WAFR 2026 Open Problems Session

Session chair: **Dan Halperin** Scribe: **Hannah Erickson**

June 17, 2026

1 Mini Introduction: Open Problems in the Age of AI

References (here and throughout the summary, click on the hyperlink to get to the source):

- G. Barber, "They Spent Years on a Math Problem. Then They Were Scooped by A.I.," The New York Times, June 8, 2026.
- N. Alon, T. F. Bloom, W. T. Gowers, D. Litt, W. Sawin, A. Shankar, J. Tsimerman, V. Wang, and M. M. Wood, "Remarks on the Disproof of the Unit Distance Conjecture," arXiv preprint arXiv:2605.20695, May 2026.
- N. Alon, Erdős Problems and Speculations about the Power of AI Models, invited lecture, Technion – Israel Institute of Technology, Haifa, Israel, June 9, 2026. YouTube video. [Online]. Available: https://www.youtube.com/watch?v=KbNctTQnVHI

2 Minimizing Makespan in Multi-Robot Coordination

Posed by Dan Halperin

Consider unit disc robots translating in the plane. Each has to move from a start position to a goal position. At any time, each robot is either standing still or moving at unit speed.

Question: Give two of the above-described unit disc robots, what is the shortest time plan (makespan) to get the robots from their start points to their goal points without the robots colliding with each other?

Background material can be found in Section 1 of the survey

- M. Abrahamsen and D. Halperin, "Ten Problems in Geobotics" arXiv preprint arXiv:2408.12657, Aug. 2024. [Online]. Available: https://arxiv.org/abs/2408.12657

In Section 5.4 of their paper

- E. D. Demaine, S. P. Fekete, P. Keldenich, H. Meijer, and C. Scheffer, "Coordinated Motion Planning: Reconfiguring a Swarm of Labeled Robots with Bounded Stretch," SIAM Journal on Computing, vol. 48, no. 6, pp. 1727–1762, 2019, doi: 10.1137/18M1194341.

Demaine et al. studied a particular instance of the problem and assumed that one of the robots moves along a circular arc. They stated that they do not have a proof that this is the optimal solution. Indeed, it has recently been shown that the optimal solution in that case does not contain motion along a circular arc [Jack Stade, personal communication, 2026]. This may hint that the solution to this problem is more complex than the solution to the same question when the optimization criterion is the minimum total length of the paths of the two robots.

3 Continuity and Computability of Minimum Pseudometrics

Posed by Vadim Weinstein

A *pseudometric* over a set X is a function $d\colon X \times X \to \mathbb{R}_{\geq 0}$ that is symmetric and obeys the triangle inequality. However, unlike a metric $d(x, y)$ can equal 0 even when x and y are different members of X.

Consider the comparison relationship over two pseudometrics which is defined by $d_1 \leq d_2$ if and only if $\forall\, x, y \ \big(d_1(x, y) \leq d_2(x, y)\big)$.

Fix a function $f\colon X \to X$. Let d be a pseudometric on X. We say that f is *Lipschitz for d* if $d(f(x), f(y)) \leq L d(x, y)$ for some real $L \geq 0$.

It is already known that, given a bounded pseudometric d and any function f as above, there exists a unique $\leq$-smallest pseudometric $d_1 \geq d$ such that f is Lipschitz for d_1. Since d_1 is uniquely determined from d and f, we can denote $d_1 = \mathrm{MSR}_f(d)$ (minimal sufficient refinement, as this generalizes the notion).

First question (Concrete): Let $X = \mathbb{R}$, $f(x) = x+1$, and let $h\colon \mathbb{R} \to \mathbb{R}_{\geq 0}$ be continuous and bounded. Let d_h be the pseudometric defined by $d_h(x, y) = |h(x) - h(y)|$. What is $\mathrm{MSR}_f(d_h)$? Is there a general way to obtain it? Even more concretely, if h is given by

$$h(x) = \max\{1 - |x|, 0\} \ .$$

What is $\mathrm{MSR}(d_h)$?

