Open Problems from CCCG 2025

Hugo Akitaya* Shahin Kamali† Seyed-Mohammad Seyed-Javadi‡ Frederick Stock§

The following is a list of the problems presented during the open problem session at the 37th Annual Canadian Conference on Computational Geometry (CCCG), held at York University in Toronto from August 13 to 15, 2025. This year's open problem session took place on the second day of the conference and featured lively participation from attendees. As in past years, the session served as a catalyst for spontaneous collaboration and discussion throughout the remainder of the meeting.

Open problem sessions have long been a tradition at CCCG, offering a venue for researchers to share intriguing questions, conjectures, and partially explored ideas. The 2025 session continued this tradition, with problems spanning classical themes such as Minkowski-type existence results and modern geometric inequalities.

To the best of our knowledge, the problems below accurately reflect the questions posed during the session. The problems are listed in the order in which they were presented at the conference. Where appropriate, references and attributions to presenters have been included.

1 Logarithmic Minkowski Problem

presented by Joseph O'Rourke

Smith College (jorourke@smith.edu)

O'Rourke posed a problem that was recently presented in the AMS Bulletin [Huang et al. (2025)]. Minkowski-type problems ask when a geometric object can be reconstructed from prescribed geometric data—such as the normals of its faces and measures associated with them (areas, lengths, or other quantities). These questions trace back to Hermann Minkowski's foundational results in convex geometry [Minkowski (1897)] and have inspired many modern extensions. Minkowski showed that a convex polyhedron is uniquely determined (up to translation) by the normals and areas of its faces. Constructive methods for explicitly finding such a polyhedron were later developed by Little [Little (1985)] and Sellaroli [Sellaroli (2017)].

*University of Massachusetts Lowell, USA, hugo_akitaya@uml.edu

†York University, Toronto, Canada, kamalis@yorku.ca

‡York University, Toronto, Canada, mohammad.sj2@gmail.com

§University of Massachusetts Lowell, USA, frederick_stock@student.uml.edu

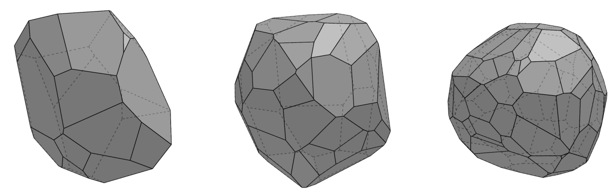

Figure 1: Examples of reconstructed polyhedra with $n = 25, 50, 100$. Figure is taken from [Sellaroli (2017)].

Theorem 1 (Minkowski's Theorem for Polyhedra) *[Minkowski (1897), Little (1985), Sellaroli (2017)] Given unit face normals $u_1, \ldots, u_n$ and corresponding face areas $A_1, \ldots, A_n$, there exists a unique convex polyhedron (up to translation) whose faces have those normals and areas.*

An illustration is provided in Figure 1.

The following theorem presents the two-dimensional version of Minkowski's theorem.

Theorem 2 (Minkowski's Theorem for Polygons) *[Huang et al. (2025)] Let $\ell_1, \ldots, \ell_m$ be positive numbers and $u_1, \ldots, u_m$ be unit vectors in $\mathbb{R}^2$. Then there exists a convex polygon P in $\mathbb{R}^2$ with side lengths $\ell_1, \ldots, \ell_m$ and corresponding outer unit normals $u_1, \ldots, u_m$ if and only if $\ell_1 u_1 + \ldots + \ell_m u_m = 0$ and the unit vectors $u_1, \ldots, u_m$ do not lie in a unit half-circle. Additionally, P is unique up to translation.*

The *logarithmic Minkowski problem for polygons*, introduced by [Huang et al. (2025)], is a variant of the classical Minkowski problem in which the prescribed data are triangle areas rather than edge lengths.

Question 1. (The logarithmic Minkowski problem for polygons) [Huang et al. (2025)] Given unit face normals $u_1, \ldots, u_n$ and triangle areas $A_1, \ldots, A_n$, construct a convex m-gon P in $\mathbb{R}^2$ containing the origin $\mathbf{0}$ such that, for each side s_i of P, u_i is its unit vector and the area of the triangle corresponding to the convex hull of $\{\mathbf{0}\} \cup s_i$ is A_i.

