ISSN 2186-7437

NII Shonan Meeting Report

No. 2016-9

Theory and Applications of Geometric Optimization

Siu-Wing Cheng
Otfried Cheong
Yoshio Okamoto

May 30–June 2, 2016

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Theory and Applications of Geometric Optimization

Organizers:
Siu-Wing Cheng (Hong Kong University of Science and Technology)
Otfried Cheong (Korea Advanced Institute of Science and Technology)
Yoshio Okamoto (The University of Electro-Communications)

May 30–June 2, 2016

There are many algorithmic, geometric, combinatorial, and implementation challenges that arise in geometric optimization. Such challenges have prompted relevant research and progress in algorithm design and analysis. Geometric optimization has wide applications in and connections to various areas, including computer graphics, computer-aided design and manufacturing, robotics, computer vision, spatial databases, geographical information systems, machine learning, and scientific computing. A lot of the theoretical results in geometric optimization fall within the area of computational geometry, which is a vibrant and mature field of research, with a flagship annual international conference and several dedicated international journals. Some of the results also appear in theoretical computer science conferences and journals. This meeting will focus on current interests and future trends in geometric optimization.

For example, data analysis is a popular research topic. The support vector machine paradigm has successfully linked many data analytics problems to rigorously defined convex optimization problems, which have a strong geometric flavor. Indeed, finding a sparse approximate solution to such a convex optimization problem is closely related to the concept of core-sets in computational geometry. It has been discovered that core-sets are strongly related to the well-known Frank-Wolfe greedy optimization method. The problem of analyzing massive data size has prompted researchers to study the computation of a good approximation by examining only a small subset of the input. This is a particularly popular research theme in the data streaming environment. For example, it is extremely useful to obtain a faithful, concise, and dynamic summary of a data stream that is provably good under some well-defined criterion.

There has been recent algorithmic progress on the shape matching problem under rigid and affine transformations. The problem calls for placing two shapes using the allowed transformations in order to maximize some similarity measure (for instance overlap or symmetric difference). Recently, fast approximation algorithms have been obtained for matching polygonal shapes under rigid motions, matching convex sets in arbitrary dimensions under scaling and translations, and finding the largest common point set under rigid motions in 2D. Good progress has also been obtained on computing the Frechet distance of two polygonal curves. There is still a lot to be done for the shape matching problem; for example, matching polyhedral shapes in 3D under rigid motions, finding the largest common point set under rigid motions in 3D, etc. There are also many GPS trajectory problems that are closely related to shape matching.

Finding a shortest path is a classical geometric optimization problem that finds applications in computer graphics, motion planning, geographical information systems, and seismic simulations. The traditional objective is to minimize the Euclidean path length, but there has been research work on modeling the varying difficulties in traversing different regions. Recently, there has been progress in the weighted region model, in combining path length with height constraints, and in handling situations in which the speed of travel is direction-sensitive. There are still many open problems; for instance whether there is an FPTAS for the weighted region problem in 3D; or whether one can handle more general cost functions for navigating a terrain. Closely related is the evacuation problem, which finds applications in planning evacuation routes for crowds of people when an emergency arises. There has been progress on the problem when the underlying network is a path. More research is needed to handle more general graph topologies.

There are many other important geometric optimization problems, in addition to the ones mentioned above. To advance the state of the art, an extensive collaboration among researchers is highly desirable. This meeting serves to trigger such collaboration on advanced research in this area.

