q_2 - p_2 \leq D[p,q] + \epsilon$ we can ensure that the distance between p and q in our solution is less than $D[p,q] + \epsilon$. Also, since there is no strong edge between p and q, the distance already given by $|p_1 - q_1|$ is good enough for a 3-approximation solution.

Let $G = (P, E)$. If p, q, w are vertices in the same connected component of G, we also add to E the weak edges between the points v and w if $p_1 \leq v_1 \leq q_1$ and if v is not part of the component:

$$E' = E \bigcup \{(v, w) | \exists p, q \text{in the same connected component as } w \text{ of } G \text{ and } v \text{ is not in the same connected component as } w \text{ and } D[v, w] > |v_1 - w_1| + \epsilon \text{ and } p_1 \leq v_1 \leq q_1\}$$

Let $G' = (P, E')$.

Definition 4.2. *For an edge $(p, q) \in E'$, $p_1 \leq q_1$, we have two cases: $p_2 - q_2 \geq 0$ or $p_2 - q_2 < 0$* [1] *We say an edge is "oriented up" if it satisfies the first inequality and "down" if it satisfies the second inequality.*

The main idea of the algorithm is the following: We partition the elements of P into connected components of G'. We first note that if we know the orientation of all the strong edges, we can compute an embedding with distortion of at most $3\epsilon^*$ via linear programming. We also note that within each connected component, if we fix the orientation of a single edge, we can determine the orientation of all the other edges. Finally, we observe that any relative orientation of the edges between the connected components suffices in computing an embedding with distortion of at most $5\epsilon^*$.

Claim 4.1. *If p and q are in the same connected component of G, and $p_1 \leq q_1$, then every k for which $p_1 \leq k_1 \leq q_1$ is part of the same connected component of G' as p and q. Moreover, every weak edge in G' is adjacent to at least one strong edge.*

Proof. Since p and q are in the same connected component, there exists a path connecting p and q in G. Therefore, there exist v and w such that $(v, w) \in E$ and $v_1 \leq k_1 \leq w_1$. Since $(v, w) \in E$ we have

$$D[v, w] > |v_1 - w_1| + 3\epsilon$$

If $(v, k) \in E'$, then k is in the same connected component of G' as p and q. Otherwise, we have

$$D[v, k] \leq |k_1 - v_1| + \epsilon$$

By the triangle inequality we have $D[v, w] \leq D[v, k] + D[k, w]$. Combining these equations we get

$$D[k, w] \geq D[v, w] - D[v, k] > |v_1 - w_1| + 3\epsilon - |k_1 - v_1| - \epsilon = |k_1 - w_1| + 2\epsilon \tag{4.1}$$

which means $(k, w) \in E'$.

Moreover, every edge added has an adjacent edge from E, which is a strong edge. Note that we do not add all the weak edges to E'.

[1] Note that the fact that $(p, q) \in E'$ and the first inequality implies $p_2 - q_2 > q_1 - p_1$. Also, $(p, q) \in E'$ and the second inequality implies $p_2 - q_2 > q_1 - p_1$.

Definition 4.3. *We say that two connected components of G' overlap if and only if a) there is no vertical line l that separates the elements of the first component from the elements of the second component, and b) l does not intersect any point.*

Claim 4.1 reveals the structure of G'. More specifically, no two connected components overlap. This structure is exactly the desired one. We do not want to have strong edges between the connected components, and we want them not to overlap, such that we can guess the orientation of each component. Note that we do not care if we can have weak edges between the components.

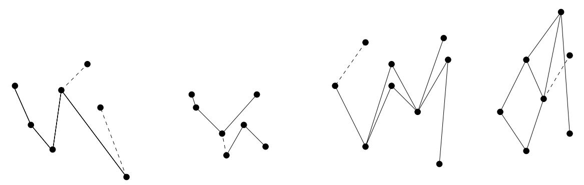

Figure 1: The structure of G': G' has 4 connected components. Strong edges are shown with solid lines, and weak edges with dotted lines. There are no strong edges between the connected components. Components do not "overlap." Each weak edge is adjacent to at least one strong edge.

For an edge $(v, w) \in E'$ which is oriented up, such that $v_1 \leq w_1$, we have

$$w_2 - v_2 + \epsilon \geq D[v, w] \geq w_2 - v_2 - \epsilon \tag{4.2}$$

Claim 4.2. *By fixing the orientation of an edge of G' we also fix the orientation of all the other edges in the same connected component of G'.*

Proof. We first show that if we know the orientation of an edge e, then we can also determine the orientation of any adjacent edge if both e and the adjacent edge are not weak edges. By setting the orientation of an edge and repeating this process we can determine the orientation of all the edges in the connected component.

