Computing Largest Common Point Sets under Approximate Congruence

Samarjit Chakraborty

Joint work with Christoph Ambühl and Bernd Gärtner
Eidgenössische Technische Hochschule Zürich

Eidgenössische
Technische Hochschule
Zürich

Largest Common Point Sets (LCP)

Given: Two point sets A, $B \subset \mathbb{R}^d$, a transformation group $\mathcal{G}$, and a real number $\varepsilon \geq 0$

Goal: Compute the maximum cardinality subset $A' \subseteq A$, for which there exists a $g \in \mathcal{G}$, such that the distance $d(g(A'), B') \leq \varepsilon$

Here $B' \subseteq B$ and the "distance" between two point sets is measured using a metric like the Hausdorff distance, bottleneck matching metric, etc.

Versions of Point Set Pattern Matching Problems

- Dimensionality d of the point sets.
- The allowed group $\mathcal{G}$ of transformations, like only translations, rotations, isometries, affine transformations, etc.
- Whether entire point sets, or only parts of them are considered, examples - congruent copy detection (CCD),LCP, etc.
- The notion of similarity between two point sets, measured using the Hausdorff distance, bottleneck matching metric, etc.
  - Notion of match between two points (ε is given; minimize ε)
  - L_∞-, L_2-norms, ...
- Exact or approximation algorithms.

This Talk

- d-dimensional point sets
- Translations, isometries, affine transformations
- Parts of the two point sets are allowed to match - LCP, with a given ε
- Bottleneck matching metric under L_2- and L_∞-norms
- Exact algorithms

Initiated by [Alt, Mehlhorn, Wagener, Welzl '88] ($d = 2$; isometries; equal cardinality point sets; ε given; bottleneck matching under L_2-norm). Running time $O(n^8)$

Exact Matching Metric		
Properties	Complexity	Paper
3-d point sets	$O(n \log n)$	Atkinson '87
d-dimensions	$O(n^{d-2} \log n)$	Alt *et al.* '88
d-dimensions, Monte-Carlo	$O(n^{(d-1)/2} \log n)$	Akutsu '98
d-dimensions, deterministic	$O(n^{d/2+O(1)})$	Matoušek
d-dimensions	$O(n^{\lceil d/3 \rceil} \log n)$	Braß, Knauer '00

Exponential dependence on d ! Asymptotic lower bound is $\Omega(n \log n)$. Therefore satisfactory solutions are yet to be found for higher dimensions.

Previous Work (Contd.)

Minimum Hausdorff Distance		
Properties	Complexity	Paper
planar, translation, L_1, L_2, L_∞ 3-d, translation, L_2	$O(n^3 \log n)$ $O(n^{5+\varepsilon})$	Huttenlocher *et al.* '93
planar, L_2, isometry	$O(n^5 \log^2 n)$	Chew *et al.* '97
d-dimensions, L_∞, translations L_2	$O(n^{5d/4} \log^2 n)$ $O(n^{\lceil 3d/2 \rceil + 1 + \delta})$	Chew, Dor Efrat, Kedem '99

d-dimensional point sets and isometry?

Bottleneck Matching (under ε-Congruence)		
Properties	Complexity	Paper
planar, isometry, L_2	$O(n^8)$	Alt, Mehlhorn, Wagener, Welzl '88

Extensions to higher dimensions?

- Existing algorithms either work only in low-dimensions, or for restricted sets of metrics and transformations.
- Exact matching is ill-posed and many applications require a one-to-one matching between points (Hausdorff metric is many-to-one).
- The LCP problem is a generalization of matching between two equal cardinality point sets.

Previous Work (Contd.)

Approximation Algorithms		
Properties	Complexity	Paper
planar, isometry, bottleneck matching	$O(n^{2.5})$	Heffernan, Schirra '92
planar, isometry, Hausdorff	$\tilde{O}(n^2)$	Indyk *et al.* '99
d-dimensions, translations isometry, Hausdorff	$O(n \log n)$ $O(n^d \log n)$	Cardoze, Schulman '99
planar, isometry *generalized bottleneck*	$\tilde{O}(n)$	Indyk, Venkatasubramanian '00

The algebraic convolution based paradigm due to [Indyk *et al.* '99] and [Cardoze, Schulman '99] does not extend to bottleneck matching.