Meeting Schedule

- May 29 (Sunday)
  - 15:00: Check in
  - 19:00 21:00 : Welcome reception
- May 30 (Monday)
  - 07:30 09:00 : breakfast
  - 09:00 09:15 : NII and Shonan Introduction
  - 09:15 10:15 : (Survey talk) Pankaj Agarwal: Algorithms for Geometric Similarity
  - 10:15 11:00 : coffee break
  - 11:00 11:30 : Ji-won Park: Obstructing Visibilities with One Obstacle.
  - 11:30 12:00 : Peyman Afshani: Approximating Simplicial Depth
  - 12:00 13:30 : lunch
  - 13:30 14:30 : (Survey talk) Anne Driemel: Two decades of algorithms for the Frechet distance
  - 14:30 15:00 : Pat Morin: Turán-Type Theorems for Triangles in Convex Point Sets
  - 15:00 15:30 : Joachim Gudmundsson: Sparse geometric networks supporting online routing
  - 15:30 16:15 : coffee break
  - 16:15 18:00 : open problem session and discussion
  - 18:00 19:30 : dinner
- May 31 (Tuesday)
  - 07:30 09:00 : breakfast
  - 09:15 10:15 : (Survey talk) Suresh Venkatasubramanian: The Shape of Learning
  - 10:15 11:00 : coffee break
  - 11:00 11:30 : Jack Snoeyink: Guarantees for Neutron Tracking: Theory in Practice
  - 11:30 12:00 : Wolfgang Mulzer: Routing in Unit Disk Graphs
  - 12:00 13:30 : lunch
  - 13:30 14:30 : (Survey talk) Helmut Alt: Computational Aspects of Packing Problems
  - 14:30 15:30 : break-out session and discussion
  - 15:30 16:15 : coffee break
  - 16:15 18:00 : break-out session and discussion
  - 18:00 19:30 : dinner

- June 1 (Wednesday)
  - 07:30 09:00 : breakfast
  - 09:15 10:15 : (Survey talk) Matias Korman: Computational geometry algorithms in memory constrained environments
  - 10:15 11:00 : coffee break
  - 11:00 12:00 : break-out session and discussion
  - 12:00 13:30 : lunch
  - Excursion starts at 13:30.
- June 2 (Thursday)
  - 07:30 09:00 : breakfast
  - 09:15 10:15 : break-out session and discussion
  - 10:15 11:00 : coffee break
  - 11:00 12:00 : wrap-up session
  - 12:00 13:30 : lunch

Survey Talks

Algorithms for Geometric Similarity

Pankaj Agarwal, Duke University

A basic problem in classifying, or searching for similar objects, in a large set of geometric objects is computing similarity between two objects. There has been extensive work on computing geometric similarity between two objects. This talk discusses some old and some new geometric-similarity algorithms, with an emphasis on transportation distance, Frechet distance, dynamic time warping, and Gromov-Hausdorff distance. It will also touch upon a few open problems in this area.

Two decades of algorithms for the Frechet distance

Anne Driemel, Eindhoven University of Technology

Time series data is one of the most common forms of big data. More generally, we can think of sequence data such as DNA-sequences, strings, trajectories of moving objects, geometric curves and speech recordings. The edit distance stands as an unchallenged distance measure for strings. Dynamic time warping is a variation of the edit distance that was developed for speech recognition and is heavily used in the data mining community for various types of time series data, alongside with LCS (longest common subsequence). The Frechet distance is the geometric counterpart which was conceived independently as a mathematical metric for curves.

In 1995, Alt and Godau described a quadratic-time algorithm to decide if two polygonal curves are similar under the Frechet distance. For almost two decades, faster algorithms seemed out of reach and it was commonly conjectured that this algorithm is optimal. In 2014, Bringmann gave complexity-theoretical evidence for this conjecture by proving that no $O(n^{2-\epsilon})$-algorithm can exists (for any $\epsilon > 0$), unless the strong exponential time hypothesis is false.

The talk will put these results into context with related distance measures such as dynamic time warping, LCS and edit distance, which share a similar story. We will review the body of work that has been developed to deal with the quadratic hardness, making near-linear time algorithms possible after all and we will go beyond single distance computation towards important topics such as data structures and clustering under the Frechet distance.

The Shape of Learning

Suresh Venkatasubramanian, Univeristy of Utah

Often lost in the fuss over big data and machine learning is the basic fact that geometric representations are at the heart of effective learning. Learning representations often resolves to finding the right geometric representation of the data (whether it be in the form of a kernel, a manifold or even the output of an auto encoder). Optimization, so much a part of learning, has natural geometric interpretations that lead to clean and easy analyses of many problems.

In this talk, I review some of the many connections between geometry and machine learning, and point out a number of places where insights from computational geometry might be brought to bear on important problems in ML. Warning: this talk will likely go well beyond low dimensions!