Without loss of generality, we assume that the edge $(v, w) \in E'$ ($w_1 \geq v_1$) is oriented up: $w_2 - v_2 > w_1 - v_1$. Also let $(w, t) \in E'$.

If (w, t) is a strong edge and (w, t) is oriented up, by using (4.2) multiple times we get

$$D[v,t] \geq |t_2 - v_2| - \epsilon \geq |w_2 - v_2| + |t_2 - w_2| - \epsilon \geq D[v,w] + D[w,t] - 3\epsilon$$

Since (w, t) is a strong edge, $D[w, t] > 3\epsilon$, therefore

$$D[v,t] > D[v,w] \tag{4.3}$$

Using the fact that (v, w) is an edge and (w, t) is a strong edge (i.e., $D[v, w] > w_1 - v_1 + \epsilon$ and $D[w, t] > |t_1 - w_1| + 3\epsilon$), we get

$$D[v,t] > (w_1 - v_1) + |t_1 - w_1| + \epsilon > |t_1 - v_1| + \epsilon \tag{4.4}$$

If (w,t) is a strong edge and (w,t) is oriented down, we have

$$D[v,t] \leq ||t-v||_\infty + \epsilon \leq \max\{|t_1 - v_1| + \epsilon, D[v,w] + \epsilon - D[w,t] + \epsilon\} \\ \leq \max\{|t_1 - v_1| + \epsilon, D[v,w]\} \tag{4.5}$$

Since equations (4.3) and (4.4) contradict equation (4.5), we can determine whether (w,t) is oriented up or down.

If (w,t) is a weak edge and (v,w) is a strong edge, then by a similar argument we can determine the orientation of (w,t). If a connected component of G' contains only one connected component of G, we can determine the orientation of all the strong edges first and then the orientation of the weak edges. If a connected component of G' contains two connected components of G, then these two components must overlap (by the way we add weak edges), which means that there is a weak edge connecting them. This means that we can determine the orientation of the strong edges in the first component, then the orientation of a weak edge between the two components, then the orientation of the strong edges of the second component, and finally the orientation of all the remaining weak edges. The same argument applies to the case when the connected component of G' is composed of several connected components of G.

CLAIM 4.3. *Given the orientation of all the strong edges, we can compute a* 3*-approximation solution via linear programming.*

Proof. We construct the following linear program:

$$\text{Min} \qquad \delta$$

subject to

$$\begin{array}{ll} D[p,q] - \delta \geq q_2 - p_2 \geq D[p,q] + \delta, & \text{if } (p,q) \in E \text{ is oriented up and } q_1 \geq p_1 \\ D[p,q] - \delta \geq p_2 - q_2 \leq D[p,q] + \delta, & \text{if } (p,q) \in E \text{ is oriented down and } q_1 \geq p_1 \\ -D[p,q] - \delta \leq q_2 - p_2 \leq D[p,q] + \delta, & \text{if } (p,q) \notin E \end{array}$$

First note, that if we were to have all the edges (including the weak ones) in E, we would get an optimal solution. However, we know only the orientation of the strong edges, and this gives a 3-approximation solution: It is clear that a 3-approximation solution that satisfies the orientation given is a feasible solution for this linear program. It is also clear that any solution of this program is an embedding with error of at most $\max(\delta^*, 3\epsilon)$ and $\delta^* \leq 3\epsilon$.

By using claims 4.2 and 4.3 we can get an approximate solution to the problem. But what if we have several connected components in G'? We will prove that no matter what relative orientation we take between the edges of the components, we will still get a constant approximation solution. So, the algorithm is as follows: we choose an arbitrary orientation to one edge from each connected component, and by using claim 4.2 and 4.3 we get an approximate solution.

CLAIM 4.4. *There is a* 5*-approximation solution for every relative orientation between the edges of the components.*

Proof. Let f be an optimal embedding. Let $C_1, C_2, C_3, \ldots, C_k$ be the connected components of G', from the ones with the smallest x coordinate to the ones with higher x coordinate: $\forall i$, if $v \in C_i$ and $w \in C_{i+1}$, then $v_1 \leq w_1$.

Choose any relative orientation of the components. Let the function $s : \{1, 2, \ldots, k\} \to \{0, 1\}$ denote the relative orientation of our arbitrary choice to the optimal embedding f: $s(i) = 1$ if the orientation of the component C_i is different in f than in our arbitrary selection; $s(i) = 0$ otherwise.

We are going to start with an optimal solution and modify it to get a feasible solution that has the given relative orientation, with error of at most 5ϵ:

We are going to modify f incrementally from the first component to the last. Without loss of generality we can assume $s(1) = 0$. (If $s(1) = 1$ we can flip (or reflect) f by the x axis, flipping each $s(i)$ and still having an optimal embedding.)