What's the problem?

- An exact solution is more difficult than an approximate one.
- Solutions from the low dimensional cases do not generalize.
- Bottleneck matching is a "global property" of the point sets.
- Believed that it would involve computing arrangements of curves in high-dimensional space.
- Would very likely be of high time complexity.

No algorithms, and no complexity results are known!

Given: A transformation L and a translation vector v

Goal: Compute the "distance" between point sets $L(A) + v$ and B

- Construct the bipartite graph $G = (A \cup B, E)$
  - for $a_i \in A$ and $b_j \in B$, $(a_i, b_j) \in E$ if $d(L(a_i) + v, b_j) < \varepsilon$
- Compute the maximum matching in G

Computing the LCP amounts to finding the optimal L and v under which the graph G has the maximum matching.

LCP under translation (L is fixed)

Problem: Compute the LCP between $L(A)$ and B

Solution: The "translation space"

$$\mathcal{T}_{ij}(L) := \{v \in \mathbb{R}^d | \; d(L(a_i) + v, \, b_j) < \varepsilon\}$$

- Compute the arrangement $\mathcal{A}(L)$ formed by the overlay of all possible $\mathcal{T}_{ij}(L)$s
- v_{opt} lies in some cell of $\mathcal{A}(L)$
- Compute the maximum matching in the graph G for each cell

Computing v_{opt} reduces to examining one vector in each cell of $\mathcal{A}(L)$, thereby resulting in a discrete version of the problem.

LCP under general transformations

Basic idea: As L varies, the LCP between $L(A)$ and B under translation remains invariant as long as the arrangement $\mathcal{A}(L)$ does not undergo any combinatorial change.

This means: The different combinatorial structures of $\mathcal{A}(L)$ partition the space $\mathcal{L}$ (of all possible transformations L) into a number of cells. Each cell is the maximal set of transformations generating combinatorially equivalent arrangements $\mathcal{A}(L)$.

- Here L is a rotation parametrized by a single angle $\theta \in [0, 2\pi)$.
- $\mathcal{T}_{ij}(\theta)$ is a circular disk of radius ε, therefore $\mathcal{A}(\theta)$ is an arrangement of $O(mn)$ such disks.
- Combinatorial structure of $\mathcal{A}(\theta)$ changes when [Helly's theorem]
  1. Two disks touch at a point - linear equation in $\sin\theta$ and $\cos\theta$
     $\Rightarrow$ quadratic equation in either $\sin\theta$ or $\cos\theta$
  2. Three disks meeting at a point - cubic equation in $\sin\theta$, $\cos\theta$
     $\Rightarrow$ algebraic equation of degree six in either $\sin\theta$ or $\cos\theta$
- [1] gives rise to $O(m^2n^2)$ intervals and [2] gives rise to $O(m^2n^3)$ intervals.
- Overall complexity of $O(n^{7.5})$ is an improvement compared to $O(n^8)$.

Affine transformations and the L_∞-norm

Consider transformation L to be a $d \times d$ matrix with variable coefficients $x_{11}, \ldots, x_{dd}$

- As L varies, the combinatorial structure of $\mathcal{A}(L)$ changes only at points where for some i, j, k, l, $\mathcal{T}_{ij}(L)$ and $\mathcal{T}_{kl}(L)$ share a coordinate in some *direction* t.

- This gives rise to linear equality constraints in $x_{t1}, \ldots, x_{td}$.

- Therefore all combinatorial changes in $\mathcal{A}(L)$ can be characterized by hyperplanes in d^2 dimensional space.

Affine Transformations and the L_∞-norm (Contd.)

