Tuesday, June 18

9:00-9:10	Welcome	
9:10-9:40	Best Paper: **Almost Tight Lower Bounds for Hard Cutting Problems in Embedded Graphs** V. Cohen-Addad, É. Colin de Verdière, D. Marx and A. de Mesmay We prove essentially tight lower bounds, conditionally to the Exponential Time Hypothesis, for two fundamental but seemingly very different cutting problems on surface-embedded graphs: the Shortest Cut Graph problem and the Multiway Cut problem. A cut graph of a graph G embedded on a surface S is a subgraph of G whose removal from S leaves a disk. We consider the problem of deciding whether an unweighted graph embedded on a surface of genus g has a cut graph of length at most a given value. We prove a time lower bound for this problem of $n^{\Omega(g/\log g)}$ conditionally to ETH. In other words, the first $n^{O(g)}$-time algorithm by Erickson and Har-Peled [SoCG 2002, Discr. Comput. Geom. 2004] is essentially optimal. We also prove that the problem is W[1]-hard when parameterized by the genus, answering a 17-year old question of these authors. A multiway cut of an undirected graph G with t distinguished vertices, called *terminals*, is a set of edges whose removal disconnects all pairs of terminals. We consider the problem of deciding whether an unweighted graph G has a multiway cut of weight at most a given value. We prove a time lower bound for this problem of $n^{\Omega(\sqrt{gt+g^2}/\log(gt))}$, conditionally to ETH, for any choice of the genus $g \geq 0$ of the graph and the number of terminals $t \geq 4$. In other words, the algorithm by the second author [Algorithmica 2017] (for the more general multicut problem) is essentially optimal; this extends the lower bound by the third author [ICALP 2012] (for the planar case). Reductions to planar problems usually involve a grid-like structure. The main novel idea for our results is to understand what structures instead of grids are needed if we want to exploit optimally a certain value g of the genus.	
9:40-10:50	Fast Forward/Coffee Break	
	Tue-2A: Data Structures I	**Tue-2B: Persistent Homology I**
10:50-11:10	**Dynamic Planar Point Location in External Memory** J. I. Munro and Y.Nekrich In this paper we describe a fully-dynamic data structure for the planar point location problem in the external memory model. Our data structure supports queries in $O(\log_B n(\log\log_B n)^3))$ I/Os and updates in $O(\log_B n(\log\log_B n)^2))$ amortized I/Os, where n is the number of segments in the subdivision and B is the block size. This is the first dynamic data structure with almost-optimal query cost. For comparison all previously known results for this problem require $O(\log_B^2 n)$ I/Os to answer queries. Our result almost matches the best known upper bound in the internal-memory model.	**DTM-based Filtrations** H. Anai, F. Chazal, M. Glisse, Y. Ike, H. Inakoshi, R. Tinarrage and Y. Umeda Despite strong stability properties, the persistent homology of filtrations classically used in Topological Data Analysis, such as, e.g. the Čech or Vietoris-Rips filtrations, are very sensitive to the presence of outliers in the data from which they are computed. In this paper, we introduce and study a new family of filtrations, the DTM-filtrations, built on top of point clouds in the Euclidean space which are more robust to noise and outliers. The approach adopted in this work relies on the notion of distance-to-measure functions and extends some previous work on the approximation of such functions.
11:10-11:30	**A Divide-and-Conquer Algorithm for Two-Point L_1 Shortest Path Queries in Polygonal Domains** Haitao Wang Let $\mathcal{P}$ be a polygonal domain of h holes and n vertices. We study the problem of constructing a data structure that can compute a shortest path between s and t in $\mathcal{P}$ under the L_1 metric for any two query points s and t. To do so, a standard approach is to first find a set of n_s "gateways" for s and a set of n_t "gateways" for t such that there exist a shortest s-t path containing a gateway of s and a gateway of t, and then compute a shortest s-t path using these gateways. Previous algorithms all take quadratic $O(n_s \cdot n_t)$ time to solve this problem. In this paper, we propose a divide-and-conquer technique that solves the problem in $O(n_s + n_t \log n_s)$ time. As a consequence, we construct a data structure of $O(n + (h^2 \log^3 h/\log\log h))$ size in $O(n + (h^2 \log^4 h/\log\log h))$ time such that each query can be answered in $O(\log n)$ time.	**Topological Data Analysis in Information Space** Herbert Edelsbrunner, Ziga Virk, Hubert Wagner Various kinds of data are routinely represented as discrete probability distributions. Examples include text documents summarized by histograms of word occurrences and images represented as histograms of oriented gradients. Viewing a discrete probability distribution as a point in the standard simplex of the appropriate dimension, we can understand collections of such objects in geometric and topological terms. Importantly, instead of using the standard Euclidean distance, we look into dissimilarity measures with information-theoretic justification, and we develop the theory needed for applying topological data analysis in this setting. In doing so, we emphasize constructions that enable usage of existing computational topology software in this context.

<table>
<tr><td>11:30-11:50</td><td>Maintaining the Union of Unit Discs under Insertions with Near-Optimal Overhead
Pankaj K. Agarwal, Ravid Cohen, Dan Halperin and Wolfgang Mulzer

We present efficient data structures for problems on unit discs and arcs of their boundary in the plane. (i) We give an output-sensitive algorithm for the dynamic maintenance of the union of n unit discs under insertions in $O(k \log^2 n)$ update time and $O(n)$ space, where k is the combinatorial complexity of the structural change in the union due to the insertion of the new disc. (ii) As part of the solution of (i) we devise a fully dynamic data structure for the maintenance of lower envelopes of pseudo-lines, which we believe is of independent interest. The structure has $O(\log^2 n)$ update time and $O(\log n)$ vertical ray shooting query time. To achieve this performance, we devise a new algorithm for finding the intersection between two lower envelopes of pseudo-lines in $O(\log n)$ time, using tentative binary search; the lower envelopes are special in that at $x = -\infty$ any pseudo-line contributing to the first envelope lies below every pseudo-line contributing to the second envelope. (iii) We also present a dynamic range searching structure for a set of circular arcs of unit radius (not necessarily on the boundary of the union of the corresponding discs), where the ranges are unit discs, with $O(n \log n)$ preprocessing time, $O(n^{1/2+\varepsilon} + \ell)$ query time and $O(\log^2 n)$ amortized update time, where ℓ is the size of the output and for any $\varepsilon > 0$. The structure requires $O(n)$ storage space</td><td>On the Metric Distortion of Embedding Persistence Diagrams into separable Hilbert spaces
M. Carrière and U. Bauer

Persistence diagrams are important descriptors in Topological Data Analysis. Due to the nonlinearity of the space of persistence diagrams equipped with their diagram distances, most of the recent attempts at using persistence diagrams in machine learning have been done through kernel methods, i.e., embeddings of persistence diagrams into Reproducing Kernel Hilbert Spaces, in which all computations can be performed easily. Since persistence diagrams enjoy theoretical stability guarantees for the diagram distances, the metric properties of the feature map, i.e., the relationship between the Hilbert distance and the diagram distances, are of central interest for understanding if the persistence diagram guarantees carry over to the embedding. In this article, we study the possibility of embedding persistence diagrams into separable Hilbert spaces with bi-Lipschitz maps. In particular, we show that for several stable embeddings into infinite-dimensional Hilbert spaces defined in the literature, any lower bound must depend on the cardinalities of the persistence diagrams, and that when the Hilbert space is finite dimensional, finding a bi-Lipschitz embedding is impossible, even when restricting the persistence diagrams to have bounded cardinalities.</td></tr>
<tr><td>11:50-12:00</td><td colspan="2">Break</td></tr>
<tr><td></td><td>Tue-3A: Combinatorial Geometry I</td><td>Tue-3B: ε-Nets and VC Dimension</td></tr>
<tr><td>12:00-12:20</td><td>On the Complexity of the k-Level in Arrangements of Pseudoplanes
M. Sharir and C. Ziv

A classical open problem in combinatorial geometry is to obtain tight asymptotic bounds on the maximum number of k-level vertices in an arrangement of n hyperplanes in $\mathbb{R}^d$ (vertices with exactly k of the hyperplanes passing below them). This is essentially a dual version of the k-set problem, which, in a primal setting, seeks bounds for the maximum number of k-sets determined by n points in $\mathbb{R}^d$, where a k-set is a subset of size k that can be separated from its complement by a hyperplane. The k-set problem is still wide open even in the plane. In three dimensions, the best known upper and lower bounds are, respectively, $O(nk^{3/2})$ and $nk \cdot 2^{\Omega(\sqrt{\log k})}$.

In its dual version, the problem can be generalized by replacing hyperplanes by other families of surfaces (or curves in the planes). Reasonably sharp bounds have been obtained for curves in the plane, but the known upper bounds are rather weak for more general surfaces, already in three dimensions, except for the case of triangles. The best known general bound, due to Chan is $O(n^{2.997})$, for families of surfaces that satisfy certain (fairly weak) properties.