We repeat the following steps for i from 2 to k:

If $s(i) = 0$, then we go to the next component C_{i+1}.

If $s(i) = 1$, then we will flip all the points in $\bigcup_{j=i}^{k} C_j$ by a certain line parallel to the x axis. This flip will change the values of $s(j)$ for all $j \geq i$. The line by which we are going to flip the points is computed as follows:

- Let p be a point in $\bigcup_{j=1}^{i-1} C_j$ that maximizes $p_1 + p_2$: $\forall v \in \bigcup_{j=1}^{i-1} C_j, p_1 + p_2 \geq v_1 + v_2$.
- Let q be a point in $\bigcup_{j=1}^{i-1} C_j$ that maximizes $q_1 - q_2$: $\forall v \in \bigcup_{j=1}^{i-1} C_j, q_1 - q_2 \geq v_1 - v_2$.
- Let l_1 be the line of slope $-\pi/4$ that passes through p, and let l_2 be the line of slope $\pi/4$ that passes through q. Let the point r denote the intersection of l_1 with l_2. Let l be the line through r that is parallel to the x axis.
- Flip the points in $\bigcup_{j=i}^{k} C_j$ by l: for all $v \in \bigcup_{j=i}^{k} C_j$, $v'_2 = 2r_2 - v_2$

It is easy to see that by performing this flip operation on f we change only the distances between v and w for $v \in \bigcup_{j=1}^{i-1} C_j$ and $w \in \bigcup_{j=i}^{k} C_j$. The question is: by how much? Since there are no strong edges between the components, for such v and w, we have:

$$D[v, w] \leq |v_1 - w_1| + 3\epsilon \tag{4.6}$$

If $v \in C_{i-1}$, then by the flips up to this step, the distance between v and w has remained the same as in the original embedding:

$$D[v, w] + \epsilon \geq ||v - w||_\infty \geq |w_2 - v_2| \tag{4.7}$$

Combining equations (4.6) and (4.7) we get

$$|w_2 - v_2| \leq D[v, w] + \epsilon \leq |w_1 - v_1| + 4\epsilon \tag{4.8}$$

Let v' be the point 4ϵ to the left of v: $v'_1 = v_1 - 4\epsilon$ and $v'_2 = v_2$. Let d' be the line of slope $\pi/4$ that goes through v' and d'' be the line of slope $-\pi/4$ that goes through v'.

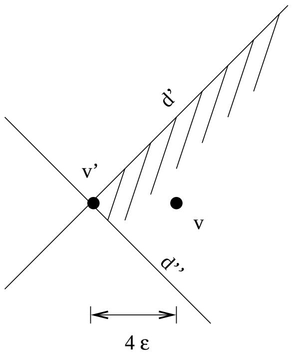

Figure 2: The points $w \in \bigcup_{j=i}^{k} C_j$ are located beneath d' and above d''.

It is easy to see that equation (4.8) implies w is located beneath d' and above d''. (See Figure 2)

This holds for each $v \in C_{i-1}$. Therefore each $w \in \bigcup_{j=i}^{k} C_j$ is located between these wedges. Now, our line l is chosen such that after we flip the points $w \in \bigcup_{j=i}^{k} C_j$, they will remain between these wedges, such that equation (4.8) remains true after the flip. [2] Since we do not change the x coordinates and the original f is a solution with error ϵ we know that

$$D[v, w] + \epsilon \geq |w_1 - v_1| \tag{4.9}$$

Combining equations (4.8) and (4.9), we get

$$|w_2 - v_2| \leq |w_1 - v_1| + 4\epsilon \leq D[v, w] + 5\epsilon$$

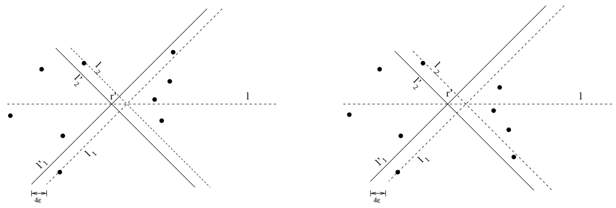

Figure 3: The points in $\bigcup_{j=i}^{k} C_j$ will remain between the two lines (wedges) l'_1 and l'_2 after the flip.

In addition, since for the next flips, the points that are being flipped are inside an even more restrictive space between two wedges, they will not leave the space between the two wedges and equation (4.8) will remain true after we have completed all the flips. Again, since

[2]The intersection of the space between pairs of wedges has the same shape as the space between 2 wedges and l divides this intersection into 2 symmetrically equal pieces.

equation (4.9) is also true, we can combine them and get that $|w_2 - v_2| \leq D[v,w] + 5\epsilon$ after all the flips are completed. By applying this argument to $v \in C_1, v \in C_2, \ldots,$ we conclude that for all points v, w, $w_1 \geq v_1$, we have $|w_2 - v_2| \leq D[v,w] + 5\epsilon$. Since we have no strong edges between the components, for $v \in C_i$, $w \in C_j$, $i \neq j$, $|v_1 - w_1| \geq D[v,w] - 3\epsilon$. This implies that for two points from different components, the distortion in our construction is at most 5ϵ. Since we preserve the distances between points from the same component, we have constructed an embedding with distortion of at most 5ϵ for an arbitrary relative orientation of the edges between the connected components.