Algorithm

1. The $O(m^2n^2)$ hyperplanes in dimension d^2 define the arrangement $\mathcal{H}$. The number of cells in the worst case is $(mn)^{2d^2}$ and can be traversed in time proportional to this using *reverse search* [Avis and Fukuda '96].

2. From each cell of $\mathcal{H}$ choose a L, and traverse the arrangement $\mathcal{A}(L)$ (time $(mn)^{O(d)}$).

3. Perform graph matching in each cell of $\mathcal{A}(L)$ (time $(mn)^{O(1)}$).

4. Overall time complexity is $(mn)^{2d^2+O(d)}$ (therefore polynomial for fixed d) and space is $O(m^2n^2)$ (independent of d).

Affine Transformations under the L_2-norm

- Here $\mathcal{A}(L)$ is an arrangement of Euclidean ε-balls.

- A combinatorial change in $\mathcal{A}(L)$ occurs iff some circumscribed ball of k $(2 \le k \le d+1)$ centers attains radius ε [Helly's theorem].

- Each such condition can be written as $f(x) = 0$, where $f(x)$ is a constant-degree (in $O(d)$) polynomial in the entries of the matrix L, i.e. in $x_{11}, \ldots, x_{dd}$.

- These give rise to a collection of $(mn)^{O(d)}$ d^2-variate polynomials with rational coefficients and maximum degree $O(d)$.

- "Traverse" the cells of this arrangement of polynomials and proceed as before.

Affine Transformations under the L_2-norm (Contd.)

Combinatorial changes in planar arrangement of circles

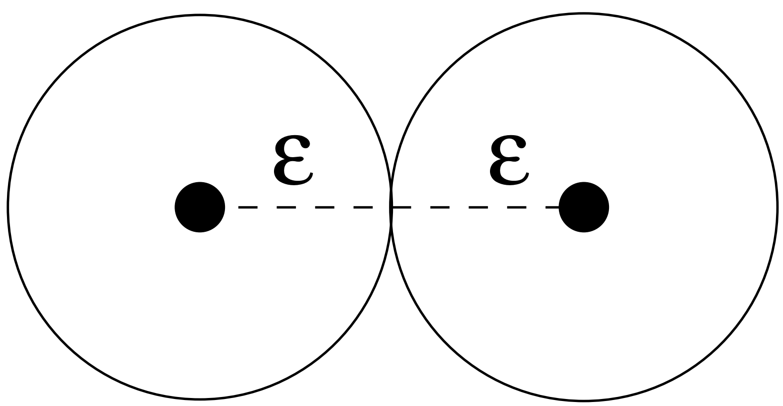

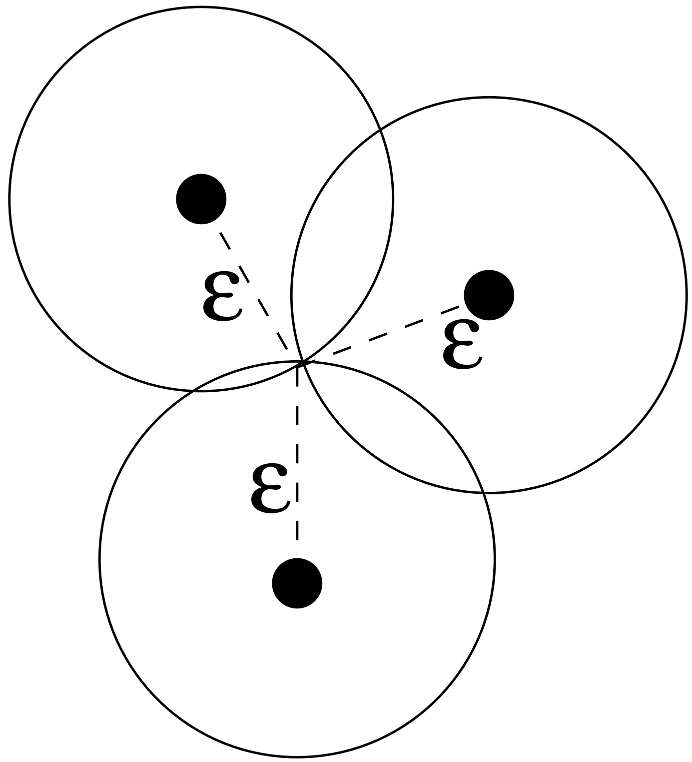

Affine Transformations under the L_2-norm (Contd.)

