

Geometric Methods for Shape Recovery from Line Drawings of Polyhedra

LLUÍS ROS AND FEDERICO THOMAS

Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Llorens Artigas 4-6, 2a Planta, 08028 Barcelona, Spain

llros@iri.upc.es

fthomas@iri.upc.es

Abstract. Using structural geometric arguments, Whiteley showed that a line drawing is a correct projection of a spherical polyhedron if and only if it has a cross-section compatible with it. We here enlarge the class of drawings to which this test applies, including those of polyhedral disks, possibly with perforations. This extension is helpful, as it makes the test applicable to verify and reconstruct drawings from usual scenes with opaque objects.

The presented results rely on geometric constructions, thus offering an alternative approach to line drawing interpretation, complementary to the algebraic-combinatorial treatment given in the classic work by Sugihara.

Keywords: shape-from-image, image understanding, line drawing interpretation, scene reconstruction, polyhedral scene

1. Introduction

Emulating the human performance in interpreting line drawings (plane projections of polyhedra) has been one of the goals of Computer Vision along the past three decades. A usual motivation behind the extensive work done in the area is helping to elucidate why humans are able to reject “impossible figures”, and recover 3D shapes from correct ones (Fig. 1), despite the reduced information they offer, without textures on the surfaces, illumination patterns, or extra views. Among the many issues to be solved, the following two have emerged as keystone problems to tackle:

- *Realizability.* Decide whether a given line drawing is *realizable*; that is, whether it is the correct projection of some 3-dimensional scene of polyhedral objects.
- *Reconstruction.* If the drawing is correct, obtain the set of its possible *reconstructions*; that is, all polyhedral scenes that project onto it. Since their number is infinite, what is asked for here is a parameterization of the whole set, in order to later select the most plausible reconstruction according to additional criteria.

Contributions to these problems range from the pioneering work by Waltz [42], Huffman [14] and Clowes [1] in the seventies, to recent advances in consistent edge labelling by Parodi et al. [24, 25], Myers and Hancock [23], and Trytten and Tuceryan [41], going through the Origami World approach by Kanade [16], the gradient-space techniques by Mackworth, Huffman and Draper [9, 15, 20], the structural geometric methods by Whiteley and Crapo [6, 43, 45] and the milestone book by Sugihara [39]. Latest developments include the extension of consistent labelling algorithms to non-polyhedral objects [2–4, 21], methods to correct incorrect pictures [13, 32] and general schemes to deal with uncertainty in the vertex locations [29, 33]. Usual applications of realizability and reconstruction techniques are to polyhedral shape-from-shading [39, Chapter 10], to 3D solid modeling from 2D free-hand sketches [11, 19] or engineering drawings [7], to compact shape representation [30], or even to toy design [40], to name a few. Chapter 3 in [28] offers a survey of the broad literature in the subject.

Overall, it is well known that the realizability problem was solved by Sugihara in his series of papers [35–38], where he reduced it to an instance of

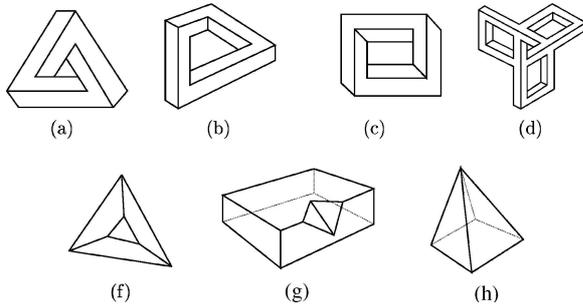


Figure 1. Some examples of line drawings. Those in the bottom row are realizable, while those at the top are not. The realizability problem is to classify a given line drawing into one of these two categories. Figure (a) is adapted from Penrose and Penrose [27], (b) from Draper [8], (c) from Huffman [14], and (d) from Ernst [10].

linear programming. However, a less known fact within the Computer Vision community, which cannot be directly inferred from this linear programming approach, is that for drawings depicting a single spherical polyhedron (showing the projections of all edges, even the hidden ones) their realizability can be decided by checking only the concurrence of groups of three lines derived from the drawing itself. (By *spherical* we mean here that the polyhedron is homeomorphic¹ to a sphere.) Although the authors independently realized this in [31], they later found that the result had already appeared in [45], due to Whiteley, who even identified a minimal number of such concurrence conditions, giving rise to the so-called *cross-section reciprocal test*, or *cross-section test* for short.² Using elementary geometric arguments, we here prove that the cross-section test is not only valid for testing drawings of spherical polyhedra, but also for those of polyhedral surfaces homeomorphic to either a disk,³ or to a disk with perforations, two classes of objects that will be hereafter referred to as *polydisks*, and *perforated polydisks*, respectively.

This constitutes a helpful extension since it makes the test applicable to drawings of scenes with opaque objects, which arise in common applications. Indeed, a necessary input to the test is the topology of the projected objects [45], or, more precisely, its *incidence structure*, as defined below. Given that hidden vertices, edges and faces are not visible on opaque objects, and that objects on top may occlude others behind, the topology of their projections is far from spherical, as shown in Fig. 2: although the drawing in the middle is a projection of the scene on the left, all we can reconstruct is a collection of surfaces such as those on the right (the visible portions from the center of projection) where each one of them is clearly homeomorphic to a polydisk, possibly perforated.

Contrarily to [45], where elaborate structural geometric arguments are used for the proof, our extension here is obtained via elementary synthetic-geometric arguments. Moreover, the given proof is constructive in nature: to verify the realizability of a drawing we explicitly construct a 3D polyhedral surface that correctly projects onto it. Following this process, one is able to obtain the parameterized infinity of all possible reconstructions, thus permitting a potential application of the result to 3D shape recovery from 2D projections.

