Low distortion maps between point sets

Sylvain Biehler, Claire Kenyon, Yuval Rabani and Alistair Sinclair

- Decide whether two sets of points in Euclidian space are isometric
- Decide whether two graphs are isomorphic
- Decide whether two hand-drawn characters represent the same character

An approximate notion of isometry

- Minimum distortion bijection between the two sets
- **Distortion:** The infimum $\alpha\beta$ such that there exists a bijection σ such that σ and σ^{-1} expand distances by at most α and β respectively.

$$\forall x, y, \quad d'(\sigma(x), \sigma(y)) \leq \alpha d(x, y)$$
$$\forall x', y', \quad d(\sigma^{-1}(x), \sigma^{-1}(y)) \leq \beta d'(x, y)$$

Also known as a biLipschitz map

Some nice properties

- Symmetric
- Invariant by scaling – people often scale so that $\beta' = 1$ and $\alpha' = \alpha\beta$
- Always greater than or equal to 1
- Equals 1 iff the two sets are isometric
- For graph vertices and shortest path distances, equals 1 iff the two graph are isomorphic

Related: low distortion embeddings

- **Embed:** a given source set of points into a target space with small distortion
- **Tool:** Embedding into a simpler metric space (small dimension) is a useful tool in algorithm design
- **Embed into line** Badoiu Dhamdere Gupta Rabinovich Racke Ravi Sidiropoulos SODA 2005, Badoiu Chuzhoy Indyk Sidiropoulos STOC 2005
- **Matching hand-written characters:** Belongie Malik Puzicha IEEE Trans. on Pattern Analysis and Machine Intelligence 2002
- **Low distortion embeddings between point sets** Kenyon Rabani Sinclair STOC 2004, Biehler +KRS in preparation

Matching hand-written characters

- Goal: robust and simple algorithm to find correspondences between shapes
- **Context of a pixel p:** Write other pixels' positions in polar coordinates with p at the origin. Discretize to define buckets. Record how many pixels are in each bucket: histogram.
- **Mapping p to q:** cost is defined using the χ^2-statistic between the distributions defined by the histograms
- **Algorithm:** min cost bipartite graph perfect matching
- **Evaluation:** fast and works well in practice, but lacks theoretical justification and is not scale-invariant
- **Observe:** that the most relevant problem is to find a correspondence when the two sets are quite similar
- **Remark:** sampling guarantees that the two sets have equal size

Work of Badoiu Dhamdere et al. (BDGRRRS)

- Focuses on embeddings into the line
- **Main result:** Given an unweighted graph and shortest path distances, a $O(\sqrt{n})$ approximation algorithm to find a low distortion embedding into the line
- **For graphs embeddable with distortion c:** An embedding with distortion $O(c^2)$ obtainable in time $O(n^3c)$
- **Main difference:** in the model: they are free to place the images wherever they want on the line

Basic Lemma (Matousek 1990)

- Any shortest path metric over an unweighted graph can be embedded into a line in linear time with distortion at most $2n - 1$
- Construct a spanning tree, double all the edges, perform an Eulerian traversal starting from some arbitrary vertex, record each vertex the first time it is visited: gives an ordering for the embedding of the vertices.
- Position of the images on the line: first point at 0. If v succeeds to u in the ordering, then place the image of v at distance $d(u, v)$ from the image of u on the line.
- The map is non-contracting; the minimum distance in the source graph is 1; the spread of the embedding is $2n - 1$, QED!

Analysis is tight: take the ladder graph for G for example.

Algorithm from BDGRRRS

- By exhaustive search, guess the graph vertices t_1 and t_2 whose images are leftmost and rightmost on the line
- Compute the shortest path $(v_1, v_2, \dots v_L)$ from t_1 to t_2 in the graph
- Partition the vertices into sets: V_i contains the vertices closest to v_i. Break ties so that V_i is connected.
- Apply Matousek's construction to each V_i starting from v_i and concatenate the resulting embeddings, leaving distance $|V_i|$ between the embedding of V_i and the embedding of V_{i+1}.

Work of Badoiu, Chuzhoy, Indyk and Sidiropoulos (BCIS)

- **Main result:** Given a **weighted** graph and shortest path distances, find an embedding with distortion $O(\Delta^{3/4}c^{11/4})$, where c is the optimal distortion and Δ is the ratio max to min distance.
- Extension of BGDRRRS

Algorithm from BCIS

- Remove all long edges, of weight $\geq L$. Contract the connected components G_i and apply Matousek's lemma to order the (yet to be defined) embeddings of the G_i's, leaving space between the embedding of G_i and of G_j equal to the maximum distance in $G_i \times G_j$.
- To embed G_i: By exhaustive search, guess the graph vertices t_1 and t_2 whose images are leftmost and rightmost on the line.
- Remove all edges of weight $\geq cL$ from G_i. Compute the shortest path $(v_1, v_2, \ldots v_L)$ from t_1 to t_2 in the resulting graph.
- Partition the vertices into sets: V_i contains the vertices closest to v_i.
- Define superclusters W_j consisting of $c^4 L$ consecutive V_is, from a random starting point
- Apply Matousek's construction to each W_j starting from v_i and concatenate the resulting embeddings, leaving distance between the embedding of W_j and the embedding of W_{j+1}.