In this paper we consider the case of pseudoplanes in $\mathbb{R}^3$ (defined in detail in the introduction), and establish the upper bound $O(nk^{5/3})$ for the number of k-level vertices in an arrangement of n pseudoplanes. The bound is obtained by establishing suitable (and nontrivial) extensions of dual versions of classical tools that have been used in studying the primal k-set problem, such as the Lovász Lemma and the Crossing Lemma.</td><td>On weak ε-nets and the Radon number
S. Moran and A. Yehudayoff

We show that the Radon number characterizes the existence of weak nets in separable convexity spaces (an abstraction of the Euclidean notion of convexity). The construction of weak nets when the Radon number is finite is based on Helly's property and on metric properties of VC classes. The lower bound on the size of weak nets when the Radon number is large relies on the chromatic number of the Kneser graph. As an application, we prove an amplification result for weak ε-nets.</td></tr>
</table>

12:20-12:40	**On grids in point-line arrangements in the plane** M. Mirzaei and A. Suk The famous Szemerédi-Trotter theorem states that any arrangement of n points and n lines in the plane determines $O(n^{4/3})$ incidences, and this bound is tight. In this paper, we prove the following Turán-type result for point-line incidence. Let $\mathcal{L}_a$ and $\mathcal{L}_b$ be two sets of t lines in the plane and let $P = \{\ell_a \cap \ell_b : \ell_a \in \mathcal{L}_a, \ell_b \in \mathcal{L}_b\}$ be the set of intersection points between $\mathcal{L}_a$ and $\mathcal{L}_b$. We say that $(P, \mathcal{L}_a \cup \mathcal{L}_b)$ forms a *natural* $t \times t$ *grid* if $\lvert P\rvert = t^2$, and $conv(P)$ does not contain the intersection point of some two lines in $\mathcal{L}_a$ and does not contain the intersection point of some two lines in $\mathcal{L}_b$. For fixed $t > 1$, we show that any arrangement of n points and n lines in the plane that does not contain a natural $t \times t$ grid determines $O(n^{\frac{4}{3}-\varepsilon})$ incidences, where $\varepsilon = \varepsilon(t) > 0$. We also provide a construction of n points and n lines in the plane that does not contain a natural 2×2 grid and determines at least $\Omega(n^{1+\frac{1}{14}})$ incidences.	**Distribution-Sensitive Bounds on Relative Approximations of Geometric Ranges** Y. Tao and Y. Wang A family $\mathcal{R}$ of ranges and a set X of points, all in $\mathbb{R}^d$, together define a range space $(X, \mathcal{R}\lvert_X)$, where $\mathcal{R}\lvert_X = \{X \cap h \mid h \in \mathcal{R}\}$. We want to find a structure to estimate the quantity $\lvert X \cap h\rvert/\lvert X\rvert$ for any range $h \in \mathcal{R}$ with the (ρ, ϵ)-*guarantee*: (i) if $\lvert X \cap h\rvert/\lvert X\rvert > \rho$, the estimate must have a relative error ϵ; (ii) otherwise, the estimate must have an absolute error $\rho\epsilon$. The objective is to minimize the size of the structure. Currently, the dominant solution is to compute a relative (ρ, ϵ)-approximation, which is a subset of X with $\tilde{O}(\lambda/(\rho\epsilon^2))$ points, where λ is the VC-dimension of $(X, \mathcal{R}\lvert_X)$, and $\tilde{O}$ hides polylog factors. This paper shows a more general bound sensitive to the content of X. We give a structure that stores $O(\log(1/\rho))$ integers plus $\tilde{O}(\theta \cdot (\lambda/\epsilon^2))$ points of X, where θ — called the *disagreement coefficient* — measures how much the ranges differ from each other in their intersections with X. The value of θ is between 1 and $1/\rho$, such that our space bound is never worse than that of relative (ρ, ϵ)-approximations, but we improve the latter's $1/\rho$ term whenever $\theta = o(\frac{1}{\rho \log(1/\rho)})$. We also prove that, in the worst case, summaries with the $(\rho, 1/2)$-guarantee must consume $\Omega(\theta)$ words even for $d = 2$ and $\lambda \leq 3$. We then constrain $\mathcal{R}$ to be the set of halfspaces in $\mathbb{R}^d$ for a constant d, and prove the existence of structures with $o(1/(\rho\epsilon^2))$ size offering (ρ, ϵ)-guarantees, when X is generated from various stochastic distributions. This is the first formal justification on why the term $1/\rho$ is not compulsory for "realistic" inputs.
12:40-1:00	**The Crossing Tverberg Theorem** R. Fulek and B. Gärtner and A. Kupavskii and P. Valtr and U. Wagner The Tverberg theorem is one of the cornerstones of discrete geometry. It states that, given a set X of at least $(d+1)(r-1)+1$ points in $\mathbb{R}^d$, one can find a partition $X = X_1 \cup \ldots \cup X_r$ of X, such that the convex hulls of the X_i, $i = 1, \ldots, r$, all share a common point. In this paper, we prove a strengthening of this theorem that guarantees a partition which, in addition to the above, has the property that the boundaries of full-dimensional convex hulls have pairwise nonempty intersections. Possible generalizations and algorithmic aspects are also discussed. As a concrete application, we show that any n points in the plane in general position span $\lfloor n/3 \rfloor$ vertex-disjoint triangles that are pairwise crossing, meaning that their boundaries have pairwise nonempty intersections; this number is clearly best possible. A previous result of Alvarez-Rebollar et al. guarantees $\lfloor n/6 \rfloor$ pairwise crossing triangles. Our result generalizes to a result about simplices in $\mathbb{R}^d, d \geq 2$.	**Journey to the Center of the Point Set** S. Har-Peled and M. Jones We revisit an algorithm of Clarkson et al. (Internat. J. Comput. Geom. Appl., 6.03 (1996) 357), that computes (roughly) a $1/(4d^2)$-centerpoint in $\widetilde{O}(d^9)$ time, for a point set in $\Re^d$, where $\widetilde{O}$ hides polylogarithmic terms. We present an improved algorithm that computes (roughly) a $1/d^2$-centerpoint with running time $\widetilde{O}(d^7)$. While the improvements are (arguably) mild, it is the first progress on this well known problem in over twenty years. The new algorithm is simpler, and the running time bound follows by a simple random walk argument, which we believe to be of independent interest. We also present several new applications of the improved centerpoint algorithm.

Wednesday, June 19		
	Wed-4A: Smallest Enclosing	**Wed-4B: Persistent Homology II**
9:00-9:20	**Probabilistic Smallest Enclosing Ball in High Dimensions via Subgradient Sampling** A. Krivošija and A. Munteanu We study a variant of the median problem for a collection of point sets in high dimensions. This generalizes the geometric median as well as the (probabilistic) smallest enclosing ball (pSEB) problems. Our main objective and motivation is to improve the previously best algorithm for the pSEB problem by reducing its exponential dependence on the dimension to linear. This is achieved via a novel combination of sampling techniques for clustering problems in metric spaces with the framework of stochastic subgradient descent. As a result, the algorithm becomes applicable to shape fitting problems in Hilbert spaces of unbounded dimension via kernel functions. We present an exemplary application by extending the support vector data description (SVDD) shape fitting method to the probabilistic case. This is done by simulating the pSEB algorithm implicitly in the feature space induced by the kernel function.	**Exact computation of the matching distance on 2-parameter persistence modules** Michael Kerber, Michael Lesnick and Steve Oudot The matching distance is a pseudometric on multi-parameter persistence modules, defined in terms of the weighted bottleneck distance on the restriction of the modules to affine lines. It is known that this distance is stable in a reasonable sense, and can be efficiently approximated, which makes it a promising tool for practical applications. In this work, we show that in the 2-parameter setting, the matching distance can be computed exactly in polynomial time. Our approach subdivides the space of affine lines into regions, via a line arrangement in the dual space. In each region, the matching distance restricts to a simple analytic function, whose maximum is easily computed. As a byproduct, our analysis establishes that the matching distance is a rational number, if the bigrades of the input modules are rational.
9:20-9:40	**Smallest k-Enclosing Rectangle Revisited** T. M. Chan and S. Har-Peled Given a set of n points in the plane, and a parameter k, we consider the problem of computing the minimum (perimeter or area) axis-aligned rectangle enclosing k points. We present the first near quadratic time algorithm for this problem, improving over the previous near-$O(n^{5/2})$-time algorithm by Kaplan, Roy, and Sharir [ESA 2017]. We provide an almost matching conditional lower bound, under the assumption that $(\min, +)$-convolution cannot be solved in truly subquadratic time. Furthermore, we present a new reduction (for either perimeter or area) that can make the time bound sensitive to k, giving near $O(nk)$ time. We also present a near linear time $(1+\varepsilon)$-approximation algorithm to the minimum area of the optimal rectangle containing k points. In addition, we study related problems including the 3-sided, arbitrarily oriented, weighted, and subset sum versions of the problem.	**Chunk Reduction for Multi-Parameter Persistent Homology** U. Fugacci and M. Kerber The extension of persistent homology to multi-parameter setups is an algorithmic challenge. Since most computation tasks scale badly with the size of the input complex, an important pre-processing step consists of simplifying the input while maintaining the homological information. We present an algorithm that drastically reduces the size of an input. Our approach is an extension of the chunk algorithm for persistent homology (Bauer et al., Topological Methods in Data Analysis and Visualization III, 2014). We show that our construction produces the smallest multi-filtered chain complex among all the complexes quasi-isomorphic to the input, improving on the guarantees of previous work in the context of discrete Morse theory. Our algorithm also offers an immediate parallelization scheme in shared memory. Already its sequential version compares favorably with existing simplification schemes, as we show by experimental evaluation.
9:40-10:00	**Computing Shapley Values in the Plane** S. Cabello and T. M. Chan We consider the problem of computing Shapley values for points in the plane, where each point is interpreted as a player, and the value of a coalition is defined by the area of usual geometric objects, such as the convex hull or the minimum axis-parallel bounding box. For sets of n points in the plane, we show how to compute in roughly $O(n^{3/2})$ time the Shapley values for the area of the minimum axis-parallel bounding box and the area of the union of the rectangles spanned by the origin and the input points. When the points form an increasing or decreasing chain, the running time can be improved to near-linear. In all these cases, we use linearity of the Shapley values and algebraic methods. We also show that Shapley values for the area of the convex hull or the minimum enclosing disk can be computed in $O(n^2)$ and $O(n^3)$ time, respectively. These problems are closely related to the model of stochastic point sets considered in computational geometry, but here we have to consider random insertion orders of the points instead of a probabilistic existence of points.	**Computing Persistent Homology of Flag Complexes via Strong Collapses** J-D. Boissonnat and S. Pritam In this article, we focus on the problem of computing Persistent Homology of a flag tower, i.e. a sequence of flag complexes connected by simplicial maps. We show that if we restrict the class of simplicial complexes to flag complexes, we can achieve decisive improvement in terms of time and space complexities with respect to previous work. We show that strong collapses of flag complexes can be computed in time $O(k^2v^2)$ where v is the number of vertices of the complex and k is the maximal degree of its graph. Moreover we can strong collapse a flag complex knowing only its 1-skeleton and the resulting complex is also a flag complex. When we strong collapse the complexes in a flag tower, we obtain a reduced sequence that is also a flag tower we call the core flag tower. We then convert the core flag tower to an equivalent filtration to compute its PH. Here again, we only use the 1-skeletons of the complexes. The resulting method is simple and extremely efficient.