In addition, the result and its proof provide alternative visual proofs to two important known properties of line drawings. On the one hand, their realizability is a projective invariant property: if a drawing is realizable, any projective transformation of itself is realizable too. On the other hand, when a drawing is realizable, it is so independently of the type of projection assumed—either orthogonal, parallel oblique or central.

Finally, it is worth mentioning that, since a cross-section test is a purely geometric tool, it constitutes an alternative approach to line drawing interpretation, complementary to the algebraic-combinatorial treatment of the classic work by Sugihara [39].

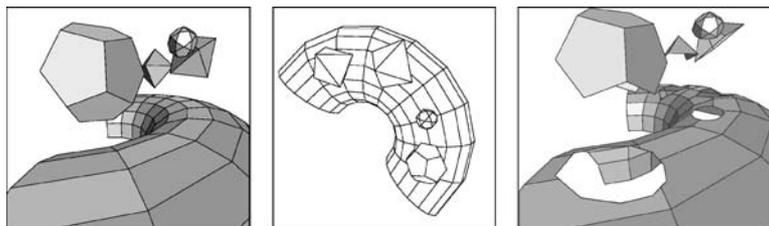


Figure 2. When a drawing (center) is a projection of an opaque scene (left), all we can reconstruct is a collection of polyhedral surfaces, each one of them topologically equivalent to a disk, possibly perforated (right).

The paper is structured as follows. Section 2 introduces some needed background and terminology. Then, Section 3 starts reviewing Whiteley’s cross-section test for spherical polyhedra, and the role of concurrence conditions to decide their realizability. This serves as an introduction to later develop the main result of this paper: the extension of the test to deal with drawings of polydisks. General scenes may be made of objects with more complicated topologies and Section 4 shows how the test still applies on perforated polydisks, but is merely a set of necessary conditions for realizability on polyhedral surfaces “with handles”, which poses an interesting open problem for future consideration. Section 5 shows how to graphically construct a cross-section for some classes of line drawings and, finally, Section 6 concludes and summarizes several points deserving further research.

2. Preliminaries

A *line drawing* is a 2D diagram made with straight line segments, called *edges*, and points where two or more segments meet, called *vertices*. Such a diagram divides the plane into several regions, called its *faces*. The term *skeleton* will be used to refer to the natural graph induced by the vertices and edges of a line drawing.

A *polyhedral surface* is a piecewise linear and continuous 2-manifold made with planar polygons, also called *faces*, glued in pairs along their edges. We allow the existence of any number of closed boundary curves on these surfaces, thus distinguishing between *boundary edges*, with just one incident face, and *interior edges*, with two incident faces. A polyhedral surface is *trihedral* if all vertices except those in the boundary have exactly three incident faces. A *spherical polyhedron*, or *polysphere* for short, is defined as a polyhedral surface homeomorphic to a sphere. Also for short, a polyhedral surface homeomorphic to a disk will be called a *polydisk*, and one homeomorphic to a disk with perforations, a *perforated polydisk*.

We say that a drawing \mathcal{D} is *correct*, or *realizable*, if we can vertically lift its vertices along the imaginary rays of projection to construct a polyhedral surface that projects onto \mathcal{D} , with distinct planes for every two adjacent faces. Such a polyhedral surface is usually called a *reconstruction*, a *spatial interpretation* or a *lifting* of the drawing.

Throughout the paper, we will also assume that the drawing is given along with its *incidence structure*, which determines the combinatorial structure of its spa-

tial interpretations—basically, which vertices will be incident to which faces. More formally, this structure is a triple $I = (V, F, R)$, where V is the set of vertices of the drawing and F is the set of its faces. We put a face in F for every subset of vertices that must be kept coplanar in the spatial interpretation. $R \subseteq V \times F$ is the *incidence set*: there is an *incidence pair* (v, f) in R if vertex v must lie on face f in 3-space. The incidence structure can be computed by applying the methods by Sugihara in [39, p. 45], after a consistent labelling of its edges has been obtained. Several techniques exist to obtain consistent labellings, like those by Huffman [14], Waltz [42], Hancock et al. [12,23] or Parodi et al. [25,26], to name a few. Since finding a consistent edge labelling is an NP-complete problem [17], the previous assumption, that the drawing is given with an identified incidence structure, implicitly means in all results below that a consistent edge-labelling has been performed yet. The realizability problem will thus be understood as that of deciding whether a pair “drawing plus incidence structure” has some plausible spatial interpretation.

The following definitions and a related lemma will be useful in the proof of Theorem 1 below. Let D be a subset of the plane homeomorphic to a disk. A *disk-partition* of D is a collection of regions, each homeomorphic to a disk, such that any two regions can only intersect along their boundaries and the union of all regions is D . Now, select m regions r_1, \dots, r_m of a disk-partition, such that the union $D' = r_1 \cup \dots \cup r_m$ is homeomorphic to a disk. We say that another region r of the disk-partition is *continuously adjacent* to D' if $r \cap D'$ is a continuous curve segment with more than one point. (Note that, under this circumstances, the union $r \cup D'$ must be homeomorphic to a disk too.) Furthermore, let P be a disk partition of D , a *disk-growing sequence* of P is an ordering of the regions in P into a sequence s so that every region in s is continuously adjacent to the union of all regions preceding it in s . Then, the following result holds:

Lemma 1. *Let D be a subset of the plane homeomorphic to a disk, and let P be a disk-partition of D . Then, for any region r of P there is at least one disk-growing sequence whose first element is r .*

Proof: A disk-growing sequence starting at r can be easily obtained as follows. (1) Mark r as visited and put it in a list L , initially empty. (2) Obtain the union D' of all regions in L . (3) Choose an arbitrary non-visited