Figure 2 provides an illustration.

Even in this planar, discrete setting, the problem remains unsolved in full generality and is known to be substantially more difficult than the classical Minkowski problem.

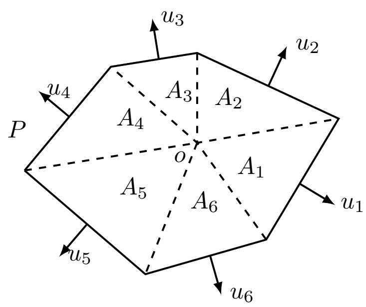

Figure 2: An illustration of the logarithmic Minkowski problem for polygons. Figure is taken from [Huang et al. (2025)].

Symmetric measures can lead to obstructions in the logarithmic Minkowski problem. For instance, if the prescribed unit normals satisfy $u_1 = -u_3$ and $u_2 = -u_4$, and the corresponding areas satisfy $A_1 = A_3$ and $A_2 = A_4$, then any solution must be a parallelogram centered at the origin. This forces $A_1 = A_2$; hence, if $A_1 \neq A_2$, no solution exists (See Figure 3). The challenge is to determine, given the u_i and A_i, when there is and when there is not a solution.

References

[Huang et al. (2025)] Yong Huang, Deane Yang, and Gaoyong Zhang. Minkowski problems for geometric measures. *Bulletin of the American Mathematical Society*, 62(3):359–425, 2025.

[Little (1985)] James J. Little. Extended gaussian images, mixed volumes, shape reconstruction. In *Proc. the First Annual Symposium on Computational Geometry (SOCG)*, pages 15–23. ACM, 1985.

[Minkowski (1897)] Hermann Minkowski. Allgemeine lehrsätze über die convexen polyeder. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, pages 198–219, 1897.

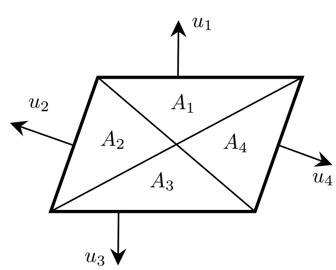

Figure 3: An obstruction in the logarithmic Minkowski problem

[Sellaroli (2017)] Giuseppe Sellaroli. An algorithm to reconstruct convex polyhedra from their face normals and areas. *CoRR*, abs/1712.00825, 2017.

2 Sibley's Convexity Measures

presented by Joseph O'Rourke

Smith College (jorourke@smith.edu)

O'Rourke posed a question recently presented in Mathematics Magazine by Thomas Sibley [Sibley (2025)]. While previous Minkowski problems focus on reconstructing convex bodies from prescribed geometric data, such as normals and areas, Sibley's conjectures are about quantitative relationships among geometric measures.

Sibley proposed two conjectures relating geometric measures that capture how far a polygon deviates from convexity and how much of it remains "visible" from within. For a simple polygon P, let kernel(P) denote the set of points from which every point of P is visible—equivalently, the intersection of the interior half-planes defined by its edges. The area of kernel(P) quantifies the region of complete visibility within P.

For a polygon P, define the following normalized quantities:

$$G(P) = \frac{\text{area}(\text{kernel}(P))}{\text{area}(P)}, \quad A(P) = \frac{\text{area}(P)}{\text{area}(\text{conv}(P))},$$

$$L(P) = \frac{\text{perimeter}(\text{conv}(P))}{\text{perimeter}(P)}.$$

Here, conv(P) denotes the convex hull of P. Intuitively, $G(P)$ measures the fraction of the polygon that is fully visible, while $A(P)$ and $L(P)$ quantify how much of the polygon's area and perimeter are "lost" due to non-convexity.

Sibley made the following conjectures, which compare the visible portion of a polygon to its geometric deviation from convexity:

Question 2 (Thomas Sibley's Conjectures). [Sibley (2025)] Do the following inequalities hold for all simple polygons P?

A: $G(P) \leq A(P)$

L: $G(P) \leq L(P)$.

Note, this question is trivial for any non-star-shaped polygon P, as the area of P's kernel will be zero.