10:00-10:30	Coffee Break	
	Wed-5A: Combinatorial Geometry II	**Wed-5B: Optimization and Approximation**
10:30-10:50	**Ham-Sandwich cuts and center transversals in subspaces** Patrick Schnider The Ham-Sandwich theorem is a well-known result in geometry. It states that any d mass distributions in $\mathbb{R}^d$ can be simultaneously bisected by a hyperplane. The result is tight, that is, there are examples of $d+1$ mass distributions that cannot be simultaneously bisected by a single hyperplane. In this abstract we will study the following question: given a continuous assignment of mass distributions to certain subsets of $\mathbb{R}^d$, is there a subset on which we can bisect more masses than what is guaranteed by the Ham-Sandwich theorem? which we answer in the affirmative. We investigate two types of subsets. The first type are linear subspaces of $\mathbb{R}^d$, i.e., k-dimensional flats containing the origin. We show that for any continuous assignment of d mass distributions to the k-dimensional linear subspaces of $\mathbb{R}^d$, there is always a subspace on which we can simultaneously bisect the images of all d assignments. We extend this result to center transversals, a generalization of Ham-Sandwich cuts. As for Ham-Sandwich cuts, we further show that for $d-k+2$ masses, we can choose $k-1$ of the vectors defining the k-dimensional subspace in which the solution lies. The second type of subsets we consider are subsets that are determined by families of n hyperplanes in $\mathbb{R}^d$. Also in this case, we find a Ham-Sandwich-type result. In an attempt to solve a conjecture by Langerman about bisections with several cuts, we show that our underlying topological result can be used to prove this conjecture in a relaxed setting.	**Packing Disks into Disks with Optimal Worst-Case Density** S. P. Fekete and P. Keldenich and C. Scheffer We provide a tight result for a fundamental problem arising from packing disks into a circular container: The critical density of packing disks in a disk is 0.5. This implies that any set of (not necessarily equal) disks of total area $\delta \leq 1/2$ can always be packed into a disk of area 1; on the other hand, for any $\varepsilon > 0$ there are sets of disks of area $1/2+\varepsilon$ that cannot be packed. The proof uses a careful manual analysis, complemented by a minor automatic part that is based on interval arithmetic. Beyond the basic mathematical importance, our result is also useful as a blackbox lemma for the analysis of recursive packing algorithms.
10:50-11:10	**On the chromatic number of disjointness graphs of curves** János Pach and István Tomon Let $\omega(G)$ and $\chi(G)$ denote the clique number and chromatic number of a graph G, respectively. The disjointness graph of a family of curves (continuous arcs in the plane) is the graph whose vertices correspond to the curves and in which two vertices are joined by an edge if and only if the corresponding curves are disjoint. A curve is called x-monotone if every vertical line intersects it in at most one point. An x-monotone curve is grounded if its left endpoint lies on the y-axis. We prove that if G is the disjointness graph of a family of grounded x-monotone curves such that $\omega(G)=k$, then $\chi(G) \leq \binom{k+1}{2}$. If we only require that every curve is x-monotone and intersects the y-axis, then we have $\chi(G) \leq \frac{k+1}{2}\binom{k+2}{3}$. Both of these bounds are best possible. The construction showing the tightness of the last result settles a 25 years old problem: it yields that there exist K_k-free disjointness graphs of x-monotone curves such that any proper coloring of them uses at least $\Omega(k^4)$ colors. This matches the upper bound up to a constant factor.	**Preconditioning for the Geometric Transportation Problem** A. B. Khesin, A. Nikolov, and D. Paramonov In the geometric transportation problem, we are given a collection of points P in d-dimensional Euclidean space, and each point is given a supply of $\mu(p)$ units of mass, where $\mu(p)$ could be a positive or a negative integer, and the total sum of the supplies is 0. The goal is to find a flow (called a transportation map) that transports $\mu(p)$ units from any point p with $\mu(p) > 0$, and transports $-\mu(p)$ units into any point p with $\mu(p) < 0$. Moreover, the flow should minimize the total distance traveled by the transported mass. The optimal value is known as the transportation cost, or the Earth Mover's Distance (from the points with positive supply to those with negative supply). This problem has been widely studied in many fields of computer science: from theoretical work in computational geometry, to applications in computer vision, graphics, and machine learning. In this work we study approximation algorithms for the geometric transportation problem. We give an algorithm which, for any fixed dimension d, finds a $(1+\varepsilon)$-approximate transportation map in time nearly-linear in n, and polynomial in ε^{-1} and in the logarithm of the total supply. This is the first approximation scheme for the problem whose running time depends on n as $n \cdot \mathrm{polylog}(n)$. Our techniques combine the generalized preconditioning framework of Sherman [SODA 2017], which is grounded in continuous optimization, with simple geometric arguments to first reduce the problem to a minimum cost flow problem on a sparse graph, and then to design a good preconditioner for this latter problem.

11:10-11:30	**Semi-algebraic colorings of complete graphs** J. Fox, J. Pach, and A. Suk We consider m-colorings of the edges of a complete graph, where each color class is defined semi-algebraically with bounded complexity. The case $m = 2$ was first studied by Alon et al., who applied this framework to obtain surprisingly strong Ramsey-type results for intersection graphs of geometric objects and for other graphs arising in computational geometry. Considering larger values of m is relevant, e.g., to problems concerning the number of distinct distances determined by a point set. For $p \geq 3$ and $m \geq 2$, the classical Ramsey number $R(p; m)$ is the smallest positive integer n such that any m-coloring of the edges of K_n, the complete graph on n vertices, contains a monochromatic K_p. It is a longstanding open problem that goes back to Schur (1916) to decide whether $R(p; m) = 2^{O(m)}$, for a fixed p. We prove that this is true if each color class is defined semi-algebraically with bounded complexity, and that the order of magnitude of this bound is tight. Our proof is based on the Cutting Lemma of Chazelle *et al.*, and on a Szemerédi-type regularity lemma for multicolored semi-algebraic graphs, which is of independent interest. The same technique is used to address the semi-algebraic variant of a more general Ramsey-type problem of Erdős and Shelah.	**Algorithms for Metric Learning via Contrastive Embeddings** D. Ihara, N. Mohammadi and A. Sidiropoulos We study the problem of supervised learning a metric space under *discriminative* constraints. Given a universe X and sets $\mathcal{S}, \mathcal{D} \subset \binom{X}{2}$ of *similar* and *dissimilar* pairs, we seek to find a mapping $f : X \to Y$, into some target metric space $M = (Y, \rho)$, such that similar objects are mapped to points at distance at most u, and dissimilar objects are mapped to points at distance at least ℓ. More generally, the goal is to find a mapping of maximum *accuracy* (that is, fraction of correctly classified pairs). We propose approximation algorithms for various versions of this problem, for the cases of Euclidean and tree metric spaces. For both of these target spaces, we obtain fully polynomial-time approximation schemes (FPTAS) for the case of perfect information. In the presence of imperfect information we present approximation algorithms that run in quasi-polynomial time (QPTAS). We also present an exact algorithm for learning line metric spaces with perfect information in polynomial time. Our algorithms use a combination of tools from metric embeddings and graph partitioning, that could be of independent interest.
11:30-11:40	Break	
11:40-12:40	Invited Talk: **A Geometric Data Structure from Neuroscience** Sanjoy Dasgupta **Abstract:** An intriguing geometric primitive, "expand-and-sparsify", has been found in the olfactory system of the fly and several other organisms. It maps an input vector to a much higher-dimensional sparse representation, using a random linear transformation followed by winner-take-all thresholding. I'll show that this representation has a variety of formal properties, such as locality preservation, that make it an attractive data structure for algorithms and machine learning. In particular, mimicking the fly's circuitry yields algorithms for similarity search and for novelty detection that have provable guarantees as well as having practical performance that is competitive with state-of-the-art methods. This talk is based on work with Saket Navlakha (Salk Institute), Chuck Stevens (Salk Institute), and Chris Tosh (Columbia). **Bio:** Sanjoy Dasgupta is a Professor of Computer Science and Engineering at UC San Diego, where he has been since 2002. He works on algorithmic statistics, with a particular focus on unsupervised and minimally supervised learning. He is author of a textbook, "Algorithms" (with Christos Papadimitriou and Umesh Vazirani).	

See next page ...