Finally, we apply claim 4.3 to produce a 5-approximation solution. We know that for our orientation there is an embedding with error 5ϵ and this is a feasible solution for our linear program.

5 The Final Algorithm

So far, we assumed we know certain points (the pair of points that give the diameter, etc). To satisfy this assumption, we will iterate over all possible choices (a polynomial number of such choices). The total running time of the algorithm is $O(\log \frac{diam}{\epsilon^*} n^6 LP)$ where $diam$ is the value of the diameter of P and LP is the time to solve a linear program with n variables and $O(n)$ constraints. Thus, the running time is polynomial in n.

6 Conclusions

In this paper, we show how to approximate within a constant factor an embedding of an arbitrary metric into a two-dimensional space where distances are computed using the l_1 norm with the notion of an additive error ϵ. Our constant is 30, but by combining the general case with the special case more carefully, we can get the constant down to 19.

Future Work We believe the distortion and the running time can be improved further. We also believe that the same technique might be extended to get the same result for other norms (i.e., l_2) or multiplicative error. It would also be of interest to extend this result for higher dimensions - let's say three-dimensional space. It should also be noticed that in the case of the l_2 norm, while we don't know how to prove claim (4.4) with additive error, it is easy to prove it with multiplicative error and we can obtain an embedding $f : P \to \Re^2$, for which $D[p,q] \leq ||f(p) - f(q)||_2 \leq aD[p,q] + b\epsilon^*$, where a and b are absolute constants.

Acknowledgements The author would like to thank Piotr Indyk and Sariel Har-Peled for numerous comments on early drafts of this paper and in particular to **Piotr Indyk** for useful discussions.

References

[ABFC+96] R. Agarwala, V. Bafna, M. Farach-Colton, B. Narayanan, M. Paterson, and M. Thorup. On the approximability of numerical taxonomy: (fitting distances by tree metrics). *7th Symposium on Discrete Algorithms*, 1996.

[Bou85] J. Bourgain. On lipschitz embedding of finite metric spaces into hilbert space. *Isreal Journal of Mathematics*, 52:46–52, 1985.

[FCK96] M. Farach-Colton and S. Kannan. Efficient algorithms for inverting evolution. *28th Symposium on Theory of Computing*, 1996.

[HIL98] J. Hastad, L. Ivansson, and J. Lagergren. Fitting points on the real line and its application to rh mapping. *Lecture Notes in Computer Science*, 1461:465–467, 1998.

[Ind01] P. Indyk. Algorithmic applications of low-distortion geometric embeddings. *Proceedings of the 42nd Annual IEEE Symposion on Foundations of Computer Science*, pages 10–31, 2001.

[Iva00] L. Ivansson. Computational aspects of radiation hybrid. *Doctoral Dissertation, Department of Numerical Analysis and Computer Science, Royal Institute of Technology*, 2000.

[Kru64a] J.B. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29:1–27, 1964.

[Kru64b] J.B. Kruskal. Nonmetric multidimensional scaling: A numerical method. *Psychometrika*, 29:115–129, 1964.

[Mat90] J. Matoušek. Bi-lipschitz embeddings into low-dimensional euclidean spaces. *Comment. Math. Univ. Carolinae*, 31:589–600, 1990.

[Mat02] J. Matoušek. *Lectures in Discrete Geometry.* Springer-Verlag, 2002.

[MDS] Algorithms for multidimensional scaling.
http://dimacs.rutgers.edu/SpecialYears/2001_Data/Algorithms/MDSdescription.html.

[She62a] R. N. Shepard. The analysis of proximities: Multidimensional scaling with an unknown distance function 1. *Psychometrika*, 27:125–140, 1962.

[She62b] R. N. Shepard. The analysis of proximities: Multidimensional scaling with an unknown distance function 2. *Psychometrika*, 27:216–246, 1962.

A The general case

The main idea is to fix the x coordinates of the points and then to use the algorithm from the previous section. We do not need to guess the x coordinates exactly. If we guess them within $c\epsilon$ for a constant c, it will be enough to get a constant approximation algorithm for the general case.

Let p and q denote the diameter. Let f be an optimal embedding with error ϵ. Without loss of generality we have $p = (0,0)$ and we assume that the diameter is given by $q_1 - p_1$.