Computational Aspects of Packing Problems

Helmut Alt, Freie Universität Berlin

Packing problems are concerned with positioning geometric objects so that they do not overlap and require an amount of space as small as possible. They have been investigated within mathematics for centuries starting with the famous Kepler conjecture. There are many surprising properties and open problems in connection with packing. The lecture will give a short survey about these classical issues but then concentrate on algorithms for packing. Since already the most simple variants are NP-hard, it makes sense to look for efficient approximation algorithms. We will present constant factor approximations for packing rectangles and convex polygons into containers which are minimum area rectangles or convex sets. Algorithms for analogous problems concerning three-dimensional objects will be presented, as well.

This is joint research with Mark de Berg, Christian Knauer, Léonard von Niederhäusern, and Nadja Scharf.

Computational geometry algorithms in memory constrained environments

Matias Korman, Tohoku University

An s-workspace algorithm has read-only access to the elements of the input and can only use $O(s)$ words of working space (for some small value $s < n$). In this talk we will give a general survey on the computational geometry algorithms that have been designed for such workspaces, with a special emphasis towards the possible lines of research.

Talks

Obstructing Visibilities with One Obstacle.

Ji-won Park, KAIST

An obstacle representation of a graph G is a drawing of G in the plane with polygons called obstacles; two points are adjacent iff the straight line segment connecting them does not intersect any obstacles. Obstacle number of a graph is the smallest number of obstacles which allows an obstacle representation of the graph.

Even a class of graphs of obstacle number 1 is not known completely. There is a nice characterization for graphs which have a representation with 1 convex obstacle: non-double covering circular arc graphs. Also it is known that every outerplanar graph has a representation with 1 outside obstacle. And as far as I know, they are all results about graphs of obstacle number 1.

In this talk, some recent results are presented:

1. Every graph of circumference at most 6 has an outside obstacle representation.
2. A smallest graph of obstacle number 2. It has 8 vertices and it is tight. The smallest graph of obstacle number 2 known so far had 10 vertices.
3. A class of graphs with an outside obstacle and a a class of graphs without an outside obstacle are different. It was one of main questions on the obstacle number of graphs.

Approximating Simplicial Depth

Peyman Afshani, Aarhus University

The simplicial depth of a point q with respect to a given set P of n points is the number of simplicies formed by the points of P that contain q. The exact computation of simplicial depth is costly and it is conjectured to require at least $\Omega(n^{d-1})$ time. However, there is very little known about methods to approximate the depth. In this talk, we look at some results and open questions in this direction.

This is joint work with Donald Sheehy and Yannik Stein

Turán-Type Theorems for Triangles in Convex Point Sets

Pat Morin, Carleton University

Originally motivated by a problem of Erdös on the maximum number of maximum-area triangles determined by an nn-point set, we will discuss some new and old results on the following family of problems: Given a set of n points in convex position, what is the maximum number triangles one can create having these points as vertices while avoiding certain forbidden configurations. As forbidden configurations we consider all 8 ways in which a pair of triangles in such a point set can interact by sharing and/or interleaving vertices.

Sparse geometric networks supporting online routing

Joachim Gudmundsson, University of Sydney

Online routing in a planar embedded graph is central to a number of fields and has been studied extensively in the literature. For most graphs no competitive online routing algorithm exists. A notable exception is the Delaunay triangulation for which Bose and Morin showed that there exists an online routing algorithm that is c-competitive. However, a Delaunay triangulation might be expensive to build; it can have linear degree and a total weight that is a linear factor greater than the weight of a minimum spanning tree.

We approach the problem from a different direction. We tried to prove that given a set V of n points in the Euclidean plane and two positive constants r and $t < \pi/8$, one can construct a plane geometric graph of V in $O(n \log n)$ time that has (i) weight at most $(2r + 1)(1 + 1/\cos(2t)) \cdot \mathrm{wt}(\mathrm{MST}(V))$, (ii) degree at most $15\pi/t$ and (iii) admits a local routing strategy that is $O(1)$-competitive. To the best of our knowledge this is the first time such a construction has been attempted. Unfortunately we have been unable to prove the third property. We can prove and that for every edge (u, v) in a Delaunay triangulation of V, u and v has a $(1 + 1/\cos(2t))$-competitive face-path in G.

Joint work with Christos Levcopoulos and Bengt J. Nilsson.

