

34 GEOMETRIC RECONSTRUCTION PROBLEMS

Yann Disser and Steven S. Skiena

INTRODUCTION

Many problems from mathematics and engineering can be described in terms of reconstruction from geometric information. *Reconstruction* is the algorithmic problem of combining the results of measurements of some aspect of a physical or mathematical object to obtain desired information about the object.

In this chapter, we consider three different classes of geometric reconstruction problems. In Section 34.1, we examine static reconstruction problems, where we are given a geometric structure derived from an original structure, and seek to invert this transformation. In Section 34.2, we consider interactive reconstruction problems, where we are permitted to repeatedly “probe” an unknown object at arbitrary places and seek to reconstruct the object using the fewest such probes. Finally, in Section 34.3, we turn to exploration and mapping of unknown environments, where we take the perspective of a mobile agent that locally observes its surroundings and aims to infer information about the global layout of the environment.

Our focus in this chapter is on *exact, theoretical* reconstruction problems from a perspective of computational geometry. In contrast, a significant body of theoretical work in computational geometry is concerned with the *approximate* reconstruction of shapes and surfaces (see Chapter 35). Practical reconstruction problems beyond the scope of this chapter arise in many fields, with examples including computer vision, computer-aided tomography, and the reconstruction of 3D objects from 2D images. In these settings, quality criteria for good solutions are not always mathematically well-defined and may rely on aesthetics, practical applicability, or consistency with reference data.

34.1 STATIC RECONSTRUCTION PROBLEMS

Here we consider inverse problems of the following type. Let A be a geometric structure, and T a transformation such that $T(A) \rightarrow B$, where B is some different geometric structure. Now, given T and B , construct a structure A' such that $T(A') \rightarrow B$. If T is one-to-one, then $A = A'$. If not, we may be interested in finding or counting all solutions.

GLOSSARY

Gabriel graph: A graph whose vertices are points in \mathbb{R}^2 , with an edge (x, y) if points x and y define the diameter of an empty circle.

Relative neighborhood graph: A graph whose vertices are points in \mathbb{R}^2 , with

an edge (x, y) if there exists no point z such that z is closer to x than y is and z is closer to y than x is. See Section 32.1.

Interpoint distance: Distance between a pair of points in \mathbb{R}^2 . The distance is *labeled* if the identities of the two points defining the distance are associated with the distance, and *unlabeled* otherwise.

Stabbing Information: For every vertex v of a polygon the (at most) two edges that are first intersected by the rays wv and uv , where w, u are the neighbors of v along the boundary.

Line cross-sections: A set of lines \mathcal{L} together with the line segments that constitute the intersection of \mathcal{L} with a polygon P .

Visibility polygon: The subset of points in a polygon P visible to a fixed point $x \in P$, i.e., all points $y \in P$ for which the line segment xy is contained in P .

Direction edge/face count: A vector d together with a number k of edges/faces visible in an orthogonal projection in direction d .

Cross-ratio in a triangulation: The ratio $\frac{bd}{ce}$, where a, b, c and a, d, e are the lengths of the edges (in ccw order) of two touching triangles.

Vertex visibility graph: The graph with a node for every vertex of a polygon, and with edges between pairs of vertices that mutually see each other, i.e., whose straight-line connection lies inside the polygon. See Section 33.3.

Point visibility graph: The graph for a set of points with an edge between two points that mutually see each other, i.e., whose straight-line connection does not contain other points. See Section 33.3.

Angle measurement: The ordered list of angles in counter-clockwise order between the edges of the visibility graph at a vertex.

Distance measurement: The ordered list of distances in counter-clockwise order to the vertices visible from a given vertex.

Corner: A vertex of a polygon with an interior angle different from π .

Complex moment of order k : The complex value $\iint_B z^k dx dy$ for a region $B \subset \mathbb{C}$ and $z = x + iy$.

Extended Gaussian image: A transform that maps each face of a convex polyhedron to a vector normal to the face whose length is proportional to the area of the face.

X-ray projection: The length of the intersection of a line with a convex body.

Determination: A class of sets is determined by n directions if there are n fixed directions such that all sets can be reconstructed from X-ray projections along these directions.

Verification: A class of sets is verified by n directions if, for each particular set, there are n X-ray projections that distinguish this set from any other.

MAIN RESULTS

An example of an important class of reconstruction problems is visibility graph reconstruction, i.e., given a graph G , construct a polygon P whose visibility graph is G (see Section 33.3). Results for this and other static reconstruction problems are

summarized in Table 34.1.1. We characterize each problem by its input and the inverted structure we wish to reconstruct. We also specify whether the corresponding transformation is one-to-one, i.e., the result of the reconstruction is unique.

TABLE 34.1.1 Static reconstruction problems.