References

[Sibley (2025)] Thomas Sibley. Comparing convexity measures. *Mathematics Magazine*, 98(3):162–174, 2025.

3 Separating Convex Sets

presented by Shakhar Smorodinsky
Ben-Gurion University (shakhar@math.bgu.ac.il)

Note: A very recent result [Chen et al. (2025)] shows that this problem has now been resolved. We include it here for historical context.

Separating families of points in the plane is a fundamental concept in computational geometry (and other domains such as machine learning). Let X be a finite set. A family $\mathcal{F}$ of subsets of X is called a *separating system* if, for every pair of distinct elements $a, b \in X$, there exists a set $F \in \mathcal{F}$ that contains exactly one of a and b [Biniaz et al. (2025)]. Equivalently, every pair of elements in X is distinguished by at least one member of $\mathcal{F}$. The smallest possible size of such a family is known as the *separation number* of X.

For instance, when X is a set of n points in general position in the plane and $\mathcal{F}$ consists of halfplanes defined by lines, it is known that $\lceil n/2 \rceil$ lines are always sufficient and sometimes necessary [Boland and Urrutia(1995)]. When $\mathcal{F}$ consists of disks instead of lines, and no four points are cocyclic, the tight bound becomes $2\lceil n/6 \rceil + 1$ [Gledel and Parreau (2019)].

The problem considered here extends this framework to two sets of convex sets, $A = \{A_1, \ldots, A_n\}$ and $B = \{B_1, \ldots, B_n\}$, where $\mathcal{F}$ is formed by lines and the separation condition applies only to pairs (A_i, B_j).

Question 3 *Let $\mathcal{A} = \{A_1, \ldots, A_n\}$ and $\mathcal{B} = \{B_1, \ldots, B_n\}$ be two collections of convex sets in the plane such that for every $i, j \in [n]$, the sets A_i and B_j are disjoint. Define $f(n)$ to be the smallest integer such that, for any such pair of collections $\mathcal{A}$ and $\mathcal{B}$, there exists a set L of at most $f(n)$ lines with the property that for every $i, j \in [n]$, some line $\ell \in L$ separates A_i and B_j.*
Determine (or bound) the value of $f(n)$.

It is known that if, in addition, the family $A \cup B$ consists of pairwise disjoint sets, then $f(n) = O(n)$. This problem arises in the context of the recent work in [Alon and Smorodinsky (2026)], which relates to an extension of the classical Radon's and Tverberg's theorems in discrete geometry.

References

[Alon and Smorodinsky (2026)] Noga Alon and Shakhar Smorodinsky. Extended VC-dimension, and Radon and Tverberg type theorems for unions of convex sets. In *Proc. Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (to appear), 2026.

[Biniaz et al. (2025)] Ahmad Biniaz, Prosenjit Bose, Jean-Lou De Carufel, Anil Maheshwari, Babak Miraftab, Saeed Odak, Michiel Smid, Shakhar Smorodinsky, and Yelena Yuditsky. On separating path and tree systems in graphs. *Discret. Math. Theor. Comput. Sci.*, 27(2), 2025.

[Boland and Urrutia(1995)] Ralph P. Boland and Jorge Urrutia. Separating collections of points in Euclidean spaces. *Inf. Process. Lett.*, 53(4):177–183, 1995.

[Chen et al. (2025)] Wenchong Chen, Gennian Ge, Yang Shu, Zhouningxin Wang, and Zixiang Xu. A Tverberg-type problem of Kalai: Two negative answers to questions of alon and smorodinsky, and the power of disjointness, November 2025.

[Gledel and Parreau (2019)] Valentin Gledel and Aline Parreau. Identification of points using disks. *Discret. Math.*, 342(1):256–269, 2019.

4 Geometric Graph Covering

presented by Takumi Shiota
University of Hyogo (takumi_shiota@gsis.u-hyogo.ac.jp)

Covering problems in geometric unit-grid graphs arise frequently in combinatorial geometry and discrete tilings. Let G denote a unit grid graph whose vertices are a subset of the integer lattice points in a rectangular region of $\mathbb{Z}^2$ and whose edges are the unit horizontal and vertical grid segments. In this problem, we ask whether the entire edge set of G can be expressed as the union of the boundary edges of a small number of 2×2 tiles.