Let A be the set of points $v \in P - \{p,q\}$ for which the following equation is true:

$$D[p,q] + k\epsilon \geq D[p,v] + D[v,q]$$

$$A = \{v \in P - \{p,q\} : D[p,q] + k\epsilon \geq D[p,v] + D[v,q]\}$$

for a fixed constant k which we will chose later. [3]

Note that for $v \in A$, we have the following two inequalities:

$$v_1 \leq D[p,v] + \epsilon$$

$$v_1 \geq D[p,q] - D[v,q] - 2\epsilon \geq D[p,v] - (k+2)\epsilon$$

It follows that by fixing $v_1 = \frac{2D[p,v]-(k+1)\epsilon}{2} = D[p,v] - \frac{(k+1)\epsilon}{2}$ we are within an additive factor of $\frac{(k+3)\epsilon}{2}$ from the value of v_1 in the optimal embedding f.

It is clear that if $A = P - \{p,q\}$ then we can guess the x coordinate of all the points in A and then by reducing the problem to the special case with $\epsilon' = \frac{(k+3)\epsilon}{2}$ we get a $\frac{5(k+3)}{2}$-approximation algorithm.

But what happens if $P - A \neq \emptyset$? We break our analysis into two cases:

A.1 Case 1

For this case we assume that for all points $v \in P - A - \{p,q\}$ we have either $||p-v||_\infty = |p_1 - v_1|$ or $||q-v||_\infty = |q_1 - v_1|$.

Partition the points of $P - A - \{p,q\}$ into three sets:

$$B = \{v \in P - A - \{p,q\} : ||p-v||_\infty = |p_1 - v_1|\}$$

$$C = \{v \in P - A - \{p,q\} : ||q-v||_\infty = |q_1 - v_1| \text{ and } v_2 - p_2 \geq 0\}$$

$$D = \{v \in P - A - \{p,q\} : ||q-v||_\infty = |q_1 - v_1| \text{ and } v_2 - p_2 < 0\}$$

Note that this is a partition: If $B \bigcap C \neq \emptyset$ then there exists $v \in B \bigcap C$, and we have $D[p,q] \geq q_1 - p_1 - \epsilon = q_1 - v_1 + v_1 - p_1 - \epsilon = ||q-v||_\infty + ||v-p||_\infty - \epsilon \geq D[q,v] + D[p,v] - 3\epsilon$ which implies $v \in A$, which is a contradiction.

The idea is to choose certain points to decide for each point $v \in P - A - \{p,q\}$ if $v \in B$ or $v \in C \bigcup D$. If we can decide that, we can approximate its x coordinate within an additive error of $c\epsilon$ for some constant c.

Let the point $p' \in C$ be such that $p' = \min_{v \in C} v_1 + v_2$ and the point $p'' \in D$ such that $p'' = \min_{v \in D} v_1 - v_2$. Of course, p' or p'' may not even exist, but those cases are easier to handle and the proof is basically the same for them as well.

[3] k will be chosen to be any constant greater than 9.

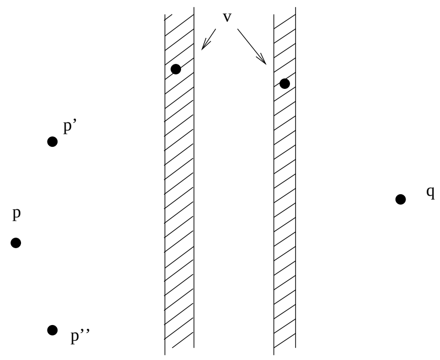

Figure 4: If $v \in B$ then v is restricted to the right stripe of width 2ϵ. If $v \in C \bigcup D$ then v is restricted to the left stripe of width 4ϵ. Since v is not in A, we know that the distance between the stripes is at least $(k-2)\epsilon$

Since $p' \notin A$, we have

$$D[p', q] > D[p, q] + k\epsilon - D[p, p'] \tag{1.10}$$

Since $p' \in C$, we have the following inequalities

$$p'_2 - \epsilon \leq D[p', p] \leq p'_2 + \epsilon \tag{1.11}$$

$$D[p, q] - D[p', q] - 2\epsilon \leq p'_1 \leq D[p, q] - D[p', q] + 2\epsilon \tag{1.12}$$

By combining (1.10) with (1.12) we get that $p'_1 < D[p, p'] - (k-2)\epsilon$. Using (1.11) we get that

$$p'_1 < p'_2 - (k-3)\epsilon \tag{1.13}$$

For a point $v \in B$, we have

$$\begin{aligned} D[p', v] &\geq |v_1 - p'_1| - \epsilon \\ &\geq D[p, v] - 2\epsilon - p'_1 \\ &> D[p, v] - p'_2 + (k-5)\epsilon && \text{by (1.13)} \\ &\geq D[p, v] - D[p, p'] + (k-6)\epsilon && \text{by (1.11)} \end{aligned} \tag{1.14}$$