INPUT	INVERTED STRUCTURE	RESULT	UNIQ	SOURCE
MST with degree ≤ 5	point embedding in \mathbb{R}^2	always realizable	no	[MS91]
MST with degree 6	point embedding in \mathbb{R}^2	NP-hard	no	[EW96]
MST with degree ≥ 7	point embedding in \mathbb{R}^2	never realizable	–	[MS91]
Gabriel graph	point embedding in \mathbb{R}^2	partial charact.	no	[MS80]
rel. neighborhood graph	point embedding in \mathbb{R}^2	partial charact.	no	[LS93]
Delaunay triangulation	point embedding in \mathbb{R}^2	partial charact.	no	[Di190][Sug94]
Voronoi diagram	point embedding in \mathbb{R}^2	partial charact.	no	[AB85]
point visibility graph	point embedding in \mathbb{R}^2	$\exists \mathbb{R}$ -complete	no	[GR15][CH17]
labeled interpoint distances	points realizing these in \mathbb{R}^d	NP-hard	no	[Sax79]
all unlab. interpoint dists	points realizing these in \mathbb{R}^1	$O(2^n n \log n)$	no	[LSS03]
all unlab. interpoint dists	points realizing these in \mathbb{R}^d	NP-hard	no	[LSS03]
vertex visibility graph	polygon realizing it	$\in \text{PSPACE}$	no	[Eve90]
distance visibility graph	polygon realizing it	$O(n^2)$	yes	[CL92]
endpoints $V \subset \mathbb{R}^2$	orthogonal line segments	$O(n \log n)$	no	[RW93]
endpoints $V \subset \mathbb{R}^2$	disjoint orth. line segments	NP-hard	no	[RW93]
corners $V \subset \mathbb{R}^2$	orthogonal polygon	$O(n \log n)$	yes	[O'R88]
corners $V \subset \mathbb{R}^2$, 3 slopes	polygon with these slopes	NP-hard	no	[FW90]
vertices $V \subset \mathbb{R}^2$	orthogonal polygon	NP-hard	no	[Rap89]
angle measurements	compatible polygon	$O(n^2)$	yes	[DMW11][CW12]
stabbing information	compatible orthogonal poly	$O(n \log n)$	no	[JW02]
set of s line cross-sections	all compatible polygons	$O(s \log s)$ /poly	no	[SBG06]
cross-ratios, bound. angles	compatible polygon	uniqueness	yes	[Sno99]
set of visibility polygons	compatible polygon	NP-hard	no	[BDS11]
direction edge counts	convex polygon	charact., algo.	no	[BHL11]
direction face counts	convex 3D polyhedron	NP-hard	no	[BHL11]
ext. Gaussian image	convex 3D polyhedron	$O(n \log n)$ /iter.	yes	[Lit85]
4 complex moments	triangle in \mathbb{C}	uniqueness	yes	[Dav77]
$2n$ complex moments	vertices of polygon in \mathbb{C}	algorithm	yes	[MVK ⁺ 95]
X-ray projections	convex body if unique	algorithm	yes	[GK07]


Another class of problems concerns proximity drawability. Given a graph G , we seek a set of points corresponding to vertices of G such that two points are “sufficiently” close if and only if there is an edge in G for the corresponding vertices. Examples of proximity drawability problems include finding points to realize graphs as minimum spanning trees (MST), Delaunay triangulations (Chapter 29), Gabriel graphs, and relative neighborhood graphs (RNGs) (Chapter 32). Although many of the results are quite technical, Liotta [Lio13] provides an excellent survey of results on these and other classes of proximity drawings; see also Chapter 55.

To provide some intuition about the minimum spanning tree results, observe that low degree graphs are easily embedded as point sets. If the maximum degree is 2, i.e., the graph is a simple path, then any straight line embedding will work. One can show that any two line segments vu , vw corresponding to adjacent edges of the tree need to form an angle not smaller than $\pi/3$, since the segment uw cannot be shorter than vu or vw . This implies that degrees larger than 6 cannot be realized, and forces the neighbors of degree 6 vertices to be spaced at equal angles of $\pi/3$, a very restrictive condition leading to the hardness result.

Other typical reconstruction problems are concerned with constructing polygons that are compatible with given geometrical parameters. See Figure 34.1.1 for three such examples (taken from [CD⁺13b]). In the first example, the angles between lines-of-sight are known at each vertex, and it turns out that this information uniquely determines the polygon (up to scaling and rotation) [DMW11]. In the second example, a polygon has to be constructed from its intersection with a set of lines. In this case, the polygon is not uniquely determined, but all compatible polygons can be enumerated efficiently [SBG06]. The third example shows the “orthogonal connect-the-dots” problem, where an orthogonal polygon has to be recovered from the coordinates of its vertices. This is uniquely and efficiently possible if vertices of degree π are forbidden [O’R88], and otherwise it is NP-hard to find any compatible polygon [Rap89].

FIGURE 34.1.1

From left to right: Reconstruction from angle measurements, line cross-sections, and vertices.



Another important set of problems concern reconstructing objects from a fixed set of X-ray projections, conventionally called Hammer’s X-ray problem [Ham63]. Different problems arise depending upon whether the X-rays originate from a point or line source, and whether we seek to verify or determine the object. A selection of results on parallel X-rays (line sources) are listed in Table 34.1.2. For example, parallel X-rays in certain sets of four directions suffice to determine any convex body if the directions are not a subset of the edges of an affinely regular polygon.

If the directions do form such a subset, then there exist noncongruent polygons that are not distinguished by any number n of parallel X-rays in these directions. Nevertheless, any pair of nonparallel directions suffice to determine “most” (in the sense of Baire category) convex sets.

There is also a collection of results on *point source X-rays*. For example, convex sets in \mathbb{R}^2 are determined by directed X-rays from three noncollinear point sources. The substantial literature on such X-ray problems is very well covered by Gardner’s monograph [Gar06], from which two of the open problems listed below are drawn.

The related field of *discrete tomography* is inspired by the use of electron microscopy to reconstruct the positions of atoms in crystal structures. A typical problem is placing integers in a matrix so as to realize a given set of row and column sums. The problem becomes more complex when the reconstructed body must satisfy connectivity constraints or simultaneously satisfy row/column sums of multiple colors. Collections of survey articles on discrete tomography include Herman and Kuba [HK99, HK07].

TABLE 34.1.2 Selected results on Hammer’s X-ray problem.

DIM	PROBLEM	SETS	RESULT	SOURCE
2	verify	convex polygons	2 parallel X-rays do not suffice	[Gar83]
	verify	convex set	3 parallel X-rays suffice	[Gie62]
	determine	convex set	4 parallel X-rays suffice	[GM80][GG97]
	determine	convex set	n arb. paral. X-rays do not suffice	[Gie62]
	determine	star-shaped poly. convex body	finite num. paral. X-rays insufficient 3 point X-rays suffice	[Gar92] [Vol86]
3	determine	convex body	4 parallel, coplanar X-rays suffice	[Gar06]
	determine	convex body	4 arb., paral. X-rays do not suffice	[Gar06]
d	determine	convex body	2 parallel X-rays “usually” suffice	[VZ89]
	verify	compact sets	no finite number of directions suffice	[Gar92]

OPEN PROBLEMS

1. Give an efficient algorithm to reconstruct a set of n points on the line from the set of $\binom{n}{2}$ unlabeled interpoint distances it defines: see [LSS03]. Note that the problem indeed remains open as of this writing, despite published comments to the contrary: see [DGN05].
2. Is a polygon uniquely determined by its distance measurements?
3. Give an algorithm to determine whether a graph is the visibility graph of a simple polygon [GG13, Problem 29].
4. Characterize the convex sets in \mathbb{R}^2 that can be determined by two parallel X-rays [Gar06, Problem 1.1].
5. Are convex bodies in \mathbb{R}^3 determined by parallel X-rays in some set of five directions [Gar06, Problem 2.2]?