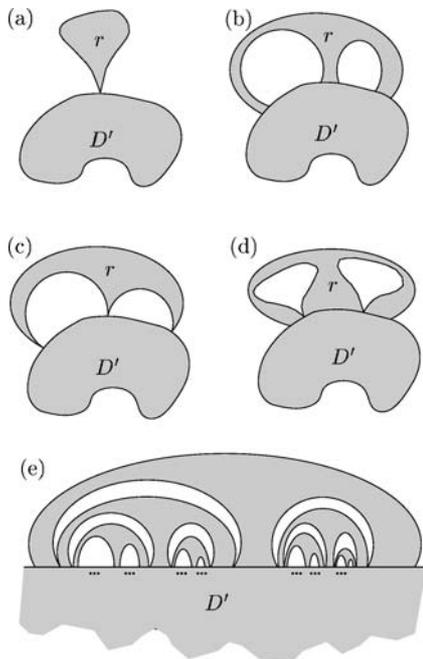


Figure 3. (a, b, c) Possible shapes of a region r in contact with D' . (d) An impossible shape for region r . (e) A possible covering of the boundary of D' with infinitely many regions of the type in (b), recursively placed inside their holes.

region continuously adjacent to D' , mark it as visited, and add it to the tail of L . (4) Repeat step 2 until no progress is possible.

At the end of this process, L will certainly contain all regions in D and, hence, the desired disk-growing sequence. To prove this, assume contrarily that the algorithm ends, but L does not contain all regions in P . This means that, upon termination, any region r not in L that is in contact with D' is not continuously adjacent to D' ; Otherwise the algorithm could progress through it. This implies that either,

- r is in contact with D' at just one point (Fig. 3(a)), or
- r is in contact with D' along p disjoint curve segments, creating $p - 1$ holes. (Figure 3(b) shows the case $p = 3$.) Observe that, while some of these segments may be isolated points (as in Fig. 3(c)), a hole, on the contrary, cannot contact the boundary of D' at a single point only (as in Fig. 3(d)), since r must be homeomorphic to a disk.

The only way to prevent the algorithm from progressing is, thus, to cover the boundary of D' with regions of any

of these two types. This is only possible, though, by using an infinite number of such regions, as illustrated in Fig. 3(e), which violates the premise that the number of regions in P is finite. Hence the lemma. \square

Remark 1. If D is not a subset of the plane homeomorphic to a disk, but is instead a subset of 3-space homeomorphic to a sphere, then we can analogously define the concepts of a disk partition of D , and of a disk-growing sequence for such partition. Then, clearly, the previous lemma also holds if P is a disk-partition of a surface D homeomorphic to a sphere.

For the sake of simplicity, we will continue by assuming that the drawing, if correct, is produced by orthogonally projecting a single polyhedral surface onto the XY plane. Proposition 3 below will extend the results to other types of projection, and [28, Section 2.3] shows how they are readily extended to drawings depicting several polydisks, possibly perforated.

3. The Cross-Section Test

Some examples suggest that a drawing's correctness is checkable solely using concurrence conditions. The truncated tetrahedron in Fig. 4(a) is only correct when its three edges l , m , and n meet at a common point. The 4-calotte in Fig. 4(b), a configuration of a quadrilateral face and its four neighboring faces, is only correct when the three lines l , m , n are concurrent or, equivalently, when the three bold points are aligned, since they all lie in the line m of intersection of the planes α and β . Likewise, the 5-calotte in Fig. 4(c) is only correct when l , m_1 , n_1 , and l , m_2 , n_2 are concurrent too. These are all necessary conditions for realizability since they hold in any spatial reconstruction of the drawing (Fig. 4, center) and projection preserves incidence relations between points and lines. Although these conditions are already useful to discard some incorrect drawings (Fig. 4, right), the challenge was to characterize a set of concurrence conditions also sufficient for realizability. Whiteley's cross-section test, which is next described, identifies one such set for spherical polyhedra.

Consider a spherical polyhedron in 3-space such as, for example, the tetrahedron in Fig. 5, top. Now, obtain the intersections of the planes of its faces with an external plane ϕ in a general position, not coincident with any other face plane. The resulting arrangement of lines is called a *cross-section* of the polyhedron. It is

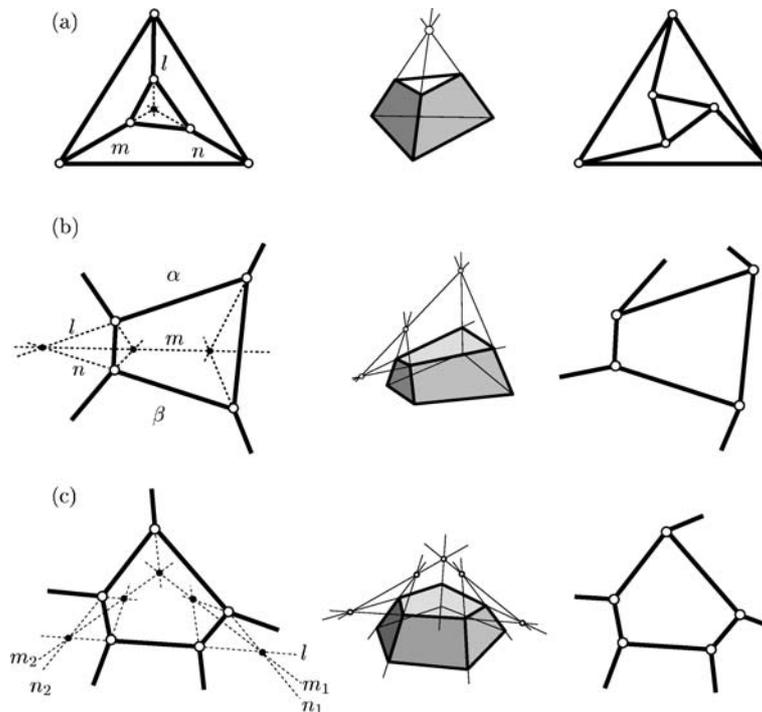


Figure 4. To be correct, the drawings on the left must verify the indicated concurrence conditions, since they hold on any of their spatial interpretations (center). These conditions allow to classify the drawings on the right as incorrect. The incidence structure implicitly considered in each example is that of its corresponding spatial reconstruction shown in the center.