12:40-2:30	Lunch on Your Own	
	Wed-6A: Graph Drawing I	**Wed-6B: Matching and Partitioning**
2:30-2:50	**Efficient Algorithms for Ortho-Radial Graph Drawing** B. Niedermann, I. Rutter, and M. Wolf Orthogonal drawings, i.e., embeddings of graphs into grids, are a classic topic in Graph Drawing. Often the goal is to find a drawing that minimizes the number of bends on the edges. A key ingredient for bend minimization algorithms is the existence of an *orthogonal representation* that allows to describe such drawings purely combinatorially by only listing the angles between the edges around each vertex and the directions of bends on the edges, but neglecting any kind of geometric information such as vertex coordinates or edge lengths. Barth et al. [2017] have established the existence of an analogous *ortho-radial representation* for *ortho-radial drawings*, which are embeddings into an ortho-radial grid, whose gridlines are concentric circles around the origin and straight-line spokes emanating from the origin but excluding the origin itself. While any orthogonal representation admits an orthogonal drawing, it is the circularity of the ortho-radial grid that makes the problem of characterizing valid ortho-radial representations all the more complex and interesting. Barth et al. prove such a characterization. However, the proof is existential and does not provide an efficient algorithm for testing whether a given ortho-radial representation is valid, let alone actually obtaining a drawing from an ortho-radial representation. In this paper we give quadratic-time algorithms for both of these tasks. They are based on a suitably constrained left-first DFS in planar graphs and several new insights on ortho-radial representations. Our validity check requires quadratic time, and a naive application of it would yield a quartic algorithm for constructing a drawing from a valid ortho-radial representation. Using further structural insights we speed up the drawing algorithm to quadratic running time.	**A Weighted Approach to the Maximum Cardinality Bipartite Matching Problem with Applications in Geometric Settings** N. Lahn and S. Raghvendra We present a weighted approach to compute a maximum cardinality matching in an arbitrary bipartite graph. Our main result is a new algorithm that takes as input a weighted bipartite graph $G(A \cup B, E)$ with edge weights of 0 or 1. Let $w \leq n$ be an upper bound on the weight of any matching in G. Consider the subgraph induced by all the edges of G with a weight 0. Suppose every connected component in this subgraph has $O(r)$ vertices and $O(mr/n)$ edges. We present an algorithm to compute a maximum cardinality matching in G in $\tilde{O}(m(\sqrt{w}+\sqrt{r}+\frac{wr}{n}))$ time. When all the edge weights are 1 (symmetrically when all weights are 0), our algorithm will be identical to the well-known Hopcroft-Karp (HK) algorithm, which runs in $O(m\sqrt{n})$ time. However, if we can carefully assign weights of 0 and 1 on its edges such that both w and r are sub-linear in n and $wr = O(n^{\gamma})$ for $\gamma < 3/2$, then we can compute maximum cardinality matching in G in $o(m\sqrt{n})$ time. Using our algorithm, we obtain a new $\tilde{O}(n^{4/3}/\varepsilon^4)$ time algorithm to compute an ε-approximate bottleneck matching of $A, B \subset \mathbb{R}^2$ and an $\frac{1}{\varepsilon^{O(d)}} n^{1+\frac{d-1}{2d-1}} \operatorname{poly} \log n$ time algorithm for computing ε-approximate bottleneck matching in d-dimensions. All previous algorithms take $\Omega(n^{3/2})$ time. Given any graph $G(A \cup B, E)$ that has an easily computable balanced vertex separator for every subgraph $G'(V', E')$ of size $\|V'\|^{\delta}$, for $\delta \in [1/2, 1)$, we can apply our algorithm to compute a maximum matching in $\tilde{O}(mn^{\frac{\delta}{1+\delta}})$ time improving upon the $O(m\sqrt{n})$ time taken by the HK-Algorithm.
2:50-3:10	**Bounded degree conjecture holds precisely for c-crossing-critical graphs with $c \leq 12$** D. Bokal, Z. Dvořák, P. Hliněný, J. Leaños, B. Mohar, T. Wiedera We study c-crossing-critical graphs, which are the minimal graphs that require at least c edge-crossings when drawn in the plane. For every fixed pair of integers with $c \geq 13$ and $d \geq 1$, we give first explicit constructions of c-crossing-critical graphs containing a vertex of degree greater than d. We also show that such unbounded degree constructions do not exist for $c \leq 12$, precisely, that there exists a constant D such that every c-crossing-critical graph with $c \leq 12$ has maximum degree at most D. Hence, the bounded maximum degree conjecture of c-crossing-critical graphs, which was generally disproved in 2010 by Dvořák and Mohar (without an explicit construction), holds true, surprisingly, exactly for the values $c \leq 12$.	**An Efficient Algorithm for Generalized Polynomial Partitioning and Its Applications** P. K. Agarwal, B. Aronov, E. Ezra, and J. Zahl In 2015, Guth proved that if S is a collection of n g-dimensional semi-algebraic sets in $\mathbb{R}^d$ and if $D \geq 1$ is an integer, then there is a d-variate polynomial P of degree at most D so that each connected component of $\mathbb{R}^d \setminus Z(P)$ intersects $O(n/D^{d-g})$ sets from S. Such a polynomial is called a *generalized partitioning polynomial*. We present a randomized algorithm that computes such polynomials efficiently—the expected running time of our algorithm is linear in $\|S\|$. Our approach exploits the technique of *quantifier elimination* combined with that of *ε-samples*. We present four applications of our result. The first is a data structure for answering point-enclosure queries among a family of semi-algebraic sets in $\mathbb{R}^d$ in $O(\log n)$ time, with storage complexity and expected preprocessing time of $O(n^{d+\varepsilon})$. The second is a data structure for answering range search queries with semi-algebraic ranges in $O(\log n)$ time, with $O(n^{t+\varepsilon})$ storage and expected preprocessing time, where $t > 0$ is an integer that depends on d and the description complexity of the ranges. The third is a data structure for answering vertical ray-shooting queries among semi-algebraic sets in $\mathbb{R}^d$ in $O(\log^2 n)$ time, with $O(n^{d+\varepsilon})$ storage and expected preprocessing time. The fourth is an efficient algorithm for cutting algebraic planar curves into pseudo-segments.

3:10-3:30	**$\mathbb{Z}_2$-Genus of Graphs and Minimum Rank of Partial Symmetric Matrices** R. Fulek and J. Kynčl The *genus* $g(G)$ of a graph G is the minimum g such that G has an embedding on the orientable surface M_g of genus g. A drawing of a graph on a surface is *independently even* if every pair of nonadjacent edges in the drawing crosses an even number of times. The $\mathbb{Z}_2$*-genus* of a graph G, denoted by $g_0(G)$, is the minimum g such that G has an independently even drawing on M_g. By a result of Battle, Harary, Kodama and Youngs from 1962, the graph genus is additive over 2-connected blocks. In 2013, Schaefer and Štefankovič proved that the $\mathbb{Z}_2$-genus of a graph is additive over 2-connected blocks as well, and asked whether this result can be extended to so-called 2-amalgamations, as an analogue of results by Decker, Glover, Huneke, and Stahl for the genus. We give the following partial answer. If $G = G_1 \cup G_2$, G_1 and G_2 intersect in two vertices u and v, and $G - u - v$ has k connected components (among which we count the edge uv if present), then $\|g_0(G) - (g_0(G_1) + g_0(G_2))\| \leq k + 1$. For complete bipartite graphs $K_{m,n}$, with $n \geq m \geq 3$, we prove that $\frac{g_0(K_{m,n})}{g(K_{m,n})} = 1 - O(\frac{1}{n})$. Similar results are proved also for the Euler $\mathbb{Z}_2$-genus. We express the $\mathbb{Z}_2$-genus of a graph using the minimum rank of partial symmetric matrices over $\mathbb{Z}_2$; a problem that might be of independent interest.	**Efficient Algorithms for Geometric Partial Matching** Pankaj K. Agarwal, Hsien-Chih Chang, Allen Xiao Let A and B be two point sets in the plane of sizes r and n respectively (assume $r \leq n$), and let k be a parameter.A matching between A and B is a family of pairs in $A \times B$ so that any point of $A \cup B$ appears in at most one pair. Given two positive integers p and q, we define the cost of matching M to be $c(M) = \sum_{(a,b)\in M} \|a - b\|_p^q$ where $\|\cdot\|_p$ is the L_p-norm. The geometric partial matching problem asks to find the minimum-cost size-k matching between A and B. We present efficient algorithms for geometric partial matching problem that work for any powers of L_p-norm matching objective: An exact algorithm that runs in $O((n + k^2) \operatorname{polylog} n)$ time, and a $(1 + \varepsilon)$-approximation algorithm that runs in $O((n + k\sqrt{k}) \operatorname{polylog} n \cdot \log \varepsilon^{-1})$ time. Both algorithms are based on the primal-dual flow augmentation scheme; the main improvements involve using dynamic data structures to achieve efficient flow augmentations. With similar techniques, we give an exact algorithm for the planar transportation problem running in $O(\min\{n^2, rn^{3/2}\} \operatorname{polylog} n)$ time.
3:30-4:00	Coffee/Snack Break	
	Wed-7A: Topology	**Wed-7B: Algorithm Complexity**
4:00-4:20	**Topologically Trivial Closed Walks in Directed Surface Graphs** Jeff Erickson and Yipu Wang Let G be a directed graph with n vertices and m edges, embedded on a surface S, possibly with boundary, with first Betti number β. We consider the complexity of finding closed directed walks in G that are either contractible (trivial in homotopy) or bounding (trivial in integer homology) in S. Specifically, we describe algorithms to determine whether G contains a simple contractible cycle in $O(n + m)$ time, or a contractible closed walk in $O(n + m)$ time, or a bounding closed walk in $O(\beta(n + m))$ time. Our algorithms rely on subtle relationships between strong connectivity in G and in the dual graph $G^\star$; our contractible-closed-walk algorithm also relies on a seminal topological result of Hass and Scott. We also prove that detecting simple bounding cycles is NP-hard. We also describe three polynomial-time algorithms to compute shortest contractible closed walks, depending on whether the fundamental group of the surface is free, abelian, or hyperbolic. A key step in our algorithm for hyperbolic surfaces is the construction of a context-free grammar with $O(g^2L^2)$ non-terminals that generates all contractible closed walks of length at most L, and only contractible closed walks, in a system of quads of genus $g \geq 2$. Finally, we show that computing shortest simple contractible cycles, shortest simple bounding cycles, and shortest bounding closed walks are all NP-hard.	**The One-Way Communication Complexity of Dynamic Time Warping Distance** V. Braverman, M. Charikar, W. Kuszmaul, D. P. Woodruff, and L. F. Yang We resolve the randomized one-way communication complexity of Dynamic Time Warping (DTW) distance. We show that there is an efficient one-way communication protocol using $\widetilde{O}(n/\alpha)$ bits for the problem of computing an α-approximation for DTW between strings x and y of length n, and we prove a lower bound of $\Omega(n/\alpha)$ bits for the same problem. Our communication protocol works for strings over an arbitrary metric of polynomial size and aspect ratio, and we optimize the logarithmic factors depending on properties of the underlying metric, such as when the points are low-dimensional integer vectors equipped with various metrics or have bounded doubling dimension. We also consider linear sketches of DTW, showing that such sketches must have size $\Omega(n)$.