Guarantees for Neutron Tracking: Theory in Practice

Jack Snoeyink, UNC Chapel Hill (currently at the US National Science Foundation)

Geometry often serves as a boundary object that enables communication within interdisciplinary teams: a CAD model of nuclear reactor is viewed differently by a nuclear physicist, a machinist, and a safety engineer, since each brings his or her specialized knowledge, but the common geometry lets them share ideas for design optimization. Sometimes the computation geometry supports this practical endeavor not by phrasing an optimization question or speeding up an algorithm, but by providing ideas that can make the tools for manipulating this boundary object robust and predictable.

I illustrate this with a problem suggested by David Griesheimer of Bettis Labs: neutron tracking along segments in hierarchical Constructive Solid Geometry (CSG) with quadratic primitives represented in floating point. By incorporating geometric rounding and degree-driven analysis into the algorithm design we phrase the problem in a way that can give both topological and geometric guarantees. (With former students David Millman and Michael Deakin.)

Routing in Unit Disk Graphs

Wolfgang Mulzer, Freie Universität Berlin

Let S be a set of n sites in the plane. The unit disk graph $UD(S)$ on S has vertex set S and an edge between two distinct sites $s, t \in S$ if and only if s and t have Euclidean distance $|st| \leq 1$. A routing scheme R for $UD(S)$ assigns to each site s in S a label $l(s)$ and a routing table $\rho(s)$. For any two sites $s, t \in S$, the scheme R must be able to route a packet from s to t in the following way:

given a current site r (initially, $r = s$), a header h (initially empty), and the target label $l(t)$, the scheme R may consult the current routing table $\rho(r)$ to compute a new site r and a new header h, where r is a neighbor of r. The packet is then routed to r, and the process is repeated until the packet reaches t. The resulting sequence of sites is called the routing path. The stretch of R is the maximum ratio of the (Euclidean) length of the routing path produced by R and the shortest path in $UD(S)$, over all pairs of distinct sites in S.

For any given $\epsilon > 0$, we show how to construct a routing scheme for $UD(S)$ with stretch $1 + \epsilon$ using labels of $O(\log n)$ bits and routing tables of $O(\epsilon^{-5} \log^2 n \log^2 D)$ bits, where D is the (Euclidean) diameter of $UD(S)$. The header size is $O(\log n \log D)$ bits.

Based on joint work with Haim Kaplan, Liam Roditty, and Paul Seiferth.

Open Problems

This section contains the list of open problems posed at the workshop by several participants.

Jack Snoeyink. You are given n points in $\mathbb{R}^3$ in general position. You are also given a planar triangulation with n vertices. If you bijectively map the triangulation vertices to the n points uniformly at random, *what is the probability that there is no self-intersection in the resulting triangulated surface in* $\mathbb{R}^3$*?* A lower bounded of this probability is sought. The question is already interesting for special triangulations, like an n-wheel for small n, which is the n-vertex graph with one vertex of degree $(n-1)$ and all other vertices of degree 3.

Pat Morin. You are given a graph G embedded (possibly with crossing edges) in the plane. For each vertex v of G, we associate a disk D_v such that D_v centered at v and is just large enough to cover the neighbors of v in G. For every point $q \in \mathbb{R}^2$, define the *interference* at q, denoted by $\mathrm{intf}(q)$, to be the number of these disks that contains q. Define the interference of G, denoted by $\mathrm{intf}(G)$, to be $\max_{q\in\mathbb{R}^2} \mathrm{intf}(q)$.

Let P be a set of n points uniformly distributed in a unit square. Let $\mathrm{MST}(P)$ denote the minimum spanning tree of P. It is known that $\mathrm{E}[\mathrm{intf}(\mathrm{MST}(P))] = O(\sqrt{\log n})$. There exists a graph H on n points such that $\mathrm{E}[\mathrm{intf}(H)] = O(\log^{1/3} n)$. For every graph G on n points, $\mathrm{E}[\mathrm{intf}(G)] = \Omega(\log^{1/4} n)$. *Can we narrow or close the gap between the upper and lower bounds?*

Peyman Afshani. Consider a set of n points in $\mathbb{R}^3$. *Is there a data structure for nearest neighbor queries with $O(\log^{O(1)} n)$ query time using $O(n^{2-\alpha})$ space for some $\alpha > 0$?*