clear that the edge line e between two faces, say f_i and f_j , must be concurrent to the point of intersection of the lines L_{f_i} and L_{f_j} of intersection of the planes of f_i and f_j with ϕ , respectively. These trivial concurrence conditions that hold in 3-space will clearly hold too when projecting the whole construction onto the plane ϕ , because projection preserves collinearity of points and all incidence relations (Fig. 5, central). Hence, we have a set of concurrence constraints that are necessary for a drawing to correctly represent the projection of a spherical polyhedron: the drawing can only be correct if we are able to draw a *compatible* cross-section diagram, one where these concurrences hold (Fig. 5, bottom). Whiteley's theorem states that the converse is also true: if the concurrences hold, a spherical polyhedron can be reconstructed from the drawing.

We will also prove the following extension for polydisks, where the word "compatible" now has a slightly different meaning: that only the interior edges of the polydisk must be concurrent to their corresponding point in the cross-section (leaving the boundary ones unconstrained). Proposition 3 below will extend its applicability to perforated polydisks.

Theorem 1 (Cross-section test for polydisks). *Let \mathcal{D} be a line drawing of a polydisk. \mathcal{D} is realizable if and only if it has a compatible cross-section such that the cross-section lines L_{f_i} and L_{f_j} of every two adjacent faces f_i and f_j are non-coincident.*

Proof: (\Leftarrow) We will prove the "if" part using the drawing \mathcal{D} and its compatible cross-section to explicitly construct a polydisk \mathcal{L} ; one whose face planes generate the given cross-section when intersected with the plane of the drawing.

Let us first assume that all faces of \mathcal{D} are topological disks. (The existence of other face types will pose no extra difficulty, as shown below.) Now, consider a subset R of the plane, homeomorphic to a disk, and embed the skeleton of \mathcal{D} in it, so that the regions of R induced by the embedding are in one-to-one correspondence with the faces of \mathcal{D} . The induced regions define a disk-partition P of R , and Lemma 1 can be applied to find a disk-growing sequence s of P . The faces of \mathcal{D} can be ordered then, according to the order in which their corresponding regions of P appear in s . Let f_1, f_2, \dots, f_n be this ordering.

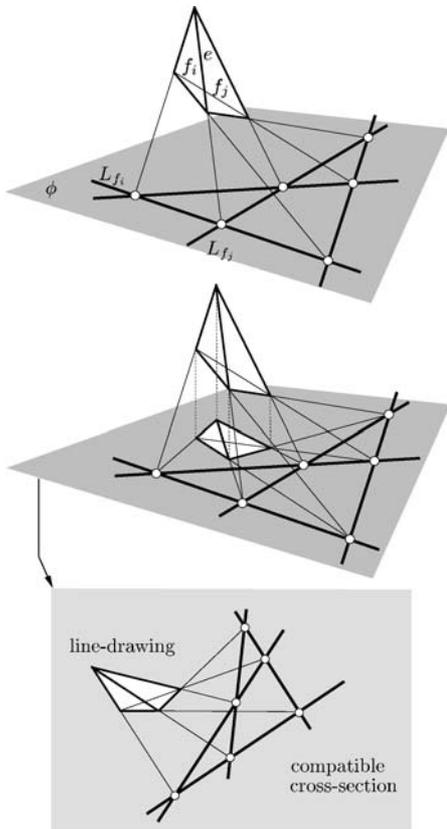


Figure 5. The cross-section test.

The proof now proceeds using induction on the sequence f_1, \dots, f_n . To start with, take f_1 and consider its cross-section line. A lifting of this face can be fixed by giving an arbitrary height to any one of its vertices not in the cross-section line. The cross-section line and the lifted vertex define the plane of f_1 , and all edges and vertices of f_1 may then be lifted vertically to lie on this plane.

Now, as induction hypothesis, we assume that a polydisk \mathcal{L}_{k-1} made up with f_1, \dots, f_{k-1} has already been correctly lifted to 3-space. We will then prove that f_k can be properly lifted too, to form a lifted polydisk \mathcal{L}_k with k faces.

Observe that, since \mathcal{L}_{k-1} has been correctly lifted, the edge-line between any two of its adjacent faces will meet the point where the cross-section lines of these two faces intersect. This applies to all interior edges of \mathcal{L}_{k-1} , but also to those on its boundary that do not lie on the boundary of \mathcal{L} .

Let us now lift f_k . Given the way f_k was chosen, it is clear that it will be adjacent to p

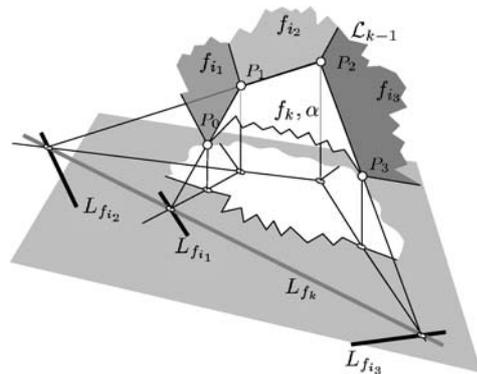


Figure 6. Lifting an intermediate face.

faces of \mathcal{L}_{k-1} , say f_{i_1}, \dots, f_{i_p} , through q edges $P_0P_1, P_1P_2, \dots, P_{q-1}P_q$, and that these edges will be sequentially linked, forming a continuous polygonal line. (This is illustrated in Fig. 6, where f_k is adjacent to $p = 3$ faces of \mathcal{L}_{k-1} , through $q = 3$ edges.) We will prove that this polygonal line and the cross-section line L_{f_k} of f_k are coplanar, and define a plane α that is different from all planes assigned to the faces f_{i_1}, \dots, f_{i_p} . This will be true even if any one of the faces f_{i_1}, \dots, f_{i_p} is in contact with this polygonal line through more than one edge.