Question 4 *Let G be a unit grid graph. Given an integer k, determine whether there exist k (possibly overlapping) induced 2×2* blocks*—each block being the cycle C_8 induced by the vertices on the boundary of a 2×2 square—such that the union of the edges of these blocks is* exactly *the full edge set $E(G)$.*

Figure 4 illustrates an example.

This problem can be solved in polynomial time if 2×2 blocks are unlimited in number, because it is sufficient to check whether they can be placed starting from the top-left. However, when the constraint of exactly k blocks is imposed, it is not clear whether this problem can be solved in polynomial time or if it is NP-complete.

Covering all edges as a problem setting has not been extensively studied in existing research. A related problem is the Border Drawing Problem (BDP) in Excel [Iwama et al. (2009)], which has been shown to be NP-hard.

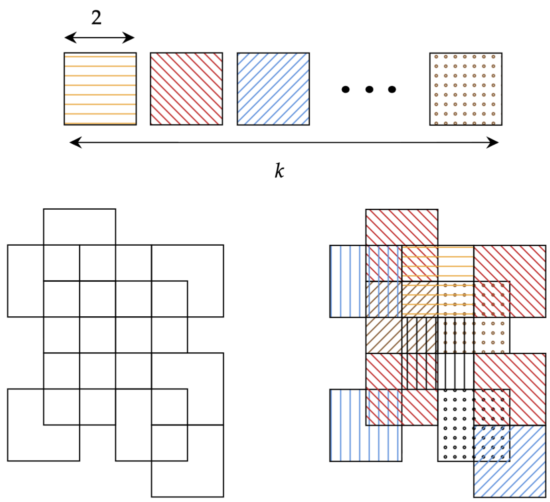

Figure 4: (Top) A set of $k = 12$ blocks of size 2×2, (Bottom left) A unit grid graph G, (Bottom right) A covering of G using the k blocks.

References

[Iwama et al. (2009)] Kazuo Iwama, Eiji Miyano, and Hirotaka Ono. Drawing borders efficiently. *Theory of Computing Systems*, 44:230–244, 2009. doi: 10.1007/s00224-008-9117-y.

5 Shapes of Vertex Sets of Planar Graphs

presented by David Epstein

University of California, Irvine (eppstein@uci.edu)

A recurring theme in graph drawing is understanding how restricted a point set can be while still supporting crossing-free straight-line drawings of various graph classes. Existing results include:

- If all vertices are required to lie in convex position, then the graphs that admit such drawings are precisely the outerplanar graphs. [folklore]
- Restricting the vertex set to lie on a constant number of parallel lines does not accommodate all planar graphs. In particular, the graphs that can be drawn with vertices on k parallel lines have pathwidth $\leq k-1$. Lemma 1 in [Dujmovic et al. (2008)].
- Even allowing a constant number of lines in arbitrary orientations does not suffice to represent all planar graphs [Chaplick et al. (2020)], that is, for every fixed k, there exists a planar graph that cannot be drawn with its vertices on k lines (not assumed to be parallel). More strongly, drawing an n-vertex planar graph may require $\Omega(k^{1/3})$ lines. This was shown in [Eppstein (2021)]. Its references provide earlier but weaker bounds on the required number of lines and on the NP-hardness of finding a drawing on two crossing lines.

In short, several natural geometric constraints on vertex placement are too restrictive to support all planar graphs. See [Bannister et al. (2019), Kryven et al. (2019), Eppstein (2021), Chaplick et al. (2023)] for some related work on limitations arising from constrained vertex locations.

The following question asks whether this limitation persists when the vertex set is allowed to lie on a fixed number of more flexible geometric supports.

Question 5 *Does there exist a constant k such that every planar graph admits a straight-line plane drawing in which all vertices lie on the union of k pairwise disjoint convex curves?*

This is closely related to open problem 16.14 from [Eppstein (2018)] which asks whether, for each n, there is a family of $O(1)$ convex curves that can support all n-vertex planar graphs. This question is also the motivation for Eppstein's upcoming paper [Eppstein (2025)].