4:20-4:40	**3-Manifold Triangulations with Small Treewidth** K. Huszár and J. Spreer Motivated by fixed-parameter tractable (FPT) problems in computational topology, we consider the treewidth $\mathrm{tw}(\mathcal{M})$ of a compact, connected 3-manifold $\mathcal{M}$, defined to be the minimum treewidth of the face pairing graph of any triangulation $\mathcal{T}$ of $\mathcal{M}$. In this setting the relationship between the topology of a 3-manifold and its treewidth is of particular interest. First, as a corollary of work of Jaco and Rubinstein, we prove that for any closed, orientable 3-manifold $\mathcal{M}$ the treewidth $\mathrm{tw}(\mathcal{M})$ is at most $4\mathfrak{g}(\mathcal{M}) - 2$, where $\mathfrak{g}(\mathcal{M})$ denotes Heegaard genus of $\mathcal{M}$. In combination with our earlier work with Wagner, this yields that for non-Haken manifolds the Heegaard genus and the treewidth are within a constant factor. Second, we characterize all 3-manifolds of treewidth one: These are precisely the lens spaces and a single other Seifert fibered space. Furthermore, we show that all remaining orientable Seifert fibered spaces over the 2-sphere or a non-orientable surface have treewidth two. In particular, for every spherical 3-manifold we exhibit a triangulation of treewidth at most two. Our results further validate the parameter of treewidth (and other related parameters such as cutwidth or congestion) to be useful for topological computing, and also shed more light on the scope of existing FPT-algorithms in the field.	**Upward Book Embeddings of st-Graphs** C. Binucci, G. Da Lozzo, E. Di Giacomo, W. Didimo, T. Mchedlidze, M. Patrignani We study *k-page upward book embeddings* (kUBEs) of *st*-graphs, that is, book embeddings of single-source single-sink directed acyclic graphs on k pages with the additional requirement that the vertices of the graph appear in a topological ordering along the spine of the book. We show that testing whether a graph admits a kUBE is NP-complete for $k \geq 3$. A hardness result for this problem was previously known only for $k = 6$ [Heath and Pemmaraju, 1999]. Motivated by this negative result, we focus our attention on $k = 2$. On the algorithmic side, we present polynomial-time algorithms for testing the existence of 2UBEs of planar *st*-graphs with branchwidth β and of plane *st*-graphs whose faces have a special structure. These algorithms run in $O(f(\beta) \cdot n + n^3)$ time and $O(n)$ time, respectively, where f is a singly-exponential function on β. Moreover, on the combinatorial side, we present two notable families of plane *st*-graphs that always admit an embedding-preserving 2UBE.
4:40-5:00	**When Convexity Helps Collapsing Complexes** D. Attali, A. Lieutier, and D. Salinas This paper illustrates how convexity hypotheses help collapsing simplicial complexes. We first consider a collection of compact convex sets and show that the nerve of the collection is collapsible whenever the union of sets in the collection is convex. We apply this result to prove that the Delaunay complex of a finite point set is collapsible. We then consider a convex domain defined as the convex hull of a finite point set. We show that if the point set samples sufficiently densely the domain, then both the Čech complex and the Rips complex of the point set are collapsible for a well-chosen scale parameter. A key ingredient in our proofs consists in building a filtration by sweeping space with a growing sphere whose center has been fixed and studying events occurring through the filtration. Since the filtration mimics the sublevel sets of a Morse function with a single critical point, we anticipate this work to lay the foundations for a non-smooth, discrete Morse Theory.	

Thursday, June 20

	Thu-8A: Contact and Surface Graphs	Thu-8B: Frechet Distance
9:00-9:20	**Near-optimal Algorithms for Shortest Paths in Weighted Unit-Disk Graphs** H. Wang, J. Xue We revisit a classical graph-theoretic problem, the *single-source shortest-path* (SSSP) problem, in weighted unit-disk graphs. We first propose an exact (and deterministic) algorithm which solves the problem in $O(n \log^2 n)$ time using linear space, where n is the number of the vertices of the graph. This significantly improves the previous deterministic algorithm by Cabello and Jejčič [CGTA'15] which uses $O(n^{1+\delta})$ time and $O(n^{1+\delta})$ space (for any small constant $\delta > 0$) and the previous randomized algorithm by Kaplan et al. [SODA'17] which uses $O(n \log^{12+o(1)} n)$ expected time and $O(n \log^3 n)$ space. More specifically, we show that if the 2D offline insertion-only (additively-)weighted nearest-neighbor problem with k operations (i.e., insertions and queries) can be solved in $f(k)$ time, then the SSSP problem in weighted unit-disk graphs can be solved in $O(n \log n + f(n))$ time. Using the same framework with some new ideas, we also obtain a $(1+\varepsilon)$-approximate algorithm for the problem, using $O(n \log n + n \log^2(1/\varepsilon))$ time and linear space. This improves the previous $(1+\varepsilon)$-approximate algorithm by Chan and Skrepetos [SoCG'18] which uses $O((1/\varepsilon)^2 n \log n)$ time and $O((1/\varepsilon)^2 n)$ space. Because of the $\Omega(n \log n)$-time lower bound of the problem (even when approximation is allowed), both of our algorithms are almost optimal.	**The VC Dimension of Metric Balls under Fréchet and Hausdorff Distances** A. Driemel, J. M. Phillips, I. Psarros The Vapnik-Chervonenkis dimension provides a notion of complexity for systems of sets. If the VC dimension is small, then knowing this can drastically simplify fundamental computational tasks such as classification, range counting, and density estimation through the use of sampling bounds. We analyze set systems where the ground set X is a set of polygonal curves in $\mathbb{R}^d$ and the sets $\mathcal{R}$ are metric balls defined by curve similarity metrics, such as the Fréchet distance and the Hausdorff distance, as well as their discrete counterparts. We derive upper and lower bounds on the VC dimension that imply useful sampling bounds in the setting that the number of curves is large, but the complexity of the individual curves is small. Our upper bounds are either near-quadratic or near-linear in the complexity of the curves that define the ranges and they are logarithmic in the complexity of the curves that define the ground set.
9:20-9:40	**Morphing Contact Representations of Graphs** Patrizio Angelini, Steven Chaplick, Sabine Cornelsen, Giordano Da Lozzo, Vincenzo Roselli We consider the problem of morphing between contact representations of a plane graph. In a contact representation of a plane graph, vertices are realized by internally disjoint elements from a family of connected geometric objects. Two such elements touch if and only if their corresponding vertices are adjacent. These touchings also induce the same embedding as in the graph. In a morph between two contact representations we insist that at each time step (continuously throughout the morph) we have a contact representation of the same type. We focus on the case when the geometric objects are triangles that are the lower-right half of axis-parallel rectangles. Such RT-representations exist for every plane graph and right triangles are one of the simplest families of shapes supporting this property. Thus, they provide a natural case to study regarding morphs of contact representations of plane graphs. We study piecewise linear morphs, where each step is a linear morph moving the endpoints of each triangle at constant speed along straight-line trajectories. We provide a polynomial-time algorithm that decides whether there is a piecewise linear morph between two RT-representations of a plane triangulation, and, if so, computes a morph with a quadratic number of linear morphs. As a direct consequence, we obtain that for 4-connected plane triangulations there is a morph between every pair of RT-representations where the "top-most" triangle in both representations corresponds to the same vertex. This shows that the realization space of such RT-representations of any 4-connected plane triangulation forms a connected set.	**Walking the Dog Fast in Practice: Algorithm Engineering of the Fréchet Distance** K. Bringmann, M. Künnemann and A. Nusser The Fréchet distance provides a natural and intuitive measure for the popular task of computing the similarity of two (polygonal) curves. While a simple algorithm computes it in near-quadratic time, a strongly subquadratic algorithm cannot exist unless the Strong Exponential Time Hypothesis fails. Still, fast practical implementations of the Fréchet distance, in particular for realistic input curves, are highly desirable. This has even lead to a designated competition, the ACM SIGSPATIAL GIS Cup 2017: Here, the challenge was to implement a near-neighbor data structure under the Fréchet distance. The bottleneck of the top three implementations turned out to be precisely the decision procedure for the Fréchet distance. In this work, we present a fast, certifying implementation for deciding the Fréchet distance, in order to (1) complement its pessimistic worst-case hardness by an empirical analysis on realistic input data and to (2) improve the state of the art for the GIS Cup challenge. We experimentally evaluate our implementation on a large benchmark consisting of several data sets (including handwritten characters and GPS trajectories). Compared to the winning implementation of the GIS Cup, we obtain running time improvements of up to more than two orders of magnitude for the decision procedure and of up to a factor of 30 for queries to the near-neighbor data structure.