To see this, note that, as the edge line P_0P_1 is incident with L_{f_k} (by induction hypothesis), P_0P_1 and L_{f_k} are coplanar. Let us call α the plane they define. Moreover, the edge line P_1P_2 is also coplanar with α , as it contains two points of this plane: the point where it intersects with L_{f_k} and the point P_1 of line P_0P_1 . The same applies to line P_2P_3 as it is incident with P_2 and L_{f_k} . Clearly, the argument can be iterated to prove that all other edges between f_k and faces of \mathcal{L}_{k-1} are coplanar with α .

With the plane for f_k already fixed, all other edges between f_k and faces not in \mathcal{L}_{k-1} can be fixed too by lifting them vertically to lie on this plane. For every such edge, say between faces f_k and f_i , we must prove that its line of support meets the point T where the cross-section lines L_{f_i} and L_{f_k} meet. Clearly, this line of support is defined by the intersection of the plane α and a vertical plane containing the projection of the edge. But both planes meet T : α meets T because T is a point of L_{f_k} , and the vertical plane meets T because the projection of the edge meets T in the cross-section.

It remains to prove that the plane α is different from all the planes given to the faces f_{i_1}, \dots, f_{i_p} . But this is trivially true, as the only way for α to coincide with one of such planes would be that their corresponding

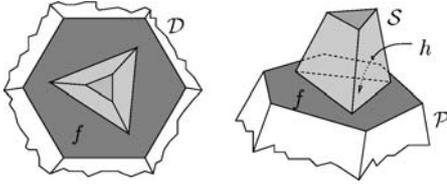


Figure 7. A drawing \mathcal{D} with a face f not homeomorphic to a disk (left) and its spatial interpretation (right) regarded as a polysphere S glued to a polydisk \mathcal{P} through face h .

cross-section lines were identical, which is not the case by the premises of the theorem.

Finally, assume that some face f of \mathcal{D} is not homeomorphic to a disk and has, for example, one hole. Then, the arrangement of faces interior to that hole can always be regarded as a polysphere S adhered to face f —thought of as with no holes now—through an imaginary face h whose contour coincides with that of the hole (Fig. 7). Assume for the moment that all faces of S are homeomorphic to a disk. The drawing \mathcal{D} may then be viewed as a compound object: a polysphere S glued to a polydisk \mathcal{P} through its face h , both objects with all faces homeomorphic to a disk. A lifting of \mathcal{D} can then be obtained in two stages: (1) Using a disk-growing sequence for \mathcal{P} , we obtain a lifting of the faces in \mathcal{P} . This fixes a plane for face f in 3-space. (2) We generate a disk-growing sequence for S , starting at f , and use it to lift all faces of S to 3-space. The generalization of this process to drawings with recursively many polyspheres adhered to S , and to faces with more than one hole is straightforward.

(\Rightarrow) Conversely, if the line drawing can be lifted to a polydisk, then we can construct a cross-section just by extending the face planes of the lifted polydisk and intersecting them with the plane of the drawing. The resulting lines define a compatible cross-section since, clearly, any interior edge of the spatial polydisk, when extended, will be incident to the intersection of the cross-section lines of its two adjacent faces, and the same concurrence will hold when extending the projections of these interior edges. \square

By following an analogous reasoning, we can come up with a synthetic-geometric proof of the cross-section test for spherical polyhedra, as follows.

Theorem 2 (Whiteley, 1991). *Let \mathcal{D} be a line drawing of a spherical polyhedron. \mathcal{D} is realizable if and only if it has a compatible cross-section such that the cross-*

section lines L_{f_i} and L_{f_j} of every two adjacent faces f_i and f_j are non-coincident.

Proof: We can embed the skeleton of \mathcal{D} onto a sphere, taking care that the delimited regions on the sphere are in one-to-one correspondence with the faces of the drawing. This defines a disk-partition P of the sphere and, thus, by Remark 1, it is possible to find a disk-growing sequence of P . Following this sequence, it is possible to find an ordering of the faces of \mathcal{D} , say f_1, \dots, f_n , so that the boundary of each f_i intersects the boundary of the union $f_1 \cup \dots \cup f_{i-1}$ along a continuous polygonal segment. The previous proof can be applied then, using induction on the number of lifted faces of the sequence f_1, \dots, f_n . \square

So far, the cross-section test has only been proved under orthogonal projection. However, its validity is easily extended to any type of projection due to the following known property.

Proposition 1 (*Independence of the assumed projection*). *A drawing is realizable under orthogonal projection if and only if it is realizable under any central or parallel oblique projection.*

Proof: A proof for general drawings was given by Sugihara in [39, Theorem 3.4]. If the drawing is a projected polydisk, though, this result can be easily deduced from the proof of Theorem 1 above. If we follow it while observing Fig. 6, we see that the same proof works for central or parallel oblique projection, the only difference being that the vertices P_0, P_1, \dots, P_q , are lifted along lines that meet at the center of projection in the former case, or at a point at infinity in a direction oblique to the plane in the latter. Clearly, if the drawing is a projected polysphere, the same argument applies. \square

We close the section by noting that Theorems 1 and 2 also allow a visual proof to another important property of line drawings.