References

[Bannister et al. (2019)] Michael J. Bannister, William E. Devanny, Vida Dujmovic, David Eppstein, and David R. Wood. Track layouts, layered path decompositions, and leveled planarity. *Algorithmica*, 81(4):1561–1583, 2019.

[Chaplick et al. (2020)] Steven Chaplick, Krzysztof Fleszar, Fabian Lipp, Alexander Ravsky, Oleg Verbitsky, and Alexander Wolff. Drawing graphs on few lines and few planes. *J. Comput. Geom.*, 11 (1):433–475, 2020.

[Chaplick et al. (2023)] Steven Chaplick, Krzysztof Fleszar, Fabian Lipp, Alexander Ravsky, Oleg Verbitsky, and Alexander Wolff. The complexity of drawing graphs on few lines and few planes. *J. Graph Algorithms Appl.*, 27(6):459–488, 2023.

[Dujmovic et al. (2008)] Vida Dujmovic, Michael R. Fellows, Matthew Kitching, Giuseppe Liotta, Catherine McCartin, Naomi Nishimura, Prabhakar Ragde, Frances A. Rosamond, Sue Whitesides, and David R. Wood. On the parameterized complexity of layered graph drawing. *Algorithmica*, 52(2):267–292, 2008.

[Eppstein (2018)] David Eppstein. *Forbidden Configurations in Discrete Geometry.* Cambridge University Press, Cambridge, 2018.

[Eppstein (2021)] David Eppstein. Cubic planar graphs that cannot be drawn on few lines. *Journal of Computational Geometry*, 12(1):178–197, 2021. doi: 10.20382/v12i1a8.

[Eppstein (2025)] David Eppstein. Stabbing faces by a convex curve. In *Proc. the 33rd Int. Symp. Graph Drawing and Network Visualization*, volume 357, pages 29:1–29:8, 2025.

[Kryven et al. (2019)] Myroslav Kryven, Alexander Ravsky, and Alexander Wolff. Drawing graphs on few circles and few spheres. *J. Graph Algorithms Appl.*, 23(2):371–391, 2019.

6 Spanning Trees of Maps

presented by Arash Rafiey

Indiana State University (Arash.Rafiey@indstate.edu)

Consider a rectangular region subdivided by a finite collection of line segments into a set of connected subregions; we refer to this subdivision as a *map*. Each region contains a finite (possibly empty) set of designated points, called *special points*. Let m denote such a map. We define a graph G_m whose vertex set consists of all special points in all regions. There is an edge between two special points whenever the regions containing them share a boundary point (i.e., the regions are adjacent). Each edge uv is assigned a weight equal to the Euclidean distance between the corresponding special points.

A *representative spanning tree* of G_m is a tree that selects one special point from each non-empty region and includes only those selected vertices as its vertex set. Equivalently, the tree induces a connected selection of one representative per region, with edges corresponding to region adjacencies. See Figure 6 for an illustration.

Question 6 *Given a map m and its associated graph G_m, find a representative spanning tree of minimum total weight. What is the computational complexity of this problem?*

This problem is related to the Group Steiner Tree problem, where one is given a collection of groups of vertices and seeks a minimum-weight subtree that contains at least one vertex from each group. The Group Steiner Tree problem is NP-hard, even when the underlying graph is a tree [Garg et al.(2000)], and it is hard to approximate within polylogarithmic factors [Halperin and Krauthgamer(2003)]. In contrast, the problem considered here requires selecting exactly one vertex from each region, and the objective is to compute a minimum-weight *spanning* tree on the selected representatives, rather than a Steiner tree.

Another related line of work considers minimum-weight spanning trees under *neighborhood constraints*, where a collection of geometric neighborhoods (e.g., disjoint disks) is given, and one representative point must be chosen from each neighborhood before computing the tree (see, e.g., [Dorrigiv et al. (2015)]). The setting in Question 6 differs in that the representative points are predetermined rather than chosen arbitrarily, and the underlying structure of the map is fixed by its line-segment subdivision.