34.2 INTERACTIVE RECONSTRUCTION PROBLEMS

In *static* reconstruction problems all available data about the structure that has to be reconstructed is revealed in a one-shot fashion. In contrast, *interactive* reconstruction allows to request data in multiple rounds, and allows each request to depend on the data gathered so far. This process is generally modeled via *geometric probing*, which defines access to the unknown geometric structure via a mathematical or physical measuring device, a *probe*. A variety of problems from robotics, medical instrumentation, mathematical optimization, integral and computational geometry, graph theory, and other areas fit into this paradigm.

The model of geometric probing was introduced by Cole and Yap [CY87] and inspired by work in robotics and tactile sensing. A substantial body of work has followed, which is extensively surveyed in [Ski92]. A collection of open problems in probing appears in [Ski89a]. More recent probing models include proximity probes [ABG15], wedge probes [BCSS15], and distance probes [AM15].

GLOSSARY

Determination: The algorithmic problem of computing how many probes of a certain type are necessary to completely determine or reconstruct an object drawn from a particular class of objects.

Verification: The algorithmic problem of, given a supposed description of an object, computing how many probes of a certain type are necessary to test whether the description is valid.

Model-based: A problem where any object is constrained to be one of a known, finite set of m possible objects.

Point probe: An oracle that tests whether a given point is within the object.

Finger probe: An oracle that returns the first point of intersection between a directed line and the object.

Hyperplane probe: An oracle that returns the first time when a hyperplane moving parallel to itself intersects the object.

X-ray probe: An oracle that measures the length of the intersection between a line and the object.

Silhouette probe: An oracle that returns a $(d-1)$ -dimensional projection (in a given direction) of the d -dimensional object.

Halfspace probe: An oracle that measures the area or volume of the intersection between a halfspace and the object.

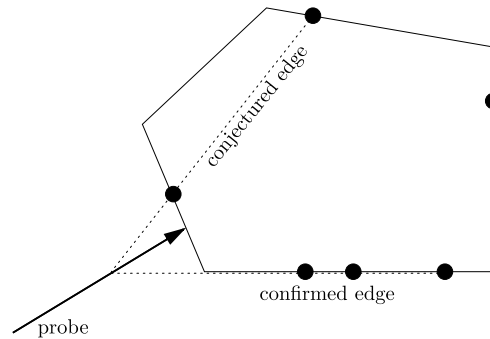
Cut-set probe: An oracle that, for a specified graph and partition of the vertices, returns the size of the cut-set determined by the partition.

Proximity probe: An oracle that returns the nearest point of the object to a specified origin point.

Wedge probe: An oracle that, for a specified origin point and translation direction, returns the first contact points between the object and a moving wedge with angle ω .

Distance probe: An oracle that returns the distance between two named points.
Fourier probe: An oracle that for a given vector $\xi \in \mathbb{R}^2$ returns $\int_D e^{-i\langle \xi, x \rangle} dx$ with respect to a region $D \subset \mathbb{R}^2$.

FIGURE 34.2.1
 Determining the next edge of P using finger probes.



MAIN RESULTS

For a particular probing model, the determination problem asks how many probes are sufficient to completely reconstruct an object from a given class. For example, Cole and Yap's strategy for reconstructing a convex polygon P from finger probes is based on the observation that three collinear contact points must define an edge. The strategy, illustrated in Figure 34.2.1, repeatedly aims a probe at the intersection point between a confirmed edge (defined by three collinear points) and a conjectured edge (defined by two contact points). If this intersection point is indeed a contact point, another vertex is determined due to convexity; if not, the existence of another edge can be inferred. Since we avoid probing the interior of any edge that has been determined, roughly $3n$ probes suffice in total, since not more than one edge can be hit four times. Table 34.2.1 summarizes probing results for a wide variety of models. In the table, f_i denotes the number of i -dimensional faces of P .

Cole and Yap's finger probing model is not powerful enough to determine nonconvex objects. There are three major reasons for this. A tiny crack in an edge can go forever undetected, since no finite strategy can explore the entire surface of the polygon. Second, it is easy to construct nonconvex polygons whose features cannot be entirely contacted with straight-line probes originating from infinity. Finally, for nonconvex polygons there exists no constant k such that k collinear probes determine an edge. To generalize the class of objects, enhanced finger probes have been considered. One such probe [ABY90] returns surface normals as well as contact points, eliminating the second problem. When restricted to polygons with no two edges defined by the same supporting line, the first and third problems are eliminated as well.

In the verification problem, we are given a description of a putative object, and charged with using a small number of probes to prove that the description is correct. Verification is clearly no harder than determination, since we are free to ignore the description in planning the probes, and could simply compare the determined object to its description. Sometimes significantly fewer probes suffice for verification. For example, we can verify a putative convex polygon with $2n$

TABLE 34.2.1 Upper and lower bounds for determination for various probing models.