Proposition 2 (*Projective invariance of realizability*). *Let \mathcal{D} and \mathcal{D}^* be two line drawings with the same incidence structure and such that \mathcal{D}^* has been obtained from \mathcal{D} by applying a non-singular projective transformation to its vertices and edges. Then, \mathcal{D} is realizable if and only if \mathcal{D}^* is realizable.*

Proof: This was proved by Whiteley and Crapo [5, 43] for more general drawings. However, if \mathcal{D} is a projected polydisk or polysphere, this can be readily seen as a consequence of Theorems 1 and 2. Indeed, since non-singular projective transformations of the plane map lines to lines and points of intersection of two lines to the points of intersection of the transformed two lines [18, page 92], if a drawing \mathcal{D} has a compatible cross-section, the transformed drawing \mathcal{D}^* will also have one, obtained by transforming the lines of the original. \square

4. Realizability for Other Topologies

So far, the sufficiency of the cross-section test has been proved for drawings of polyhedral surfaces homeomorphic to a disk (Theorem 1) or a sphere (Theorem 2). Can we extend the test to surfaces with other topologies? The following considerations will depict its full range of applicability. They are summarized in Table 1.

First, there is no trouble in dealing with self-intersecting surfaces, or with self-intersecting faces (Fig. 8) if the underlying incidence structure still has the topology of a disk or a sphere. Thus, as long as

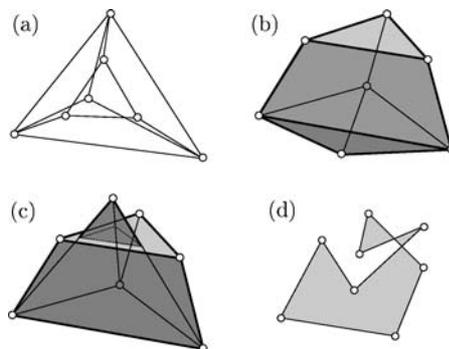


Figure 8. (a) A line drawing whose incidence structure has the topology of a sphere. (b and c) Two spatial reconstructions of (a), one of them with self-intersections. (d) A self-intersecting face.

these intersections do not produce any new faces, vertices, or edges, all of the proofs above will be valid for these objects. From now on, hence, we will concentrate on the topology of the combinatorial structure, rather than that of the spatial object itself.

At a topologic level, an orientable surface S can be fully characterized by two quantities: the number μ of its closed boundary curves, and the maximum number p of closed cuts that do not separate it into parts—also called the *genus* of S . In fact, a necessary and

Table 1. Applicability of cross-sections for several topologies.

$C_{0,0}$	0	0		Yes, with Theorem 2.
$C_{\mu,0}$ $\mu \geq 1$	1	0		Yes, with Theorem 1.
	2	0		Yes, after triangulating $\mu - 1$ holes.
	
$C_{0,p}$ $p \geq 1$	0	1		No. Only a necessary condition.
	0	2		
	
$C_{\mu,p}$ $\mu \geq 1$ $p \geq 1$	1	1		No. Only a necessary condition. But reducible to $C_{0,p}$ by triangulating the boundaries.
	1	2		
	

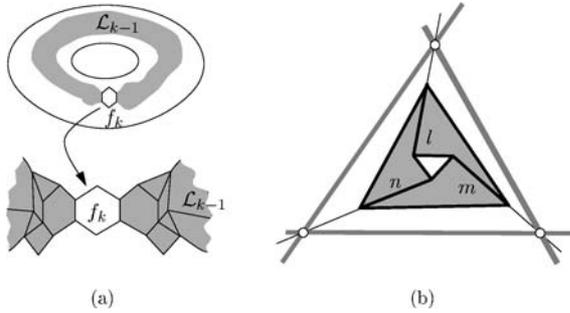


Figure 9. The cross-section test is invalid for objects other than polyspheres and polydisks. (a) If the polyhedral surface has perforations, then there will be some face f_k not continuously adjacent to a previously lifted polydisk \mathcal{L}_{k-1} . (b) An incorrect drawing of a perforated polydisk, with a compatible cross-section.

sufficient condition for two orientable surfaces to be homeomorphic is that these two numbers shall be the same for both surfaces [18, p. 106].

Note that, if $C_{\mu,p}$ denotes the class of all orientable homeomorphic surfaces with genus p and μ boundaries, what we have done so far is to prove the validity of the cross-section test for surfaces in the classes $C_{0,0}$ and $C_{1,0}$.

For polyhedral surfaces in $C_{\mu,0}$, $\mu \geq 2$ (perforated polydisks), the cross-section test cannot be directly applied. For such objects, the proof of Theorem 1 would fail, as their drawings cannot be lifted by iteratively adding faces that are continuously adjacent to a previously lifted polydisk. This can be seen with the help of Fig. 9(a), for a topologic disk with one perforation. As a specific counterexample, consider three pairwise adjacent faces with a triangular perforation (Fig. 9(b)): this drawing always has a compatible cross-section, but it is not realizable unless the three interior edges, l , m and n , are concurrent. A useful modification, however, allows the use of cross-sections even in such cases.