Also,

$$\begin{aligned} D[p', v] &\geq |v_1 - p'_1| - \epsilon \\ &\geq D[p, v] - 2\epsilon - p'_1 \\ &\geq D[p, v] - D[p, q] + D[p', q] - 4\epsilon && \text{by (1.12)} \end{aligned} \tag{1.15}$$

For a point $v \in C$, we have $D[p', v] \leq ||p', v||_\infty + \epsilon \leq \max(v_1 - p'_1, v_2 - p'_2) + \epsilon$ by the way p' was chosen.

If $||p', v||_\infty = v_2 - p'_2$, we have

$$\begin{aligned} D[p', v] &\le v_2 - p'_2 + \epsilon \le D[p, v] - p'_2 + 2\epsilon \\ &\le D[p, v] - D[p, p'] + 3\epsilon \qquad \text{by (1.11)} \end{aligned} \tag{1.16}$$

Note that if $k > 9$, equations (1.14) and (1.16) are contradictory.

If $D[p', v] = v_1 - p'_1 + \epsilon$, we have

$$\begin{aligned} D[p', v] = v_1 - p'_1 + \epsilon &\le D[p, q] - D[q, v] - p'_1 + 3\epsilon \\ &< D[p, v] - (k-3)\epsilon - p'_1 && \text{by (1.10)} \\ &\le D[p, v] - (k-3)\epsilon - (D[p, q] - D[p', q] - 2\epsilon) && \text{by (1.12)} \\ &= D[p, v] - D[p, q] + D[p', q] - (k-5)\epsilon \end{aligned} \tag{1.17}$$

Note that if $k > 9$, equations (1.17) and (1.15) are contradictory.

We can also obtain similar equations for $v \in D$ by replacing p' with p'' in the above argument. We use these observations to prove the following claim:

Claim A.1. *We can determine which points are in B and which ones are in $C \bigcup D$ if $k > 9$.*

Proof. If $v \in B$, by (1.14) and (1.15), the following equations are true:

$$D[p', v] \ge D[p, v] - D[p, p'] + (k-6)\epsilon$$

$$D[p', v] \ge D[p, v] - D[p, q] + D[p', q] - 4\epsilon$$

$$D[p'', v] \ge D[p, v] - D[p, p''] + (k-6)\epsilon$$

$$D[p'', v] \ge D[p, v] - D[p, q] + D[p'', q] - 4\epsilon$$

If $v \in C \bigcup D$ then either (1.16) or (1.17) is true (for p' if $v \in C$ or for p'' if $v \in D$) which implies that at least one of the above 4 equations is false. Therefore we say that $v \in B$ if all the about 4 equations are true, and $v \in C \bigcup D$ otherwise.

A.2 Case 2 For this case we assume that there exists a point $r \in P - \{p, q\}$ for which the following is true: $||p - r||_\infty > |p_1 - r_1|$ and $||q - r||_\infty > |q_1 - r_1|$. It follows that $||p - r||_\infty = |p_2 - r_2|$ and $||q - r||_\infty = |q_2 - r_2|$.

Let r be such a point that maximizes $|r_2 - p_2|$:

$$r = \max_{\substack{v \in P \\ ||p-r||_\infty = |p_2 - r_2| \\ ||q-r||_\infty = |q_2 - r_2|}} |r_2 - p_2|$$

For this case, we will fix the y coordinate instead. The method we will use is going to be very similar to the method used for case 1. We partition the points of $P - \{p, q, r\}$ into four sets as follows:

$$A = \{v \in P - \{r, p, q\} : D[r, p] + k\epsilon \ge D[r, v] + D[v, p] \text{ or } D[r, q] + k\epsilon \ge D[r, v] + D[v, q]\}$$

$$B = \{v \in P - A - \{r, p, q\} : ||r - v||_\infty = |r_2 - v_2|\}$$

$$C = \{v \in P - A - \{r, p, q\} : ||r - v||_\infty = |r_1 - v_1| \text{ and } v_1 - r_1 \ge 0\}$$

$$D = \{v \in P - A - \{r, p, q\} : ||r - v||_\infty = |r_1 - v_1| \text{ and } v_1 - r_1 < 0\}$$