Jack Snoeyink. Peyman started his question with the motivation of finding useful structure in 3d Delaunay complexes that may be large. His question became algorithmic, but one might also be interested in this combinatorial question of Herbert Edelsbrunner. Let M be a simplicial complex (tetrahedralization) whose union is the hull of n points in $\mathbb{R}^3$. Define a *2-tree* T of M as a sequence of triangles $t_1, \ldots, t_k$ such that, for all $1 < i \leq k$, triangle t_i of M that joins an edge from a triangle t_j for $1 \leq j < i$ to a new vertex v_i not in any triangle $t_1, \ldots, t_{i-1}$. A *spanning 2-tree* contains all vertices of M, and thus has $k = n-2$ triangles. *Does the 3D Delaunay triangulation of n points in general position always have a spanning 2-tree? Does any triangulation on points in 3D have a spanning 2-tree?*

Sang Won Bae and Ji-won Park. You are given n disjoint line segments in $\mathbb{R}^2$. It is possible to place a set P of points on the line segments such that for every $p_i \in P$, the Voronoi cell of p_i does not intersect any line segment other than the one containing p_i. We say that P separates the set of line segments. For a set of vertical and horizontal line segments, it is known that $\Theta(n^2)$ points are necessary and sufficient. *How large is P in general? Is there an algorithm to find P with the minimum cardinality?*

Note from Jack Snoeyink: Results of Chris Bishop on non-obtuse triangulation appear to imply an $O(n^{2.5})$ upper bound on $|P|$. The lower bound is $\Omega(n^2)$.

Jeff Phillips. You are given set $\mathcal{P} = \{P_1, P_2, \ldots, P_n\}$ of n uncertain points on the real line $\mathbb{R}$. Each uncertain point P_i has k possible locations $\{p_{i,1}, p_{i,2}, \ldots, p_{i,k}\}$. A set $Q = \{q_1, \ldots, q_n\}$ of n points is an instantiation of $\mathcal{P}$, denoted $Q \Subset \mathcal{P}$, if for all $i \in [1, n]$, $q_i = p_{i,j}$ for some $j \in [1, k]$. We use $\text{med}(Q)$ to denote the median of Q. Define $\text{cost}(p, Q) = \frac{1}{n} \sum_{i=1}^{n} \|p - q_i\|$. *For a given p, calculate*

$$\text{cost}(p) = \min_{\substack{Q \Subset P \\ p = \text{med}(Q)}} \text{cost}(p, Q),$$

given that $p = p_{i,j}$ for some $i \in [1, n]$ and $j \in [1, k]$, and there exists some Q' such that $p = \text{med}(Q')$. Can this be calculated in $\text{poly}(n, k)$ time?

Note: Jeff sketched a solution to the above open problem that may run in $O(n\text{polylog}n)$ time.

List of Participants

Helmut Alt, Freie Universität Berlin
Peyman Afshani, Aarhus University
Pankaj Agarwal, Duke University
Hee-Kap Ahn, POSTECH
Boris Aronov, Tandon School of Engineering, NYU
Sang Won Bae, Kyonggi University
Prosenjit Bose, Carleton University
Kevin Buchin, TU Eindhoven
Siu-Wing Cheng, Hong Kong University of Science and Technology
Otfried Cheong, Korea Advanced Institute of Science and Technology
Man Kwun Chiu, National Institute of Informatics, Japan
Jinhee Chun, Tohoku University
Anne Driemel, TU Eindhoven
Vida Dujmović, University of Ottawa
Joachim Gudmundsson, University of Sydney
Christian Knauer, Universität Bayreuth
Matias Korman, Tohoku University
Anil Maheshwari, Carleton University
Pat Morin, Carleton University
Wolfgang Mulzer, Freie Universität Berlin
Yoshio Okamoto, The University of Electro-Communications
Yota Otachi, JAIST
Ji-won Park, KAIST
Jeff Phillips, University of Utah
Marcel Roeloffzen, National Institute of Informatics, Japan
Vera Sacristán, Universitat Politècnica de Catalunya
Michiel Smid, Carleton University
Jack Snoeyink, UNC Chapel Hill and National Science Foundation, USA
Fabian Stehn, Universität Bayreuth
Ryuhei Uehara, JAIST
Andre van Renssen, National Institute of Informatics, Japan
Suresh Venkatasubramanian, University of Utah