Proposition 3 (*Cross-section test for perforated polydisks*). *If \mathcal{D}_1 is a drawing of a perforated polydisk \mathcal{P} , we can triangulate every perforation to produce a new drawing \mathcal{D}_2 whose incidence structure has the topology of a disk. Then, \mathcal{D}_1 is realizable if and only if \mathcal{D}_2 is realizable.*

Proof: If \mathcal{D}_1 is realizable, it has at least one lifting \mathcal{L}_1 , and a lifting of \mathcal{D}_2 can be found by spatially triangulating the perforations of \mathcal{L}_1 . Conversely, a lifting of \mathcal{D}_1 can be found from one of \mathcal{D}_2 by removing the triangles that cover the perforations. This argument also

shows that, actually, there is a one-to-one correspondence between the liftings of \mathcal{D}_1 and those of \mathcal{D}_2 . \square

In sum, the cross-section test can be used for objects in the class $C_{\mu,0}$, $\mu \geq 2$, as long as they are converted to the class $C_{1,0}$, by the triangulation of $\mu - 1$ boundaries.

No synthetic geometric test using cross-sections has been found for the rest of topologic objects with at least one handle ($C_{\mu,p}$, $\mu \geq 0$, $p \geq 1$). Actually, Crapo and Whiteley prove in [6] that the cross-section test is not valid for checking drawings of a prismatic torus, made up of three triangular prisms glued together in pairs through their bases. They show it can have a compatible cross-section that does not correspond to a correct spatial lifting of this torus.

5. Graphical Construction of Cross-Sections

Although the existence of a compatible cross-section is a necessary and sufficient condition for realizability, we still need some process to find all compatible cross-sections or show that none exists. For drawings of trihedral polyspheres or polydisks we give a method in Section 5.1 below, called the *incremental construction*. As explained in Section 5.2, this construction can also be used on perforated trihedral polydisks, provided that their inner boundaries have no more than four “relevant” edges. This and the trihedrality limitation make these methods unapplicable on general drawings, posing some open problems that will be discussed in Section 5.3.

5.1. Trihedral Polyspheres and Polydisks

To obtain the incremental construction, first note that, when lifting a correct drawing, one can always choose the heights of four vertices independently; namely, those of the two vertices of an interior edge e , and the height of one vertex in each of the adjacent faces of e . Thus, these two faces can receive arbitrary planes, provided that their intersection line projects onto their edge. This means that, when constructing a cross-section, the lines of two adjacent faces can be chosen with arbitrary orientation, as long as they are different and concurrent to a same point on their common edge line. For drawings of trihedral polyspheres and polydisks these two initial lines completely determine the rest of the cross-section. We see this in Fig. 10(a), where, after fixing the lines L_{f_1} and L_{f_2} for faces f_1 and

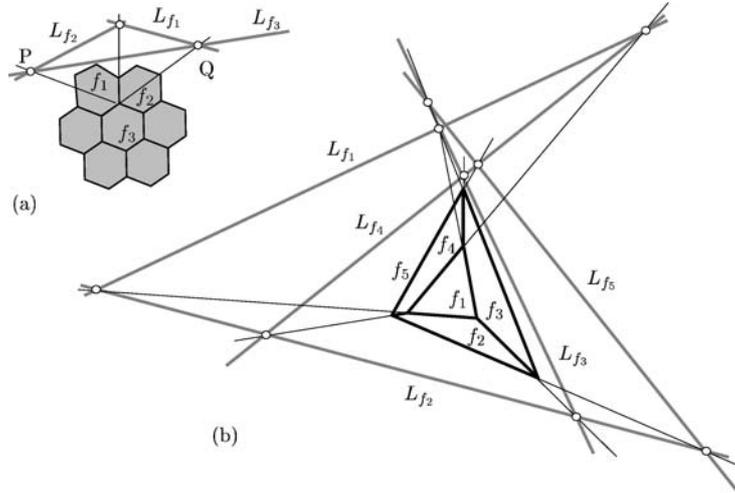


Figure 10. Testing a trihedral drawing with pencil and straightedge.

f_2 , L_{f_3} can be automatically deduced as it must contain P and Q , the points where the edge lines between f_3 and f_1 , and f_3 and f_2 meet with the cross-section lines of f_1 and f_2 , respectively. We can obtain lines for the other faces iterating this process, and construct the whole cross-section. If at some point a concurrence condition does not hold, we conclude that the drawing is incorrect. If we complete the construction, then it is correct by Theorems 1 or 2. Clearly, this algorithm takes a linear time in the number of faces, since it adds one cross-section line at a time, and each line is determined by two previously obtained points. This proves the following.

Proposition 4. *Let \mathcal{D} be a line drawing of a trihedral polysphere or a polydisk, with an incidence structure identified on it. Then, the realizability of \mathcal{D} can be checked with pencil and straightedge in linear time in the number of faces.*

Figure 10(b) illustrates the incremental construction on a truncated tetrahedron. We consider the outer triangle as a face too and, hence, the incidence structure is that of a spherical polyhedron. We start fixing L_{f_1} and L_{f_2} , with arbitrary orientation, making them concurrent to the edge line between faces f_1 and f_2 . The rest of lines are then fixed in this order: L_{f_3} , L_{f_4} , L_{f_5} .

5.2. Trihedral Perforated Polydisks

As shown in Section 4, when we have perforations, we can simply triangulate them and apply Theorem 1 to

the newly derived drawing. Nevertheless, after adding the triangles, the new drawing will not be fully trihedral, thus making the incremental construction probably unsuitable to obtain its possible cross-sections. The following considerations, though, delimit some cases where it is still a valid tool.