9:40-10:00	**Lower Bounds for Electrical Reduction on Surfaces** Hsien-Chih Chang, Marcos Cossarini, Jeff Erickson We strengthen the connections between *electrical transformations* and *homotopy* from the planar setting—observed and studied since Steinitz—to arbitrary surfaces with punctures. As a result, we improve our earlier lower bound on the number of electrical transformations required to reduce an n-vertex graph on surface in the worst case [SOCG 2016] in two different directions. Our previous $\Omega(n^{3/2})$ lower bound applies only to *facial* electrical transformations on plane graphs with *no terminals*. First we provide a stronger $\Omega(n^2)$ lower bound when the planar graph has two or more terminals, which follows from a quadratic lower bound on the number of homotopy moves in the annulus. Our second result extends our earlier $\Omega(n^{3/2})$ lower bound to the wider class of *planar* electrical transformations, which preserve the planarity of the graph but may delete cycles that are not faces of the given embedding. This new lower bound follows from the observation that the *defect* of the medial graph of a planar graph is the same for all its planar embeddings.	**Polyline Simplification has Cubic Complexity** K. Bringmann and B. R. Chaudhury In the classic polyline simplification problem we want to replace a given polygonal curve P, consisting of n vertices, by a subsequence P' of k vertices from P such that the polygonal curves P and P' are 'close'. Closeness is usually measured using the Hausdorff or Fréchet distance. These distance measures can be applied *globally*, i.e., to the whole curves P and P', or *locally*, i.e., to each simplified subcurve and the line segment that it was replaced with separately (and then taking the maximum). We provide an $O(n^3)$ time algorithm for simplification under Global-Fréchet distance, improving the previous best algorithm by a factor of $\Omega(kn^2)$. We also provide evidence that in high dimensions cubic time is essentially optimal for all three problems (Local-Hausdorff, Local-Fréchet, and Global-Fréchet). Specifically, improving the cubic time to $O(n^{3-\epsilon}\mathrm{poly}(d))$ for polyline simplification over $(\mathbb{R}^d, L_p)$ for $p = 1$ would violate plausible conjectures. We obtain similar results for all $p \in [1, \infty), p \neq 2$. In total, in high dimensions and over general L_p-norms we resolve the complexity of polyline simplification with respect to Local-Hausdorff, Local-Fréchet, and Global-Fréchet, by providing new algorithms and conditional lower bounds.
10:00-10:30	Coffee Break	
	Thu-9A: Geometric Data Structures	**Thu-9B: Robotics and Geometric Structures**
10:30-10:50	**A Spanner for the Day After** K. Buchin, S. Har-Peled and D. Oláh We show how to construct $(1+\varepsilon)$-spanner over a set P of n points in $\mathbb{R}^d$ that is resilient to a catastrophic failure of nodes. Specifically, for prescribed parameters $\vartheta, \varepsilon \in (0,1)$, the computed spanner G has $O\left(\varepsilon^{-c}\vartheta^{-6} n \log n (\log\log n)^6\right)$ edges, where $c = O(d)$. Furthermore, for *any* k, and *any* deleted set $B \subseteq P$ of k points, the residual graph $G \setminus B$ is $(1+\varepsilon)$-spanner for all the points of P except for $(1+\vartheta)k$ of them. No previous constructions, beyond the trivial clique with $O(n^2)$ edges, were known such that only a tiny additional fraction (i.e., ϑ) lose their distance preserving connectivity. Our construction works by first solving the exact problem in one dimension, and then showing a surprisingly simple and elegant construction in higher dimensions, that uses the one-dimensional construction in a black box fashion.	**General techniques for approximate incidences and their application to the camera posing problem** D. Aiger, H. Kaplan, E. Kokiopoulou, M. Sharir, B. Zeisl We consider the classical camera pose estimation problem that arises in many computer vision applications, in which we are given n 2D-3D correspondences between points in the scene and points in the camera image (some of which are incorrect associations), and where we aim to determine the camera pose (the position and orientation of the camera in the scene) from this data. We demonstrate that this posing problem can be reduced to the problem of computing ε-approximate incidences between two-dimensional surfaces (derived from the input correspondences) and points (on a grid) in a four-dimensional pose space. Similar reductions can be applied to other camera pose problems, as well as to similar problems in related application areas. We describe and analyze three techniques for solving the resulting ε-approximate incidences problem in the context of our camera posing application. The first is a straightforward assignment of surfaces to the cells of a grid (of side-length ε) that they intersect. The second is a variant of a primal-dual technique, recently introduced by a subset of the authors [ESA17] for different (and simpler) applications. The third is a non-trivial generalization of a data structure Fonseca and Mount [CGTA2010], originally designed for the case of hyperplanes. We present and analyze this technique in full generality, and then apply it to the camera posing problem at hand. We compare our methods experimentally on real and synthetic data. Our experiments show that for the typical values of n and ε, the primal-dual method is the fastest, also in practice.

10:50-11:10	**Searching for the Closest-pair in a Query Translate** J. Xue, Y. Li, S. Rahul, R. Janardan We consider a range-search variant of the closest-pair problem. Let Γ be a fixed shape in the plane. We are interested in storing a given set of n points in the plane in some data structure such that for any specified translate of Γ, the closest pair of points contained in the translate can be reported efficiently. We present results on this problem for two important settings: when Γ is a polygon (possibly with holes) and when Γ is a general convex body whose boundary is smooth. When Γ is a polygon, we present a data structure using $O(n)$ space and $O(\log n)$ query time, which is asymptotically optimal. When Γ is a general convex body with a smooth boundary, we give a near-optimal data structure using $O(n \log n)$ space and $O(\log^2 n)$ query time. Our results settle some open questions posed by Xue et al. at SoCG 2018.	**Rods and Rings: Soft Subdivision Planner for $\mathbb{R}^3 \times S^2$** C.-H. Hsu, Y.-J. Chiang and C. Yap We consider path planning for a rigid spatial robot moving amidst polyhedral obstacles. Our robot is either a rod or a ring. Being axially-symmetric, their configuration space is $\mathbb{R}^3 \times S^2$ with 5 degrees of freedom (DOF). Correct, complete and practical path planning for such robots is a long standing challenge in robotics. While the rod is one of the most widely studied spatial robots in path planning, the ring seems to be new, and a rare example of a non-simply-connected robot. This work provides rigorous and complete algorithms for these robots with theoretical guarantees. We implemented the algorithms in our open-source Core Library. Experiments show that they are practical, achieving near real-time performance. We compared our planner to state-of-the-art sampling planners in the Open Motion Planning Library (OMPL). Our subdivision path planner is based on the twin foundations of ε-exactness and soft predicates. Correct implementation is relatively easy. The technical innovations include subdivision atlases for S^2, introduction of Σ_2 representations for footprints, and extensions of our feature-based technique for "opening up the blackbox of collision detection".
11:10-11:30	**Preprocessing Ambiguous Imprecise Points** I. van der Hoog, I. Kostitsyna, M. Löffler, B. Speckmann Let $\mathcal{R} = \{R_1, R_2, \ldots, R_n\}$ be a set of regions and let $X = \{x_1, x_2, \ldots, x_n\}$ be an (unknown) point set with $x_i \in R_i$. Region R_i represents the uncertainty region of x_i. We consider the following question: how fast can we establish order if we are allowed to preprocess the regions in $\mathcal{R}$? The *preprocessing model* of uncertainty uses two consecutive phases: a preprocessing phase which has access only to $\mathcal{R}$ followed by a reconstruction phase during which a desired structure on X is computed. Recent results in this model parametrize the reconstruction time by the *ply* of $\mathcal{R}$, which is the maximum overlap between the regions in $\mathcal{R}$. We introduce the *ambiguity* $\mathcal{A}(\mathcal{R})$ as a more fine-grained measure of the degree of overlap in $\mathcal{R}$. We show how to preprocess a set of d-dimensional disks in $O(n \log n)$ time such that we can sort X (if $d = 1$) and reconstruct a quadtree on X (if $d \geq 1$ but constant) in $O(\mathcal{A}(\mathcal{R}))$ time. If $\mathcal{A}(\mathcal{R})$ is sub-linear, then reporting the result dominates the running time of the reconstruction phase. However, we can still return a suitable data structure representing the result in $O(\mathcal{A}(\mathcal{R}))$ time. In one dimension, $\mathcal{R}$ is a set of intervals and the ambiguity is linked to interval entropy, which in turn relates to the well-studied problem of sorting under partial information. The number of comparisons necessary to find the linear order underlying a poset P is lower-bounded by the graph entropy of P. We show that if P is an interval order, then the ambiguity provides a constant-factor approximation of the graph entropy. This gives a lower bound of $\Omega(\mathcal{A}(\mathcal{R}))$ in all dimensions for the reconstruction phase (sorting or any proximity structure), independent of any preprocessing; hence our result is tight. Finally, our results imply that one can approximate the entropy of interval graphs in $O(n \log n)$ time, improving the $O(n^{2.5})$ bound by Cardinal *et al.*	**Optimal algorithm for geodesic farthest-point Voronoi diagrams** Luis Barba Let P be a simple polygon with n vertices. For any two points in P, the geodesic distance between them is the length of the shortest path that connects them among all paths contained in P. Given a set S of m sites being a subset of the vertices of P, we present the first randomized algorithm to compute the geodesic farthest-point Voronoi diagram of S in P running in expected $O(n+m)$ time. That is, a partition of P into cells, at most one cell per site, such that every point in a cell has the same farthest site with respect to the geodesic distance. This algorithm can be extended to run in expected $O(n + m \log m)$ time when S is an arbitrary set of m sites contained in P.
11:30-11:55	Break + Fast Forward	

Multimedia Sessions(The abstracts are not in order of the presentations)

11:55-12:45

Fréchet View – A Tool for Exploring Fréchet Distance Algorithms

Peter Schäfer

The **Fréchet-distance** is a similarity measure for geometric shapes. Alt and Godau presented the first algorithm for computing the Fréchet-distance and introduced a key concept, the Since then, numerous variants of the Fréchet-distance have been studied.