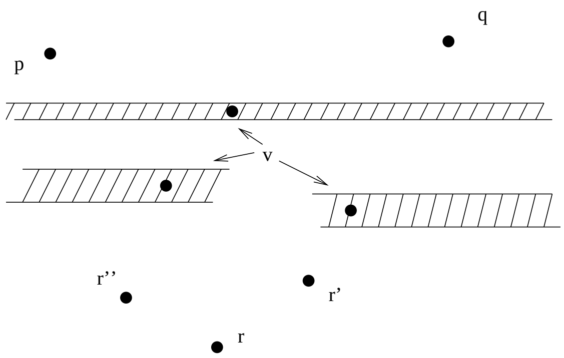

Figure 5: If $v \in B$ then v is restricted to the upper stripe of width 2ϵ. If $v \in C$ then v is restricted to the lower right stripe of width 4ϵ. If $v \in D$ then v is restricted to the lower left stripe of width 4ϵ

CLAIM A.2. *If $v \in C$ we can determine its y coordinate within an additive factor of 2ϵ*

Proof. By the way of contradiction, we suppose that $||q - v||_\infty = q_1 - v_1$. Then $D[v, q] \geq q_1 - v_1 - \epsilon$. It follows that $D[r, q] \geq q_1 - r_1 - \epsilon \geq q_1 - v_1 + v_1 - q_1 - \epsilon \geq D[r, v] + D[v, q] - 3\epsilon$. But for $k \geq 3$, this implies $v \in A$, which contradicts the fact that $v \in C$.

Therefore, we know that $||q - v||_\infty = |q_2 - v_2|$. If $v_2 > q_2$ we have that $v_1 > v_2 > q_2 > q_1$ which means that the diameter is given by the $|p_1 - v_1|$, again impossible. Therefore,

$$||q - v||_\infty = q_2 - v_2$$

Using this observation, we have the following two bounds on v_2: $D[r, q] - D[r, v] - 2\epsilon \leq v_2 \leq D[r, q] - D[r, v] + 2\epsilon$ and we can guess v_2's real value within an additive distance of 2ϵ by setting $v_2 = D[r, q] - D[r, v]$.

Similarly by replacing q with p in the above proof, if $v \in D$ we can also determine its y coordinate within an additive factor of 2ϵ.

We shift everything and flip it by the y axis if necessary such that, $r = (0, 0)$ and $p_2 > 0$. Note that this implies that $q_2 > 0$ (if $q_2 < 0$ then we would have $|p_2 - q_2| = |p_2 - r_2| + |r_2 - q_2| > |p_1 - r_1| + |r_1 - q_1| = |p_1 - q_1| = diam(P)$, which is a contradiction) and that every $v \in B$ has $v_2 \geq 0$ (because we chose r such that it maximizes $r_2 - p_2$). We proceed as before: we pick certain points to help us decide for each point in which set it belongs to. If we can decide that, we can approximate for each point its y coordinate within a constant times ϵ.

Let the point $r' \in C$ be such that $r' = \min_{v \in C} v_1 + v_2$ and the point $r'' \in D$ such that $r'' = \min_{v \in D} v_1 - v_2$. Of course, as before, r' or r'' may not even exist, but those cases are easier to handle and the proof is basically the same for them as well.

Since $r' \notin A$, we have

$$D[r', q] > D[r, q] + k\epsilon - D[r, r'] \tag{1.18}$$

Since $r' \in C$, we have the following inequalities

$$r'_1 - \epsilon \leq D[r', r] \leq r'_1 + \epsilon \tag{1.19}$$

$$D[r, q] - D[r', q] - 2\epsilon \leq r'_2 \leq D[r, q] - D[r', q] + 2\epsilon \tag{1.20}$$

By combining (1.18) with (1.20) we get that $r'_2 < D[r, r'] - (k-2)\epsilon$. Using (1.19) we get that

$$r'_2 < r'_1 - (k-3)\epsilon \tag{1.21}$$

For a point $v \in B$, we have

$$\begin{aligned} D[r', v] &\geq |v_2 - r'_2| - \epsilon \\ &\geq D[r, v] - 2\epsilon - r'_2 \\ &> D[r, v] - r'_1 + (k-5)\epsilon && \text{by (1.21)} \\ &\geq D[r, v] - D[r, r'] + (k-6)\epsilon && \text{by (1.19)} \end{aligned} \tag{1.22}$$

Also,

$$\begin{aligned} D[r', v] &\geq |v_2 - r'_2| - \epsilon \\ &\geq D[r, v] - 2\epsilon - r'_2 \\ &\geq D[r, v] - D[r, q] + D[r', q] - 4\epsilon && \text{by (1.20)} \end{aligned} \tag{1.23}$$

For a point $v \in C$, we have $D[r', v] \leq \max(v_2 - r'_2, v_1 - r'_1) + \epsilon$ by the way r' was chosen. If $||r', v||_\infty = v_1 - r'_1$, we have

$$\begin{aligned} D[r', v] \leq v_1 - r'_1 + \epsilon &\leq D[r, v] - r'_1 + 2\epsilon \\ &\leq D[r, v] - D[r, r'] + 3\epsilon && \text{by (1.19)} \end{aligned} \tag{1.24}$$