PROBE	OBJECT	LOWER	UPPER	SOURCE
finger	convex polygon	$3n$	$3n$	[CY87]
finger (n known)	convex polygon	$2n + 1$	$3n - 1$	[CY87]
finger	convex polyhedron in \mathbb{R}^d	$df_0 + f_{d-1}$	$f_0 + (d+2)f_{d-1}$	[LB88][DEY90]
finger (model based)	convex polygon	$n - 1$	$n + 4$	[JS92]
$k = 2$ or 3 fingers	convex polygon	$2n - k$	$2n$	[LB92]
4 or 5 fingers	convex polygon	$(4n - 5)/3$	$\lfloor (4n + 2)/3 \rfloor$	[LB92]
$k \geq 6$ fingers	convex polygon	n	$n + 1$	[LB92]
enh. fingers	nondegenerate polygon	$3n - 3$	$3n - 3$	[ABY90]
Line	convex polygon	$3n + 1$	$3n + 1$	[Li88]
Line (model based)	convex polygon	$2n - 3$	$2n + 4$	[JS92]
Silhouette	convex polygon	$3n - 2$	$3n - 2$	[Li88]
Silhouette	convex polyhedron in \mathbb{R}^3	$f_2/2$	$5f_0 + f_2$	[DEY90]
X-ray	convex polygon	$3n - 3$	$5n + 19$	[ES88]
Parallel X-ray	convex polygon	3	3	[ES88]
Parallel X-ray	nondegenerate polygon	$\lceil \log n \rceil - 2$	$2n + 2$	[MS96]
Halfplane	convex polygon	$2n$	$7n + 7$	[Ski91]
Proximity	convex polygon	$2n$	$3.5n + 5$	[ABG15]
Wedge ($\omega \leq \pi/2$)	convex polygon	$2n + 2$	$2n + 5$	[BCSS15]
Fourier	nondegenerate polygon		$3n$	[WP16]
Cut-set	embedded graph	$\binom{n}{2}$	$\binom{n}{2}$	[Ski89b]
Cut-set	unembedded graph	$\Omega(n^2/\log n)$	$O(n^2/\log n)$	[Ski89b]
Distance (2 rounds)	points in \mathbb{R}^1	$9n/8$	$9n/7 + O(1)$	[AM15]

probes by sending one finger probe to contact each vertex and the interior of each edge. This gives three contact points on each edge, which, by convexity, suffices to verify the polygon. Table 34.2.2 summarizes results in verification.

Of course, there are other classes of problems that do not fit so easily into the confines of these tables. Verification is closely related to approximate geometric testing; see [ABM⁺97, Rom95]. An interesting application of probing to nonconvex polygons is presented in [HP99]. See [Ric97, Ski92] for discussions of probing with uncertainty and tactile sensing in robotics.

TABLE 34.2.2 Upper and lower bounds for verification for various probing models.

PROBE	OBJECT	LOWER	UPPER	SOURCE
Finger	convex polygon	$2n$	$2n$	[CY87]
Finger (n known)	convex polygon	$3\lceil n/2 \rceil$	$3\lceil n/2 \rceil$	[Ski88]
Line	convex polygon	$2n$	$2n$	[DEY90]
X-ray	convex polygon	$3n/2$	$3n/2 + 6$	[ES88]
Halfplane	convex polygon	$2n/3$	$n + 1$	[Ski91]

OPEN PROBLEMS

1. Tighten the gap between the lower and upper bounds for determination for finger probes in higher dimensions [DEY90].
2. Tighten the bounds for determination of convex n -gons with X-ray probes. Does a finite number (i.e., $f(n)$) of parallel X-ray probes suffice to verify or determine simple n -gons? Since each parallel X-ray probe provides a representation of the complete polygon, there is hope to detect arbitrarily small cracks in a finite number of probes; see [MS96].
3. Consider generalizations of halfplane probes to higher dimensions. How many probes are necessary to determine convex (or nonconvex) polyhedra?
4. Silhouette probes return the shadow cast by a polytope in a specified direction. These dualize to *cross-section probes* that return a slice of the polytope. Tighten the current bounds [DEY90] on determination with silhouettes in \mathbb{R}^3 .

34.3 GEOMETRIC EXPLORATION AND MAPPING

In *geometric mapping* we face the problem of reconstructing a surrounding geometric structure using local perception. We take the perspective of an agent *exploring* an initially unknown environment, while trying to piece together information gathered through its *sensors* in order to (partially) infer the global structure, i.e., a *map*. Settings vary in the types of environments that are considered, the movement and sensor capabilities of the agent, its initial knowledge of the environment, and the type of map that needs to be inferred. Similarly to interactive reconstruction, the movements of the agent may depend on its past observations, and we can view the setting as geometric probing with restricted transitions between consecutive probes.

Importantly, we generally require the map to be *uniquely* reconstructed, and thus the first question when studying a specific setting is whether mapping is feasible, i.e., whether the movement and sensor capabilities suffice to uniquely infer the map at some point (irrespective of running time). If this is the case, we are interested in mapping strategies that minimize the required movement of the agent (irrespective of computing time).

GLOSSARY

Exploration: The problem of navigating and covering an initially unknown environment using local sensing.


Mapping: The exploration problem with the additional objective of (uniquely) reconstructing a representation (map) of the environment.

Graph exploration: The problem of visiting all vertices of an initially unknown graph with an agent moving between vertices along edges of the graph. The edges of the graph are labeled with locally unique labels, and, in each step, the agent chooses a label of an outgoing edge and moves to its other end.