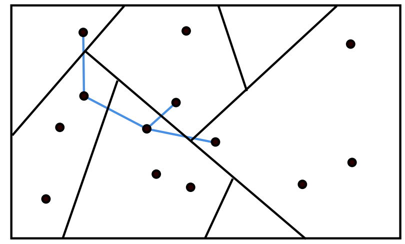

Figure 5: A map and its associated minimum-weight representative spanning tree

When the regions are unit squares, the problem is known to have a PTAS [Mirzanezhad and Rafiey (2025)].

References

[Dorrigiv et al. (2015)] Reza Dorrigiv, Robert Fraser, Meng He, Shahin Kamali, Akitoshi Kawamura, Alejandro López-Ortiz, and Diego Seco. On minimum- and maximum-weight minimum spanning trees with neighborhoods. *Theory Comput. Syst.*, 56(1):220–250, 2015.

[Garg et al.(2000)] Naveen Garg, Goran Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group steiner tree problem. *Journal of Algorithms*, 37(1):66–84, 2000.

[Halperin and Krauthgamer(2003)] Eran Halperin and Robert Krauthgamer. Polylogarithmic inapproximability. In *Proc. the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 585–594, 2003.

[Mirzanezhad and Rafiey (2025)] Majid Mirzanezhad and Arash Rafiey. Approximate and Exact Geometric Generalized Minimum Spanning Trees In *Proc. of The 37th Canadian Conference on Computational Geometry (CCCG)*, 2025

7 Bicoloured Point-Separation Rounding

presented by Jack Spalding-Jamieson

University of Waterloo (jacksj@uwaterloo.ca)

Separating red and blue point sets in the plane using geometric obstacles is an emerging theme in computational geometry. Recent work separating two points with weighted obstacles [Spalding-Jamieson and Naredla (2025)] and on approximation algorithms for the unweighted variant [Lynch and Spalding-Jamieson (2025)]. The problem below asks for a "sparsification" of an already separating family of weighted lines—that is, for a small subset that preserves a required minimum separation thickness between red and blue points.

Question 7 *You are given a set of n red points R and a set of n blue points B, and a set of n line segments C. Each segment s in C has an associated weight $0 \leq w_s \leq 1$, with a guarantee that for any path from a red point in R to a blue point in B (a "bichromatic path") will cross segments with total weight at least 1.*

Let $W := \sum_{s \in C} w_s$. Can you output a subset $S \subseteq C$ that blocks (is crossed by) all bichromatic paths, so that $|S| \in O(W\,polylog(n))$? What about $|S| \in O(W)$?

Figure 6 provides an illustration.

A positive (and algorithmic) answer would form a rounding step for an approximation algorithm solving a special case of the "generalized point-separation problem" [Kumar et al. (2022)]. This is possible with standard rounding methods in some limited cases, for example, when the maximum degree of the geometric intersection graph of C is constant, in which case $|S| = O(W)$ suffices.

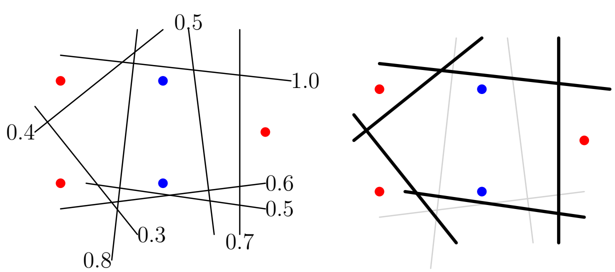

Figure 6: An input for the bicoloured point-separation rounding problem and one possible selection S

References

[Lynch and Spalding-Jamieson (2025)] Jayson Lynch and Jack Spalding-Jamieson. Subquadratic approximation algorithms for separating two points with objects in the plane. *arXiv:2507.22293*, 2025.

[Spalding-Jamieson and Naredla (2025)] Jack Spalding-Jamieson and Anurag Murty Naredla. Separating two points with obstacles in the plane: Improved upper and lower bounds. *Journal of Computational Geometry*, 2025. To appear.

[Kumar et al. (2022)] Neeraj Kumar, Daniel Lokshtanov, Saket Saurabh, Subhash Suri, and Jie Xue. Point Separation and Obstacle Removal by Finding and Hitting Odd Cycles *Proc. the 38th International Symposium on Computational Geometry (SoCG 2022)*, 2022.