We present here an interactive, graphical tool for exploring some Fréchet-distance algorithms. Given two curves, users can experiment with the free-space diagram and compute the Fréchet-distance. The Fréchet-distance can be computed for two important classes of shapes: for polygonal curves in the plane, and for simple polygonal surfaces.

Finally, we demonstrate an implementation of a very recent concept, the ***k*-Fréchet-distance**.

A manual comparison of convex hull algorithms

Maarten Löffler

We have verified experimentally that there is at least one point set on which Andrew's algorithm (based on Graham's scan) to compute the convex hull of a set of points in the plane is significantly faster than a brute-force approach, thus supporting existing theoretical analysis with practical evidence. Specifically, we determined that executing Andrew's algorithm on the point set

$$P = \{(1,4),(2,8),(3,10),(4,1),(5,7),(6,3),(7,9),(8,5),(9,2),(10,6)\}$$

takes 41 minutes and 18 seconds; the brute-force approach takes 3 hours, 49 minutes, and 5 seconds.

Packing Geometric Objects with Optimal Worst-Case Density

A. T. Becker, S. P. Fekete, P. Keldenich, S. Morr, C. Scheffer

We motivate and visualize problems and methods for packing a set of objects into a given container, in particular a set of different-size circles or squares into a square or circular container. Questions of this type have attracted a considerable amount of attention and are known to be notoriously hard. We focus on a particularly simple criterion for deciding whether a set can be packed: comparing the total area A of all objects to the area C of the container. The *critical packing density* δ^* is the largest value A/C for which any set of area A can be packed into a container of area C. We describe algorithms that establish the critical density of squares in a square ($\delta^* = 0.5$), of circles in a square ($\delta^* = 0.5390\ldots$), regular octagons in a square ($\delta^* = 0.5685\ldots$), and circles in a circle ($\delta^* = 0.5$).

Properties of Minimal-Perimeter Polyominoes

G. Barequet and G. Ben-Shachar

In this video, we survey some results concerning polyominoes, which are sets of connected cells on the square lattice, and specifically, minimal-perimeter polyominoes, that are polyominoes with the minimal-perimeter from all polyominoes of the same size.

Friday, June 21

	Fri-10A: Data Structures II	Fri-10B: Graph Drawing II
9:00-9:20	**A New Lower Bound for Semigroup Orthogonal Range Searching** Peyman Afshani We report the first improvement in the space-time trade-off of lower bounds for the orthogonal range searching problem in the semigroup model, since Chazelle's result from 1990. This is one of the very fundamental problems in range searching with a long history. Previously, Andrew Yao's influential result had shown that the problem is already non-trivial in one dimension [*Space-time tradeoff for answering range queries*, STOC 1982]: using m units of space, the query time $Q(n)$ must be $\Omega(\alpha(m,n) + \frac{n}{m-n+1})$ where $\alpha(\cdot,\cdot)$ is the inverse Ackermann's function, a very slowly growing function. In d dimensions, Bernard Chazelle [*Lower bounds for orthogonal range searching: part II. the arithmetic model*, JACM 1990] proved that the query time must be $Q(n) = \Omega((\log_\beta n)^{d-1})$ where $\beta = 2m/n$. Chazelle's lower bound is known to be tight for when space consumption is "high" i.e., $m = \Omega(n \log^{d+\varepsilon} n)$. We have two main results. The first is a lower bound that shows Chazelle's lower bound was not tight for "low space": we prove that we must have $mQ(n) = \Omega(n(\log n \log\log n)^{d-1})$. Our lower bound does not close the gap to the existing data structures, however, our second result is that our analysis is tight. Thus, we believe the gap is in fact natural since lower bounds are proven for idempotent semigroups while the data structures are built for general semigroups and thus they cannot assume (and use) the properties of an idempotent semigroup. As a result, we believe to close the gap one must study lower bounds for non-idempotent semigroups or building data structures for idempotent semigroups. We develope significantly new ideas for both of our results that could be useful in pursuing either of these directions.	**Dual Circumference and Collinear Sets** V. Dujmović and P. Morin We show that, if an n-vertex triangulation T of maximum degree Δ has a dual that contains a cycle of length ℓ, then T has a non-crossing straight-line drawing in which some set, called a *collinear set*, of $\Omega(\ell/\Delta^4)$ vertices lie on a line. Using the current lower bounds on the length of longest cycles in 3-regular 3-connected graphs, this implies that every n-vertex planar graph of maximum degree Δ has a collinear set of size $\Omega(n^{0.8}/\Delta^4)$. Very recently, Dujmović *et al* (SODA 2019) showed that, if S is a collinear set in a triangulation T then, for any point set $X \subset \mathbb{R}^2$ with $\lvert X\rvert = \lvert S\rvert$, T has a non-crossing straight-line drawing in which the vertices of S are drawn on the points in X. Because of this, collinear sets have numerous applications in graph drawing and related areas.
9:20-9:40	**Independent Range Sampling, Revisited Again** Peyman Afshani and Jeff M. Phillips We revisit the range sampling problem: the input is a set of points where each point is associated with a real-valued weight. The goal is to store them in a structure such that given a query range and an integer k, we can extract k independent random samples from the points inside the query range, where the probability of sampling a point is proportional to its weight. This line of work was initiated in 2014 by Hu, Qiao, and Tao and it was later followed up by Afshani and Wei. The first line of work mostly studied unweighted but dynamic version of the problem in one dimension whereas the second result considered the static weighted problem in one dimension as well as the unweighted problem in 3D for halfspace queries. We offer three main results and some interesting insights that were missed by the previous work: We show that it is possible to build efficient data structures for range sampling queries if we allow the query time to hold in expectation (the first result), or obtain efficient worst-case query bounds by allowing the sampling probability to be approximately proportional to the weight (the second result). The third result is a conditional lower bound that shows essentially one of the previous two concessions is needed. For instance, for the 3D range sampling queries, the first two results give efficient data structures with near-linear space and polylogarithmic query time whereas the lower bound shows with near-linear space the worst-case query time must be close to $n^{2/3}$, ignoring polylogarithmic factors. Up to our knowledge, this is the first such major gap between the expected and worst-case query time of a range searching problem.	**Cubic Planar Graphs That Cannot Be Drawn On Few Lines** David Eppstein For every integer ℓ, we construct a cubic 3-vertex-connected planar bipartite graph G with $O(\ell^3)$ vertices such that there is no planar straight-line drawing of G whose vertices all lie on ℓ lines. This strengthens previous results on graphs that cannot be drawn on few lines, which constructed significantly larger maximal planar graphs. We also find apex-trees and cubic bipartite series-parallel graphs that cannot be drawn on a bounded number of lines.