- Anonymous graph:** A graph with vertices that cannot be distinguished (unless their degrees differ). In contrast, a **labeled** graph has unique node identifiers.
- Combinatorial visibility vector (c vv):** A vector $c \in \{0, 1\}^{d-1}$ for a vertex v of degree d of a polygon, with $c_i = 1$ exactly if the i -th and $(i + 1)$ -st vertex visible from v (in ccw order) are neighbors along the boundary of the polygon.
- C vv sensor:** Provides the combinatorial visibility vector at the current vertex.
- Look-back sensor:** Provides the label of the edge leading back to the previous location of the agent.
- Pebble:** A device that can be dropped at a vertex of an anonymous graph to make the vertex distinguishable, and can be picked up and reused later.
- Angle sensor:** Provides the angle measurement (see Section 34.1) at the current vertex.
- Angle type sensor:** Provides a bit $t \in \{0, 1\}$ for each pair of vertices u, w visible from the current vertex v , with $t = 1$ exactly if the angle between the segments vu and vw is larger than π .
- Direction sensor:** Provides the angle between some globally fixed line and the line segments connecting the current vertex to each visible vertex (in ccw order).
- Distance sensor:** Provides the lengths of the line segments connecting the current vertex to each visible vertex (in ccw order). A **continuous** distance sensor provides the distance to the boundary of the environment in each direction, i.e., it provides the visibility polygon (see Section 34.1) of the current location.
- Contact sensor:** Provides a bit $c \in \{0, 1\}$, with $c = 1$ exactly if the agent's location corresponds to a point on the boundary of the environment.
- Cut:** The maximal extension vx of a boundary edge uv of a polygon P , such that v is a reflex vertex of P , and vx is collinear to uv and lies inside P .
- Cut diagram:** A graph associated with a polygon, with a node for each point where (two or more) cuts and/or boundary edges of the polygon intersect (in particular for each vertex of the polygon), and an edge between two points that are neighbors along a cut or a boundary edge.

FIGURE 34.3.1

From left to right: angle, angle type, distance, and direction sensor.



MAIN RESULTS

Research on geometric exploration and mapping has mainly considered polygonal environments, either with a focus on feasibility (weak sensors) or efficiency (strong sensors). With regards to feasibility, a key question is how minimalistic an agent model may be to still allow inferring a meaningful map of the environment. Suri et al. [SVW08] introduced such a model, where an agent moves from vertex to vertex along lines-of-sight in a simple polygon, and only observes the incident lines-of-sight in counter-clockwise (ccw) order when at a vertex. Obviously, such a minimalistic agent cannot hope to reconstruct the full geometry of the environment. Instead, the goal in this model is to infer the visibility graph that has an edge for each line-of-sight (see Section 34.1). Note that the visibility graph is a reasonable topological map, because, for example, it contains all shortest vertex-to-vertex paths in the polygon (see Chapter 31). Suri et al. [SVW08] showed that, if the agent is additionally equipped with a pebble, it can always reconstruct the visibility graph. On the other hand, Brunner et al. [BMS⁺08] showed that without pebbles the problem is infeasible, and not even the total number n of vertices can be inferred. It remains open, whether knowledge of n alone already allows mapping. Results for various extensions of the basic model are in Table 34.3.1.

TABLE 34.3.1 Summary of results on visibility graph mapping.

SENSOR	INFO	FEASIBLE	RUNTIME	SOURCE
cvv, look-back	–	no		[BMS ⁺ 08]
pebble	–	yes	poly	[SVW08]
angle	–	yes	poly	[DMW11]
look-back	n	yes	poly	[CD ⁺ 13a]
angle type	n	yes	exp	[CDD ⁺ 15]
directions	n	yes	exp	[DGM ⁺ 14]
distance	n	open	exp	
<i>none</i>	n	open	exp	

TABLE 34.3.2 Summary of results on mapping rooms with obstacles.

ROOM	OBSTACLES	COMP. RATIO	SOURCE
orthogonal polygon	none	≤ 2	[DKP98]
orthogonal polygon	none	$\geq 5/4$	[Kle94]
polygon	none	≤ 26.5	[HIKK01]
orthogonal polygon	orthogonal	$O(n)$	[DKP98]
rectangle	rectangular	$\Omega(\sqrt{n})$	[AKS02]

Another simplistic model was studied by Katsev et al. [KYT⁺11]. In their model, the agent can only move along the boundary and across cuts of the polygon, and the objective is to reconstruct the cut diagram of the environment. They show that this is possible if the agent can distinguish convex from reflex vertices and

distinguish the two cut edges at a reflex vertex in ccw order.

A much more powerful model was studied by Deng et al. [DKP98]. Here the agent has a global sense of direction, can move freely in the interior of a polygonal environment (the “room”) with polygonal obstacles, and has a continuous distance sensor that provides the exact geometry of the visible portion of the environment from the current location. Results in this model concern the competitive ratio between the length of the exploration path and an offline optimum path (of minimum length) ensuring that all interior points of the environment are visible at some point (see Table 34.3.2). Note the difference to the search problem where an object needs to be located in the environment and the offline optimum only needs to establish visibility to the corresponding location.

The general problem of mapping unknown discrete environments can be formulated in terms of graph exploration (see Table 34.3.3). In this abstract setting, the agent moves between vertices of an initially unknown, directed (strongly connected) graph, with the goal of inferring the graph up to isomorphism. For this purpose, we assume the outgoing edges at a vertex to have locally unique labels that the agent sees and uses to specify its moves. Note that, in this model, there is no immediate way to distinguish vertices with the same degrees, and, in particular, a single agent cannot hope to distinguish two 3-regular graphs, even if it knows the number of vertices. Bender and Slonim [BS94] showed that mapping is feasible for two agents in polynomial time, and Bender et al. [BFR⁺02] showed that $\Theta(\log \log n)$ pebbles are necessary and sufficient for a single agent to achieve polynomial time, i.e., “a friend is only worth $\Theta(\log \log n)$ pebbles.” The main result of Bender et al. [BFR⁺02] is that a single pebble suffices if (a bound on) n is known.

TABLE 34.3.3 Summary of results on graph exploration and mapping.