How to "traverse" an arrangement of polynomials?

Theorem (**Basu, Pollack, Roy '95**) Let $\mathcal{P} = \{f_1, \ldots, f_r\}$ be a set of p-variate polynomials with rational coefficients and maximum algebraic degree s ($p < r$). Two points $q, q' \in \mathbb{R}^p$ are *equivalent* if $\text{sign}(f_\ell(q)) = \text{sign}(f_\ell(q'))$ for $\ell = 1, \ldots, r$. The vector $(\text{sign}(f_1(q)), \ldots,$ $\text{sign}(f_\ell(q))$ is a *sign condition* of $\mathcal{P}$. In $O(r(r/p)^p s^{O(p)})$ arithmetic operations, one can compute all sign conditions determined by $\mathcal{P}$.

Affine Transformations under the L_2-norm (Contd.)

In our case all the sign conditions can be obtained in $(mn)^{O(d^3)}$ time.

- We are interested in the "full-dimensional" sign conditions (containing no zero sign).
- The combinatorial structure of $\mathcal{A}(L)$ can be determined from the sign condition.
- Overall runtime including graph matchings is bounded by $(mn)^{O(d^3)}$.

Incorporating isometries L defines an isometry iff $L^{-1} = L^T$ ($\det(L) = 1$ for only rotations). These give rise to $O(d^2)$ additional polynomial equations. Runtime is still $(mn)^{O(d^3)}$.

- A rotation in 3-space can be parametrized by ϕ_1, ϕ_2, ϕ_3. Entries of L are polynomials in $\sin\phi_i, \cos\phi_i,\ i = 1, 2, 3$.
- The combinatorial change conditions give rise to $O(m^4 n^4)$ constant degree polynomial equations $\mathcal{P}$ in $\sin\phi_i, \cos\phi_i$.
- Let $\mathcal{V}$ be the zero set of the polynomial $\sum_{i=1,2,3}(\sin\phi_i^2 + \cos\phi_i^2 - 1)^2$.
- In $O(m^{16}n^{16})$ time we can compute the $O(m^{12}n^{12})$ points in each non-empty semi-algebraically connected component of $\mathcal{P}$ over $\mathcal{V}$, along with the signs of all the polynomials of $\mathcal{P}$ at each of these points [Basu, Pollack, Roy '97].
- Given a fixed sign condition, the corresponding arrangement contains $O(m^3 n^3)$ cells.
- Overall time complexity is $O(m^{16}n^{16}\sqrt{m+n})$.

1. For unbounded dimensions, the LCP problem under approximate congruence is NP-hard (settles the issue raised by Akutsu '98)
   - Even under only translations, and also under isometries, linear transformations, ...
   - Reduction from SAT

2. *Subset matching* is NP-hard for unbounded dimensions [Akutsu '98]

3. Computing the LCP of an unbounded number of point sets even in one-dimension is hard [Akutsu and Halldórsson '00]

Conclusions and Open Problems

- Exact algorithm for LCP under approximate congruence in d dimensions

- It seems that the time complexity of $(mn)^{O(d^3)}$ is an artifact of the proof and should have been $(mn)^{O(d^2)}$

- The 3-d case is well solved from the point of view of approximation algorithms

- Approximation algorithms for high dimensions
  - Present algorithms have an exponential dependence on d
  - Dimensionality reduction techniques?
  - Recent progress in NN-search algorithms for high dimensional spaces...

Questions?

Questions?

Questions?

Questions?

Questions?

Questions?

Questions?

Questions?

Questions?