9:40-10:00	**Dynamic Geometric Data Structures via Shallow Cuttings** T. M. Chan We present new results on a number of fundamental problems about dynamic geometric data structures: 1. We describe the first fully dynamic data structures with sublinear amortized update time for maintaining (i) the number of vertices or the volume of the convex hull of a 3D point set, (ii) the largest empty circle for a 2D point set, (iii) the Hausdorff distance between two 2D point sets, (iv) the discrete 1-center of a 2D point set, (v) the number of maximal (i.e., skyline) points in a 3D point set. The update times are near $n^{11/12}$ for (i) and (ii), $n^{7/8}$ for (iii) and (iv), and $n^{2/3}$ for (v). Previously, sublinear bounds were known only for restricted "semi-online" settings [Chan, SODA 2002]. 2. We slightly improve previous fully dynamic data structures for answering extreme point queries for the convex hull of a 3D point set and nearest neighbor search for a 2D point set. The query time is $O(\log^2 n)$, and the amortized update time is $O(\log^4 n)$ instead of $O(\log^5 n)$ [Chan, SODA 2006; Kaplan et al., SODA 2017]. 3. We also improve previous fully dynamic data structures for maintaining the bichromatic closest pair between two 2D point sets and the diameter of a 2D point set. The amortized update time is $O(\log^4 n)$ instead of $O(\log^7 n)$ [Eppstein 1995; Chan, SODA 2006; Kaplan et al., SODA 2017].	**Connecting the Dots (with Minimum Crossings)** Akanksha Agrawal, Grzegorz Guśpiel, Jayakrishnan Madathil, Saket Saurabh, Meirav Zehavi We study a prototype CROSSING MINIMIZATION problem, defined as follows. Let $\mathcal{F}$ be an infinite family of (possibly vertex-labeled) graphs. Then, given a set P of (possibly labeled) n points in the Euclidean plane, a collection $L \subseteq \text{Lines}(P) = \{\ell : \ell$ is a line segment with both endpoints in $P\}$, and a non-negative integer k, decide if there is a sub-collection $L' \subseteq L$ such that the graph $G = (P, L')$ is isomorphic to a graph in $\mathcal{F}$ and L' has at most k crossings. By $G = (P, L')$, we refer to the graph on vertex set P, where two vertices are adjacent if and only if there is a line segment that connects them in L'. Intuitively, in CROSSING MINIMIZATION, we have a set of locations of interest, and we want to build/draw/exhibit connections between them (where L indicates where it is feasible to have these connections) so that we obtain a structure in $\mathcal{F}$. Natural choices for $\mathcal{F}$ are the collections of perfect matchings, Hamiltonian paths, and graphs that contain an (s, t)-path (a path whose endpoints are labeled). While the objective of seeking a solution with few crossings is of interest from a theoretical point of view, it is also well motivated by a wide range of practical considerations. For example, links/roads (such as highways) may be cheaper to build and faster to traverse, and signals/moving objects would collide/interrupt each other less often. Further, graphs with fewer crossings are preferred for graphic user interfaces. As a starting point for a systematic study, we consider a special case of CROSSING MINIMIZATION. Already for this case, we obtain NP-hardness and W[1]-hardness results, and ETH-based lower bounds. Specifically, suppose that the input also contains a collection D of d non-crossing line segments such that each point in P belongs to exactly one line in D, and L does not contain line segments between points on the same line in D. Clearly, CROSSING MINIMIZATION is the case where $d = n$—then, P is in general position. The case of $d = 2$ is of interest not only because it is the most restricted non-trivial case, but also since it corresponds to a class of graphs that has been well studied—specifically, it is CROSSING MINIMIZATION where $G = (P, L)$ is a (bipartite) graph with a so called *two-layer drawing*. For $d = 2$, we consider three basic choices of $\mathcal{F}$. For perfect matchings, we show *(i)* NP-hardness with an ETH-based lower bound, *(ii)* solvability in subexponential parameterized time, and *(iii)* existence of an $O(k^2)$-vertex kernel. Second, for Hamiltonian paths, we show *(i)* solvability in subexponential parameterized time, and *(ii)* existence of an $O(k^2)$-vertex kernel. Lastly, for graphs that contain an (s, t)-path, we show *(i)* NP-hardness and W[1]-hardness, and *(ii)* membership in XP.
10:00-10:30	Coffee Break	
	Fri-11A: Complexity	**Fri-11B: Combinatorial Geometry III**
10:30-10:50	**The Unbearable Hardness of Unknotting** A. de Mesmay, Y. Rieck, E. Sedgwick, M. Tancer We prove that deciding if a diagram of the unknot can be untangled using at most k Reidemeister moves (where k is part of the input) is **NP**-hard. We also prove that several natural questions regarding links in the 3-sphere are **NP**-hard, including detecting whether a link contains a trivial sublink with n components, computing the unlinking number of a link, and computing a variety of link invariants related to four-dimensional topology (such as the 4-ball Euler characteristic, the slicing number, and the 4-dimensional clasp number).	**An Experimental Study of Forbidden Patterns in Geometric Permutations by Combinatorial Lifting** Goaoc X., Holmsen A., and Nicaud C. We study the problem of deciding if a given triple of permutations can be realized as geometric permutations of disjoint convex sets in $\mathbb{R}^3$. We show that this question, which is equivalent to deciding the emptiness of certain semi-algebraic sets bounded by cubic polynomials, can be "lifted" to a purely combinatorial problem. We propose an effective algorithm for that problem, and use it to gain new insights into the structure of geometric permutations.

10:50-11:10	**Circumscribing Polygons and Polygonizations for Disjoint Line Segments** H. A. Akitaya, M. Korman, M. Rudoy, C. D. Tóth, and D. L. Souvaine Given a planar straight-line graph $G = (V, E)$ in $\mathbb{R}^2$, a *circumscribing polygon* of G is a simple polygon P whose vertex set is V, and every edge in E is either an edge or an internal diagonal of P. A circumscribing polygon is a *polygonization* for G if every edge in E is an edge of P. We prove that every arrangement of n disjoint line segments in the plane has a subset of size $\Omega(\sqrt{n})$ that admits a circumscribing polygon, which is the first improvement on this bound in 20 years. We explore relations between circumscribing polygons and other problems in combinatorial geometry, and generalizations to $\mathbb{R}^3$. We show that it is NP-complete to decide whether a given graph G admits a circumscribing polygon, even if G is 2-regular. Settling a 30-year old conjecture by Rappaport, we also show that it is NP-complete to determine whether a geometric matching admits a polygonization.	**A Product Inequality for Extreme Distances** Adrian Dumitrescu Let $p_1, \ldots, p_n$ be n distinct points in the plane, and assume that the minimum inter-point distance occurs $s_{\min}$ times, while the maximum inter-point distance occurs $s_{\max}$ times. It is shown that $s_{\min} s_{\max} \leq \frac{9}{8} n^2 + O(n)$; this settles a conjecture of Erdős and Pach (1990).
11:10-11:30	**Counting Polygon Triangulations is Hard** David Eppstein We prove that it is #P-complete to count the triangulations of a (non-simple) polygon.	**Convex Polygons in Cartesian Products** J.-L. De Carufel, A. Dumitrescu, W. Meulemans, T. Ophelders, C. Pennarun, C. D. Tóth, and S. Verdonschot We study several problems concerning convex polygons whose vertices lie in a Cartesian product of two sets of n real numbers (for short, *grid*). First, we prove that every such grid contains a convex polygon with $\Omega(\log n)$ vertices and that this bound is tight up to a constant factor. We generalize this result to d dimensions (for a fixed $d \in \mathbb{N}$), and obtain a tight lower bound of $\Omega(\log^{d-1} n)$ for the maximum number of points in convex position in a d-dimensional grid. Second, we present polynomial-time algorithms for computing the longest convex polygonal chain in a grid that contains no two points with the same x- or y-coordinate. We show that the maximum size of such a convex polygon can be efficiently approximated up to a factor of 2. Finally, we present exponential bounds on the maximum number of convex polygons in these grids, and for some restricted variants. These bounds are tight up to polynomial factors.
11:30-11:40	Break	

See next page ...

11:40-12:40

Invited Talk: **Some geometric and computational challenges arising in structural molecular biology**

Bruce R. Donald

Abstract: Computational protein design is a transformative field with exciting prospects for advancing both basic science and translational medical research. New algorithms blend discrete and continuous geometry to address the challenges of creating designer proteins. I will discuss recent progress in this area and some interesting open problems.

I will motivate this talk by discussing how, by using continuous geometric representations within a discrete optimization framework, broadly-neutralizing anti-HIV-1 antibodies were computationally designed that are now being tested in humans – the designed antibodies are currently in eight clinical trials (See `https://clinicaltrials.gov/ct2/results?cond=&term=vrc07&cntry=&state=&city=&dist=`), one of which is Phase 2a (NCT03721510). These continuous representations model the flexibility and dynamics of biological macromolecules, which are an important structural determinant of function.

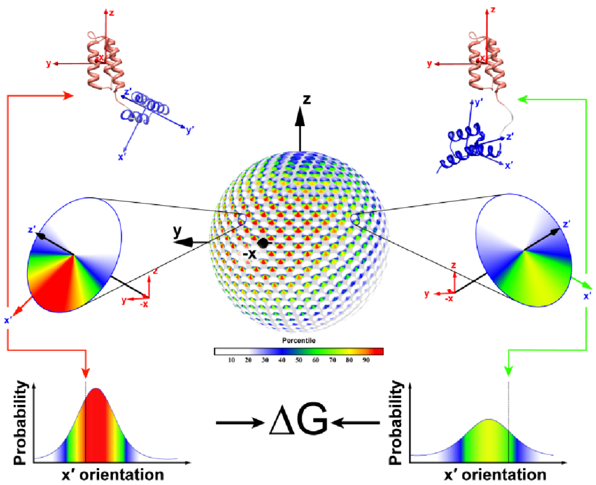

However, reconstruction of biomolecular dynamics from experimental observables requires the determination of a conformational probability distribution. These distributions are not fully constrained by the limited geometric information from experiments, making the problem ill-posed in the sense of Hadamard. The ill-posed nature of the problem comes from the fact that it has no unique solution. Multiple or even an infinite number of solutions may exist. To avoid the ill-posed nature, the problem must be regularized by making (hopefully reasonable) assumptions.

I will present new ways to both represent and visualize correlated inter-domain protein motions (See Figure). We use Bingham distributions, based on a quaternion fit to circular moments of a physics-based quadratic form. To find the optimal solution for the distribution, we designed an efficient, provable branch-and-bound algorithm that exploits the structure of analytical solutions to the trigonometric moment problem. Hence, continuous conformational PDFs can be determined directly from NMR measurements. The representation works especially well for multi-domain systems with broad conformational distributions.

Ultimately, this method has parallels to other branches of geometric computing that balance discrete and continuous representations, including physical geometric algorithms, robotics, computational geometry, and robust optimization. I will advocate for using continuous distributions for protein modeling, and describe future work and open problems.

Department of Computer Science, Department of Chemistry, and Department of Biochemistry Duke University, and Duke University Medical Center (`www.cs.duke.edu/brd/`)