Hardness of approximation from BCIS

- There is a constant $\beta > 0$ such that embedding weighted trees into the line is n^{β} hard to approximate.
- Reduction from 3SAT(5), the restriction of 3SAT where each variable participates in exactly 5 clauses.
- Uses caterpillar graph gadgets

Work of K., Rabani and Sinclair (KRS)

- **Main result:** Given two line metrics, there is an algorithm to decide whether there exists a bijection with distortion at most $3 + 2\sqrt{2} = 5.82\ldots$
- This is an exact result, not an approximation
- Dynamic program with running time $O(n^{O(1)})$.

Algorithm of KRS

- Given the bijection σ, the expansion is reached by a consecutive pair (u, v) and its image (u', v'), so by exhaustive search we can guess the optimal expansion and inverse expansion, and scale so that they are equal: $\alpha = \beta$.
- Uses the notion of **forbidden pattern**: σ, a permutation of size n, contains the pattern π of size k if there exists a subset $\{i_1, i_2, \ldots, i_k\}$ such that $\sigma(i_j) < \sigma(i_\ell)$ iff $\pi(j) < \pi(\ell)$.

KRS: basic observation

If σ contains the pattern 2413, then for any $U = \{u_1, u_2, u_3, u_4\}$ and $V = \{v_1, v_2, v_3, v_4\}$, the distortion of map σ from U to V is at least $3 + 2\sqrt{2}$.

$$
\begin{aligned}
v_4 - v_2 &\leq c(u_2 - u_1) \\
v_4 - v_1 &\leq c(u_3 - u_2) \\
v_3 - v_1 &\leq c(u_4 - u_3) \\
u_3 - u_1 &\leq (v_2 - v_1) \\
u_4 - u_1 &\leq (v_3 - v_2) \\
u_4 - u_2 &\leq (v_4 - v_3)
\end{aligned}
$$

Implies c is root of a degree two equation which solves to $c \geq 3 + 2\sqrt{2}$.

Corollary: If the distortion is less than $3 + 2\sqrt{2}$, then σ does not contain the pattern 2413.

KRS: second basic observation

- If a permutation σ avoids pattern 2413, then there exists an $i < n$ such that the image under σ of $\{1, 2 \ldots, i\}$ is either $\{1, 2, \ldots, i\}$ or $\{n - i + 1, \ldots, n - 1, n\}$.
- This suggests a divide-and-conquer approach
- Iterating, there is a hierarchical decomposition into intervals, such that σ maps subintervals to subintervals
- The dynamic program to decide whether there exists a map of distortion less than c has one table entry for each pair of intervals, one in the source set and one in the target set of points
- To combine solutions, one must check that all consecutive pairs of points have expansion or inverse expansion at most $\sqrt{c}$. This can be done if the dynamic program also memorizes the images and inverse images of the smallest and largest point of the subintervals.

Recent developments (with Biehler): first observation

- Given a permutation σ, hoe can we compute the minimum distortion achievable over all choices of $\{u_1, \ldots, u_n\}$ and images under σ $\{v_1, v_2, \ldots, v_n\}$?
- Answer: it is the spectral radius of a certain matrix A.
- Let $M(\sigma)$ denote the $n-1$ by $n-1$ 0-1 matrix defined by writing $m_{ij} = 1$ iff the interval between $\sigma(i)$ and $\sigma(i+1)$ contains $[j, j+1]$. Then $A = M(\sigma)M(\sigma^{-1})$.

Second observation

- **Simple permutation:** a permutation which maps no proper non-singleton onto an interval
- **Albert-Atkinson (to appear):** A simple permutation of size n must contain a simple permutation of size $n-1$ or $n-2$, for $n \geq 3$.
- **Corollary:** If σ avoids all simple permutations of size k or $k+1$, then σ avoids all permutations of size $\geq k$.
- **Non-simple permutations:** can be decomposed into blocks, such that the map defining which block maps to each block is a simple permutation. This suggests a divide-and-conquer approach.
- **Corollary:** If σ avoids all permutations of size $\geq k$, then there is a block decomposition with at most $k-1$ blocks. This suggests a divide-and-conquer approach.

Putting it together

- **Preprocessing:** Use the spectral radius theorem to write a computer program that finds the minimum distortion of all simple permutations of size ≤ 10
- **Result:** Simple permutations of size 9 or 10 have minimum distortion $5 + 2\sqrt{6} = 9.90\ldots$.
- Deduce a dynamic programming algorithm to decide whether there find the optimal bijection, when there exists one with distortion at most $5 + 2\sqrt{6}$

How far can we push that?

The “lattice" permutation of size n is simple and has distortion less than $7 + 4\sqrt{3} = 13.9\ldots$ so another idea is need to go beyond that.

Conclusion

Could we have a parametrized algorithm whose running time would be polynomial for fixed c?

How about more efficient algorithms which achieve some approximation factor?

Does this really help recognize manuscript characters?

How about the 2-dimensional case?