Second question (Theoretical): What do we know about the regularity or descriptive complexity of $\mathrm{MSR}_f(d)$ relative to d and f? For example, if f is continuous w.r.t. d, is $\mathrm{MSR}_f(d)$ necessarily continuous w.r.t. $d \times d$? Replace "continuous" with "piecewise continuous", "measurable", "Borel", etc. What is true?

Third question (Approximation): Can we approximate $\mathrm{MSR}_f(d)$ given f and d? Given that the answer to the second question is reasonable and we know that $\mathrm{MSR}_f(d)$ is well-behaved, then how well can we computationally approximate it given that we can query arbitrary values of f and d? A special case is when X is finite. Is the map $(d, f) \mapsto \mathrm{MSR}_f(d)$ computable in this case?

4 Part Construction Problem

Posed by Dylan Shell

Consider a *motion planning problem* consisting of a starting point (in the workspace), a goal region (subset of the workspace), and a blocked region (subset of the workspace).

Given a rigid shape and a motion planning problem, the shape *satisfies* the problem if there exists a starting configuration (in which the shape intersects the problem's starting point) and an ending configuration (in which the shape intersects the problem's goal region), and a path through the configuration space from the start to the goal that does not cause the shape to intersect the problem's blocked region.

Question (Part Construction Problem): Given a set of positive and negative motion planning problems, determine a shape that satisfies all of the positive motion planning problems but none of the negative motion planning problems.

Informally, the shape should get "stuck" in the environments corresponding to the negative problems but should "pass through" the environments corresponding to the positive problems.

For example—the following, a 2D version of the puzzle, was posed in The Guardian as "Can you solve it? An object that defies common sense". It is counter-intuitive because the negative example has trap doors that are wide open, while the positive example has the doors drawn in more narrowly. A solution exists for that puzzle, and the shape is not unique.

The problem of part construction explains something that I observed and was curious about. There is a class of ship anchor called an "Admiralty Anchor" where there's an L-shaped part called the stock. This anchor might be interpreted as the solution to a part construction problem.

In this case—a ship anchor ought to be unwieldy in order to stick on the seafloor, but also convenient to stow once it has been raised onboard. Here, one way to formulate the design problem is as an articulated structure that has multiple configurations. (In the anchor, there is a prismatic joint which allows a rotation.) This is akin to the dual problem of the 2D puzzle above: the articulated structure is the part being designed, not the trapdoor.

Related questions: (i) How would parts be efficiently parameterized or described? (ii) What is the relationship to the minimum constraint removal problem? (iii) There are also questions related to 'unwieldiness' which are not entirely geometric, but depend on physics. For instance, cereal jamming when poured out of a box with a certain sized opening.

5 POMDP Backward Matrix

Posed by Seiji Shaw

For a POMDP with state set S, the *backward matrix* of the POMDP is a $|S| \times \infty$ matrix that represents, for each state, the likelihood of each possible future action-observation sequence starting from that state. This paper of Shaw et al., https://arxiv.org/abs/2601.18930 describes the construction and probabilistic interpretation of the backward matrix.

Question (Rank of deterministic backward matrix): Consider the case where the POMDP backward matrix is all 1s and 0s (deterministic). When this POMDP backward matrix is minimized by elimination of identical rows, is the resulting matrix full rank?

Dylan Shell added: The problem of compressing a finite state–like machine representing a policy (like, e.g., Eric Hansen's "Solving POMDPs by Searching Policy Space") may be treated as a type of combinatorial filter (inputs are the observations, outputs are the actions). The paper Y. Zhang and D. A. Shell, "Cover Combinatorial Filters and their Minimization Problem", in WAFR-XIV, June 2020 shows an important fact: finite automata minimization techniques like that of Myhill–Nerode, or quotienting operations, will not allow the set of input strings to grow even though that permits additional reduction.

6 Complexity of Motion Planning with Clearance

Posed by Kiril Solovey

Piano mover problems, when the number of degrees of freedom is part of the input, are in general PSPACE-hard. Exact combinatorial representations of free space tend to be very large relative to the complexity of the obstacles. Determining them is only practical for a small number of degrees of freedom. By contrast, sampling based planning captures structure of free space when clearance is greater than 0, but not without clearance.