Note that if $k > 9$, equations (1.22) and (1.24) are contradictory.
If $||r', v||_\infty = v_2 - r'_2$, we have

$$\begin{aligned} D[r', v] \leq v_2 - r'_2 + \epsilon &\leq D[r, q] - D[q, v] - r'_2 + 3\epsilon \\ &< D[r, v] - (k-3)\epsilon - r'_2 && \text{by (1.18)} \\ &\leq D[r, v] - (k-3)\epsilon - (D[r, q] - D[r', q] - 2\epsilon) && \text{by (1.20)} \\ &= D[r, v] - D[r, q] + D[r', q] - (k-5)\epsilon \end{aligned} \tag{1.25}$$

Note that if $k > 9$, equations (1.25) and (1.23) are contradictory.
We also have very similar equations for $v \in D$ by using r''. We use these observations to prove the following claim.

As before, we can use these equations to determine which points are in B and which ones are in $C \bigcup D$. However, in this case, we should also distinguish between C and D. We make the following observation:

Claim A.3. *If* $D[q,v]+D[v,r'] \leq D[r',q]+3\epsilon$ *or* $D[p,v]+D[v,r''] \leq D[r'',p]+3\epsilon$ *we can approximate* v_2 *within a factor of* 3ϵ*. Otherwise, if* $v \in C$*, then* $D[r',v] < D[r'',v]$ *and if* $v \in D$ *then* $D[r',v] > D[r'',v]$*.*

Proof. First note that $D[r',r'']$ has to be pretty large:

$$\begin{aligned}
D[r',r''] &\geq r'_1 - r_1 + r_1 - r''_1 - \epsilon \\
&\geq D[r',r] + D[r'',r] - 3\epsilon \\
&> D[r,q] - D[r',q] + D[r,p] - D[r'',p] + (2k-3)\epsilon \\
&\geq q_2 - r_2 - \epsilon - (q_2 - r'_2 + \epsilon) + p_2 - r_2 - \epsilon - (q_2 - r''_2 + \epsilon) + (2k-3)\epsilon \\
&= r'_2 - r_2 + r''_2 - r_2 + (2k-7)\epsilon \\
&\geq (2k-7)\epsilon \qquad (1.26)
\end{aligned}$$

We break our proof into three cases:

- Case 1: $D[q,v]+D[v,r'] \leq D[r',q]+3\epsilon$

  First note that we know that

  $$v_2 \geq q_2 - D[q,v] - \epsilon \geq D[r,q] - D[q,v] - 2\epsilon \qquad (1.27)$$

  We also have that

  (1.28)
  $$v_2 \leq r'_2 + D[v,r'] + \epsilon \leq D[r,q] - D[q,r'] + D[v,r'] + 3\epsilon \leq D[r,q] - D[q,v] + 6\epsilon$$

  Therefore, as before, we can approximate v_2 within an additive error of 4ϵ by setting $v_2 = D[r,q] - D[q,v] + 2\epsilon$.

- Case 2: $D[p,v]+D[v,r''] \leq D[r'',p]+3\epsilon$

  The analysis of this case is analogous to the one above. (replace r' by r'' and q by p).

- Case 3: If $||r'-v||_\infty = v_2 - r'_1$ we have that $D[q,v]+D[v,r'] \leq D[r',q]+3\epsilon$, which falls into case 1. Therefore, we know that $||r'-v||_\infty = v_1 - r'_1$. In that case,

$$\begin{aligned}
D[r'',v] &\geq v_1 - r'_1 + r'_1 - r''_1 - \epsilon \\
&\geq D[v,r'] + r'_1 - r''_1 - 2\epsilon \\
&\geq D[v,r'] + D[r',r''] - 3\epsilon && \text{since } ||r',r''||_\infty = r'_1 - r''_1 \\
&> D[v,r'] + (2k-10)\epsilon && \text{by (1.26)} \qquad (1.29)
\end{aligned}$$

Therefore, for $v \in C$ we have $D[r'',v] > D[v,r']$ if $k \geq 5$. Symmetrically, for $v \in D$ we have $D[r'',v] < D[v,r']$ for $k \geq 5$.

We conclude that it is either the case that we can approximate v_2 within 4ϵ (case 1 or 2) or we can compare $D[v,r'']$ with $D[v,r']$ to determine if $v \in C$ or $v \in D$ which implies we can approximate v_2 within additive error 2ϵ.

Using this observation we can easily distinguish between points in C and points in D. For the points in $v \in C$ we fix the y coordinate as $v_2 = D[r,q] - D[q,v]$ and for the points $v \in D$, $v_2 = D[r,p] - D[p,v]$.