GRAPH	#AGENTS	EXTRAS	RESULT	SOURCE
anonymous digraph	1	n known	infeasible	
anonymous digraph	2	randomized	$O(n^5 \Delta^2)$ algorithm	[BS94]
anonymous digraph	1	1 pebble, n known	$O(n^8 \Delta^2)$ algorithm	[BFR ⁺ 02]
anonymous digraph	1	$O(\log \log n)$ pebbles	poly time algorithm	[BFR ⁺ 02]
anonymous digraph	1	$o(\log \log n)$ pebbles	exp time needed	[BFR ⁺ 02]
labeled graph	const		comp. ratio: $O(1)$	DFS
labeled tree	$k < \sqrt{n}$		CR: $\Omega(\log k / \log \log k)$	[DLS07]
labeled graph	$k = \sqrt{n}$	randomized	CR: $\Omega(\sqrt{\log k} / \log \log k)$	[OS12]
labeled tree	k		comp. ratio: $O(k / \log k)$	[FGKP06]
labeled graph	$n^{2+\varepsilon}$		comp. ratio: $O(1)$	[DDK ⁺ 15]
labeled graph	$\exp(n)$		comp. ratio: 1	BFS

In case the vertices of the graph are distinguishable and edges are undirected, a single agent can map any graph simply using depth-first search until every edge was visited. This strategy visits every edge at most twice, and thus trivially yields a competitive ratio of 2, compared with an offline optimal traversal that visits all edges. On the other hand, a team of exponentially many agents can execute a breadth-first search style strategy by splitting all agents at a vertex evenly among all unexplored neighbors in each step. Obviously, this strategy needs an optimal number of steps. In general, a team of k agents needs at least $O(D + n/k)$ steps,

where D is the maximum shortest path distance from the starting location to an unexplored vertex. Dereniowski et al. [DDK⁺15] showed that a constant competitive ratio can already be obtained with (roughly) quadratic team size $k = Dn^{1+\epsilon}$. The asymptotically best-possible competitive ratios for smaller, super-constant team sizes remain open. The best known lower bound on the competitive ratio of deterministic algorithms of $\Omega(\log k / \log \log k)$ for the domain $k < \sqrt{n}$ (with $n/k > D$) is due to Dynia et al. [DLS07]. This bound holds already on trees. Fraigniaud et al. [FGKP06] gave an algorithm for trees that achieves a ratio of $O(k / \log k)$.

OPEN PROBLEMS

1. Can a visibility graph be mapped by an agent without additional sensors, i.e., by observing only degrees, if the number n of vertices is known? Note that knowledge of (some bound on) n is necessary [BMS⁺08].
2. Can a visibility graph be mapped with an agent using a distance sensor?
3. Close the gaps for the mapping of rooms with/without obstacles.
4. What is the best-possible competitive ratio for mapping labeled graphs with k agents in the domain $k \in \omega(1) \cap o(n^{2+\epsilon})$?

34.4 SOURCES AND RELATED MATERIAL

SURVEYS

- [Lio13]: Survey on embedding proximity graphs (Table 34.1.1).
 [Gar06]: Survey of Hammer's X-ray problem and related work in geometric tomography (Table 34.1.2).
 [HK99, HK07]: Surveys on discrete tomography.
 [Rom95]: Survey on geometric testing.
 [Ski92]: Survey on geometric probing (Table 34.2.1).
 [CD⁺13b]: Survey on mapping polygons (Table 34.3.1).

RELATED CHAPTERS

- Chapter 31: Shortest paths and networks
- Chapter 32: Proximity algorithms
- Chapter 33: Visibility
- Chapter 35: Curve and surface reconstruction
- Chapter 50: Algorithmic motion planning
- Chapter 51: Robotics
- Chapter 55: Graph drawing
- Chapter 61: Rigidity and scene analysis

REFERENCES

- [AB85] P.F. Ash and E.D. Bolker. Recognizing Dirichlet tessellations. *Geom. Dedicata*, 19:175–206, 1985.
- [ABG15] A. Adler, F. Banahi, and K. Goldberg. Efficient proximity probing algorithms for metrology. *IEEE Trans. Autom. Sci. Eng.*, 12:84–95, 2015.
- [ABM⁺97] E.M. Arkin, P. Belleville, J.S.B. Mitchell, D.M. Mount, K. Romanik, S. Salzberg, and D.L. Souvaine. Testing simple polygons. *Comput. Geom.*, 8:97–114, 1997.
- [ABY90] P.D. Alevizos, J.-D. Boissonnat, and M. Yvinec. Non-convex contour reconstruction. *J. Symbolic Comput.*, 10:225–252, 1990.
- [AKS02] S. Albers, K. Kursawe, and S. Schuierer. Exploring unknown environments with obstacles. *Algorithmica*, 32:123–143, 2002.
- [AM15] M. Alam and A. Mukhopadhyay. Three paths to point placement. In *Proc. 1st Conf. Algorithms Discrete Appl. Math.*, vol. 8959 of *LNCS*, pages 33–44, Springer, Berlin, 2015.
- [BCSS15] P. Bose, J.-L. De Carufel, A. Shaikhet, and M. Smid. Probing convex polygons with a wedge. *Comput. Geom.*, 58:34–59, 2016.
- [BDS11] T. Biedl, S. Durocher, and J. Snoeyink. Reconstructing polygons from scanner data. *Theoret. Comput. Sci.*, 414:4161–4172, 2011.
- [BFR⁺02] M.A. Bender, A. Fernández, D. Ron, A. Sahai, and S. Vadhan. The power of a pebble: Exploring and mapping directed graphs. *Information and Computation*, 176:1–21, 2002.
- [BHL11] T. Biedl, M. Hasan, and A. López-Ortiz. Reconstructing convex polygons and convex polyhedra from edge and face counts in orthogonal projections. *Internat. J. Comput. Geom. Appl.*, 21:215–239, 2011.
- [BMS⁺08] J. Brunner, M. Mihalák, S. Suri, E. Vicari, and P. Widmayer. Simple robots in polygonal environments: A hierarchy. In *Proc. 4th Workshop on Algorithmic Aspects of Wireless Sensor Networks*, vol. 5389 of *LNCS*, pages 111–124, Springer, Berlin, 2008.
- [BS94] M.A. Bender and D.K. Slonim. The power of team exploration: Two robots can learn unlabeled directed graphs. In *Proc. 35th IEEE Sympos. Found. Comp. Sci.*, pp. 75–85, 1994.
- [CD⁺13a] J. Chalopin, S. Das, Y. Disser, M. Mihalák, and P. Widmayer. Mapping simple polygons: How robots benefit from looking back. *Algorithmica*, 65:43–59, 2013.
- [CD⁺13b] J. Chalopin, S. Das, Y. Disser, M. Mihalák, and P. Widmayer. Simple agents learn to find their way: An introduction on mapping polygons. *Discrete Appl. Math.*, 161:1287–1307, 2013.
- [CDD⁺15] J. Chalopin, S. Das, Y. Disser, M. Mihalák, and P. Widmayer. Mapping simple polygons: The power of telling convex from reflex. *ACM Trans. Algorithms*, 11:33–49, 2015.
- [CH17] J. Cardinal and U. Hoffmann. Recognition and complexity of point visibility graphs. *Discrete Comput. Geom.*, 57:164–178, 2017.
- [CL92] C.R. Coullard and A. Lubiw. Distance visibility graphs. *Internat. J. Comput. Geom. Appl.*, 2:349–362, 1992.