Question (Piano Movers with Clearance): Given a motion planning problem in d dimensions (having d degrees of freedom), with input of complexity n, and clearance $\delta > 0$, find a complete algorithm with complexity polynomial in terms of n, δ^{-1}, d.

Dan Halperin mentioned the following paper by Alt et al., where they present an analytic solution in this spirit for the case of a rectangle translating and rotating among polygons in the plane. They introduce a related concept, which they call *epsilon tightness* of a motion planning problem. For details, see

- H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra, and C. Uhrig, "Approximate motion planning and the complexity of the boundary of the union of simple geometric figures," Algorithmica, vol. 8, no. 1-6, pp. 391–406, Dec. 1992, doi: 10.1007/BF01758853.

7 Classification of Inverse Kinematic Solutions

Posed by Thomas Cohn

Inverse kinematics algorithms like the Husty-Pfurner approach will return all 16 solutions for a 6R robot arm. However, these solutions are not ordered, so for two different end-effector targets, one can not immediately match the solution sets according to which solutions come from the same "branch". Is there a more efficient approach than cylindrical algebraic decomposition to perform this sort of branch matching for arbitrary 6R manipulators? If not for arbitrary manipulators, what is the broadest class of manipulators that does admit an efficient approach?

8 Chromatic Art Gallery Problem

Posed by Hannah Erickson

Given a simply-connected (hole-free) closed polygonal region P with n vertices, let the *visibility polygon* of a point $p \in P$, denoted $\text{Vis}(p)$, be defined as $\text{Vis}(p) = \{q \in P \mid \overline{pq} \in P\}$ (informally, the set of points of P that can be "seen" by point p without being blocked by the edges of the polygon).

A set of points $S \subset P$ is a *guard set* of P if $\bigcup_{s \in S} \text{Vis}(s) = P$. (The standard Art Gallery Problem attempts to find a minimal guard set for an input polygon.) Let $\text{Guards}(P)$ be the set of guard sets for polygon P.

Two points in $p, q \in P$ *conflict* if $\text{Vis}(p) \cap \text{Vis}(q) \neq \emptyset$. For a guard set $S \in \text{Guards}(P)$ the conflict graph $\text{Conflict}(S, P)$ is a graph with a vertex set S and an edge joining two vertices if they conflict. Let $\chi(G)$ denote the chromatic number of a graph.

The *chromatic guard number* of P, denoted $\chi_G(P)$ is

$$\chi_G(P) = \min_{S \in \text{Guards(P)}} \chi(\text{Conflict}(S, P)) .$$

In other words, it is the minimum number of colors required by any guard set when two guards that see a common point cannot share a color. The motivation for this problem is to determine a relationship between the complexity of the geometry of an environment and the discriminatory power of a sensor suite required to be able to navigate in that environment while always keeping a landmark in view but not seeing two landmarks that are indistinguishable by the sensors.

Question (Chromatic Art Gallery Problem Upper Bound): For each $n \in \mathbb{N}$, what is the maximum chromatic guard number required by any n-vertex polygon?

The best known upper bound is $\lfloor \frac{n}{3} \rfloor$, which is a trivial result obtained from the standard Art Gallery Problem (find a guard set with $\lfloor \frac{n}{3} \rfloor$ guards, and give them all separate colors—see Joseph O'Rourke's book "Art Gallery Theorems and Algorithms" for proofs on this upper bound). The best known construction for a lower bound for arbitrarily large n is $\lfloor \frac{n}{4} \rfloor$ from Erickson and LaValle, Robotics: Science and Systems VII, 2011. An $\frac{n}{4} + C$ bound is conjectured to be tight for sufficiently large n.

Caveat to those searching for this problem: This problem is sometimes referred to as the *strong chromatic art gallery problem*, to differentiate it from the related *conflict-free chromatic art gallery problem* (DOI 10.1007/s00453-012-9732-5), which has weaker requirements on the colorations of the guards.