- [CW12] D.Z. Chen and H. Wang. An improved algorithm for reconstructing a simple polygon from its visibility angles. *Comput. Geom.*, 45:254–257, 2012.
- [CY87] R. Cole and C.K. Yap. Shape from probing. *J. Algorithms*, 8:19–38, 1987.
- [Dav77] P.J. Davis. Plane regions determined by complex moments. *J. Approximation Theory*, 19:148–153, 1977.
- [DDK⁺15] D. Dereniowski, Y. Disser, A. Kosowski, D. Pajak, and P. Uznański. Fast collaborative graph exploration. *Information and Computation*, 243:37–49, 2015.
- [DEY90] D. Dobkin, H. Edelsbrunner, and C.K. Yap. Probing convex polytopes. In I.J. Cox and G.T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 328–341, Springer, Berlin, 1990.
- [DGM⁺14] Y. Disser, S.K. Ghosh, M. Mihalák, P. Widmayer. Mapping a polygon with holes using a compass. *Theoret. Comput. Sci.*, 553:106–113, 2014.
- [DGN05] A. Daurat, Y. Gerard, and M. Nivat. Some necessary clarifications about the chords’ problem and the Partial Digest Problem. *Theoret. Comput. Sci.*, 347:432–436, 2005.
- [Dil90] M.B. Dillencourt. Realizability of Delaunay triangulations. *Inform. Process. Lett.*, 33:424–432, 1990.
- [DKP98] X. Deng, T. Kameda, and C. Papadimitriou. How to learn an unknown environment I: The rectilinear case. *J. ACM*, 45:215–245, 1998.
- [DLS07] M. Dynia, J. Lopuszański, and C. Schindelhauer. Why robots need maps. In *Proc. 14th Coll. Struct. Inform. Comm. Complexity (SIROCCO)*, vol. 4474 of LNCS, pages 41–50, Springer, Berlin, 2007.
- [DMW11] Y. Disser, M. Mihalák, and P. Widmayer. A polygon is determined by its angles. *Comput. Geom.*, 44:418–426, 2011.
- [ES88] H. Edelsbrunner and S.S. Skiena. Probing convex polygons with X-rays. *SIAM J. Comput.*, 17:870–882, 1988.
- [Eve90] H. Everett. *Visibility Graph Recognition*. PhD Thesis, University of Toronto, 1990.
- [EW96] P. Eades and S. Whitesides. The realization problem for Euclidean minimum spanning trees is NP-hard. *Algorithmica*, 16:60–82, 1996.
- [FGKP06] P. Fraigniaud, L. Gąsieniec, D.R. Kowalski, and A. Pelc. Collective tree exploration. *Networks*, 48:166–177, 2006.
- [FW90] F. Formann and G.J. Woeginger. On the reconstruction of simple polygons. *Bull. Eur. Assoc. Theor. Comput. Sci. ETACS*, 40:225–230, 1990.
- [Gar83] R.J. Gardner. Symmetrals and X-rays of planar convex bodies. *Arch. Math. (Basel)*, 41:183–189, 1983.
- [Gar92] R.J. Gardner. X-rays of polygons. *Discrete Comput. Geom.*, 7:281–293, 1992.
- [Gar06] R.J. Gardner. *Geometric Tomography*, second edition. Cambridge Univ. Press, 2006.
- [GG97] R.J. Gardner and P. Gritzmann. Discrete tomography: Determination of finite sets by X-rays. *Trans. Amer. Math. Soc.*, 349:2271–2295, 1997.
- [GG13] S.K. Ghosh and P.P. Goswami. Unsolved problems in visibility graphs of points, segments, and polygons. *ACM Comput. Surv.*, 46:1–29, 2013.
- [Gie62] O. Giering. Bestimmung von Eibereichen und Eikörpern durch Steiner-Symmetrisierungen. In *Sitzungsberichte der Bayerischen Akademie der Wissenschaften, Math.-Nat. Kl.*, pages 225–253, 1962.
- [GK07] R.J. Gardner and M. Kinderlen. A solution to Hammer’s X-ray reconstruction problem. *Adv. Math.*, 214:323–343, 2007.

- [GM80] R.J. Gardner and P. McMullen. On Hammer's X-ray problem. *J. London Math Soc.*, 21:171–175, 1980.
- [GR15] S.K. Ghosh and B. Roy. Some results on point visibility graphs. *Theoret. Comput. Sci.*, 575:17–32, 2015.
- [Ham63] P.C. Hammer. Problem 2. *Proc. Symposia in Pure Mathematics*, 7:498–499, AMS, Providence, 1963.
- [HIKK01] F. Hoffmann, C. Icking, R. Klein, and K. Kriegel. The polygon exploration problem. *SIAM J. Comput.*, 31:577–600, 2001.
- [HK99] G.T. Herman and A. Kuba. *Discrete Tomography: Foundations, Algorithms, and Applications*. Springer, Berlin, 1999.
- [HK07] G.T. Herman and A. Kuba, editors. *Advances in Discrete Tomography and its Applications*. Birkhäuser, Basel, 2007.
- [HP99] K. Hunter and T. Pavlidis. Non-interactive geometric probing: Reconstructing non-convex polygons. *Comput. Geom.* 14:221–240, 1999.
- [JS92] E. Joseph and S.S. Skiena. Model-based probing strategies for convex polygons. *Comput. Geom.*, 2:209–221, 1992.
- [JW02] L. Jackson and S.K. Wismath. Orthogonal polygon reconstruction from stabbing information. *Comput. Geom.*, 23:69–83, 2002.
- [Kle94] J.M. Kleinberg. On-line search in a simple polygon. In *Proc. 5th ACM-SIAM Sympos. Discrete Algorithms*, pages 8–15, 1994.
- [KYT⁺11] M. Katsev, A. Yershova, B. Tovar, R. Ghrist, and S.M. LaValle. Mapping and pursuit-evasion strategies for a simple wall-following robot. *IEEE Trans. Robotics*, 27:113–128, 2011.
- [LB88] M. Lindenbaum and A. Bruckstein. Reconstructing convex sets from support hyperplane measurements. Tech. Report 673, Dept. Electrical Engineering, Technion, 1988.
- [LB92] M. Lindenbaum and A. Bruckstein. Parallel strategies for geometric probing. *J. Algorithms*, 13:320–349, 1992.
- [Li88] R. Li Shuo-Yen. Reconstruction of polygons from projections. *Inform. Process. Lett.*, 28:235–240, 1988.
- [Lio13] G. Liotta. Proximity drawings. In R. Tamassia, editor, *Handbook of Graph Drawing*, pages. 115–154, CRC Press, Boca Raton, 2013.
- [Lit85] J.J. Little. Extended Gaussian images, mixed volumes, shape reconstruction. In *Proc. 1st Sympos. Comput. Geom.*, pages 15–23, ACM Press, 1985.
- [LS93] A. Lubiw and N. Sleumer. Maximal outerplanar graphs are relative neighbourhood graphs. In *Proc. 5th Canadian Conf. Comput. Geom.*, pages 198–203, 1993.
- [LSS03] P. Lemke, S.S. Skiena, and W.D. Smith. Reconstructing sets from interpoint distances. In *Discrete and Computational Geometry—The Goodman-Pollack Festschrift (B. Aronov et al., editors)*, pages 597–691, Springer, Heidelberg, 2003.
- [MS80] D.W. Matula and R.R. Sokal. Properties of Gabriel graphs relevant to geographic variation research and the clustering of points in the plane. *Geographical Analysis*, 12:205–222, 1980.
- [MS91] C.L. Monma, S. Suri. Transitions in geometric minimum spanning trees. *Discrete Comput. Geom.*, 8:265–293, 1992.
- [MS96] H. Meijer and S.S. Skiena. Reconstructing polygons from X-rays. *Geom. Dedicata*, 61:191–204, 1996.

- [MVK⁺95] P. Milanfar, G.C. Verghese, W.C. Karl, and A.S. Willsky. Reconstructing polygons from moments with connections to array processing. *IEEE Trans. Signal Processing*, 43:432–443, 1995.
- [O’R88] J. O’Rourke. Uniqueness of orthogonal connect-the-dots. *Mach. Intell. Pattern Recogn.*, 6:97–104, 1988.
- [OS12] C. Ortolf and C. Schindelhauer. Online multi-robot exploration of grid graphs with rectangular obstacles. In *Proc. 24th ACM Sympos. Parall. Algorithms Architectures*, pages 27–36, 2012.
- [Rap89] D. Rappaport. Computing simple circuits from a set of line segments is NP-complete. *SIAM J. Comput.*, 18:1128–1139, 1989.
- [Ric97] T. Richardson. Approximation of planar convex sets from hyperplane probes. *Discrete Comput. Geom.*, 18:151–177, 1997.
- [Rom95] K. Romanik. Geometric probing and testing—A survey. DIMACS Tech. Report 95-42, Rutgers University, 1995.
- [RW93] F. Rendel and G. Woeginger. Reconstructing sets of orthogonal line segments in the plane. *Discrete Math.*, 119:167–174, 1993.
- [Sax79] J.B. Saxe. Embeddability of weighted graphs in k -space is strongly NP-hard. In *Proc. 17th Allerton Conference on Communication, Control, and Computing*, pages 480–489, 1979.
- [SBG06] A. Sidlesky, G. Barequet, and C. Gotsman. Polygon reconstruction from line cross-sections. In *Proc. 18th Canadian Conf. Comput. Geom.*, pp. 81–84, 2006.
- [Ski88] S.S. Skiena. *Geometric Probing*. PhD thesis, University of Illinois, Department of Computer Science, 1988.
- [Ski89a] S.S. Skiena. Problems in geometric probing. *Algorithmica*, 4:599–605, 1989.
- [Ski89b] S.S. Skiena. Reconstructing graphs from cut-set sizes. *Inform. Process. Lett.*, 32:123–127, 1989.
- [Ski91] S.S. Skiena. Probing convex polygons with half-planes. *J. Algorithms*, 12:359–374, 1991.
- [Ski92] S.S. Skiena. Interactive reconstruction via geometric probing. *Proc. IEEE*, 80:1364–1383, 1992.
- [Sno99] J. Snoeyink. Cross-ratios and angles determine a polygon. *Discrete Comput. Geom.*, 22:619–631, 1999.
- [Sug94] K. Sugihara. Simpler proof of a realizability theorem on Delaunay triangulations. *Inform. Process. Lett.*, 50:173–176, 1994.
- [SVW08] S. Suri, E. Vicari, and P. Widmayer. Simple robots with minimal sensing: from local visibility to global geometry. *I. J. Robotics Res.*, 27:1055–1067, 2008.
- [Vol86] A. Volčič. A three-point solution to Hammer’s X-ray problem. *J. London Math. Soc.*, 34:349–359, 1986.
- [VZ89] A. Volčič and T. Zamfirescu. Ghosts are scarce. *J. London Math. Soc.*, 40:171–178, 1989.
- [WP16] M. Wischerhoff and G. Plonka. Reconstruction of polygonal shapes from sparse Fourier samples. *J. Comput. Appl. Math.*, 297:117–131, 2016.