Let \mathcal{D} be a line drawing of a trihedral polydisk with perforations. A 2-valent vertex of \mathcal{D} is one that is incident to exactly one face, thus having two incident edges. We realize that every 2-valent vertex lying on the boundary of a perforation can be safely removed, and its two adjacent vertices directly linked by an edge, without altering the realizability of \mathcal{D} (Fig. 11(a)). We will say that \mathcal{D} has been *regularized* when all 2-valent vertices of this kind have been removed, and will denote by $Reg(\mathcal{D})$ the resulting new drawing. Clearly, \mathcal{D} is realizable if and only if $Reg(\mathcal{D})$ is, as every lifting of \mathcal{D} can be easily converted to one of $Reg(\mathcal{D})$, and vice versa. The following holds:

Proposition 5. *If \mathcal{D} is a drawing of a perforated polydisk, and each perforation of $Reg(\mathcal{D})$ has three or four boundary edges, then the correctness of \mathcal{D} can be checked with the incremental construction.*

Proof: To check the realizability of \mathcal{D} we proceed as follows. First, we simplify the perforations as much as possible by regularizing \mathcal{D} . Then, the idea is to construct a “trihedral roof” over each perforation of $Reg(\mathcal{D})$, until we get a new drawing with the topology of a trihedral polydisk. We act

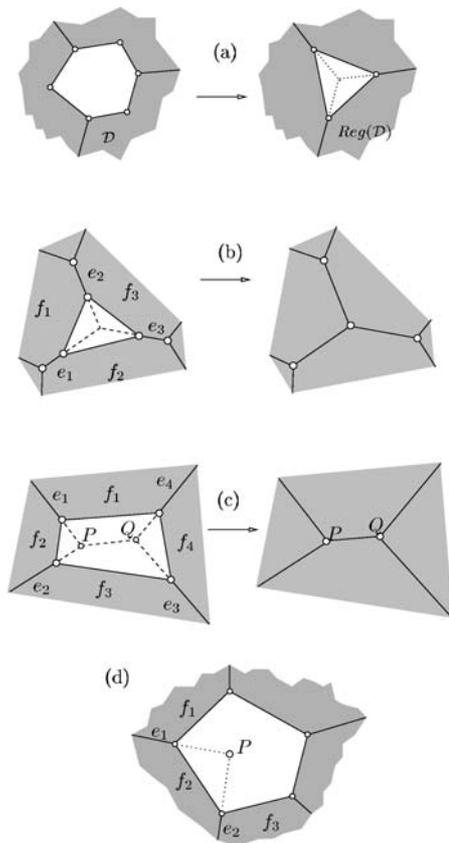


Figure 11. Testing perforated polydisks.

differently, depending on the number of edges of these perforations:

- If we have a triangular perforation, we proceed as in Fig. 11(b). We cover the perforation by extending its three faces, f_1 , f_2 and f_3 , until the common point of intersection of the edge lines e_1 , e_2 and e_3 . If these lines do not meet at a common point we can reject the drawing as incorrect. If they meet, we get a transformed drawing that is realizable if and only if the original one is.
- If we have a quadrilateral perforation, we proceed as in Fig. 11(c). The point P of intersection of the faces f_1 , f_2 and f_3 must lie on the intersection of the edge lines e_1 and e_2 . Analogously the point where f_1 , f_3 and f_4 meet is Q , on the intersection of e_3 and e_4 . Clearly, PQ is the line of intersection of f_1 with f_3 . The perforation is covered by extending f_1 , f_2 , f_3 and f_4 until they hit the dotted lines. Again, the resulting drawing is realizable if and only if the original one is.

Once all perforations have been covered using these two transformations, we end up with a drawing of a trihedral polydisk and, thus, by Proposition 4, it can be checked with the incremental construction. \square

Unfortunately, the construction of trihedral roofs is not possible on perforations with more than four boundary edges, because the intersection lines between pairs of faces in these roofs cannot be determined from the edge lines incident to the perforation. In the pentagonal perforation of Fig. 11(d), for example, one point of the intersection line between faces f_1 and f_3 is P , where e_1 and e_2 meet. To determine another point on this roof line we would need an additional face simultaneously adjacent to f_1 and f_3 , as it happened with f_4 in Fig. 11(c), but none is available.

5.3. General Drawings

The incremental construction is possible on drawings of trihedral polydisks and trihedral polyspheres precisely because they have a *determined cross-section*, one where after the initial choice of two lines, the remaining ones are fully determined. In fact, it would be useful to characterize the whole class of drawings with determined cross-sections, to delimit the full range of applicability of the incremental construction. Note that this class is not only restricted to drawings of trihedral polydisks and polyspheres. The reader can check, for example, that the two projected polyspheres in Fig. 12 also exhibit this property, but they have vertices with more than three incident faces. To our knowledge, this characterization has not been pursued yet and remains an interesting open problem.

Moreover, if we try to apply the incremental construction on a general drawing, it may happen that an intermediate line remains undetermined. If this is the

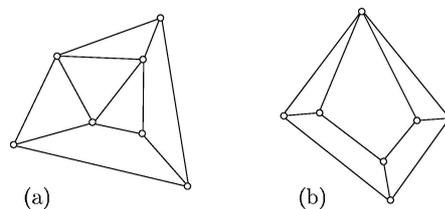


Figure 12. Non-trihedral drawings with determined cross-sections. Both drawings must be thought of as “projected polyspheres”, that is, with the outer contour forming a face too. While (a) is realizable independently of the (x, y) positions assigned to its vertices, (b) is not.

case, we can always choose one among all possible positions for this line, and continue with the construction. If we are able to complete a cross-section, the drawing is certainly correct. If we do not succeed, though, we cannot reject it as incorrect, since other locations for the undetermined lines might still yield a compatible cross-section. Obviously, we could try all possibilities in a generate-and-test fashion with backtracking, but this is clearly non-viable. Unfortunately, no general method has been devised yet to generate all compatible cross-sections, or otherwise show that none exists, using pencil and straightedge alone. This constitutes an additional challenging open problem (if solvable at all).

6. Conclusions

Traditionally, the Machine Vision approach to line drawing interpretation has been mainly algebraic. Although computer scientists have discovered graphical techniques like the dual diagrams of the *gradient-space approach* [9, 15, 20], these have not been fully exploited or sometimes left aside, arguing that they only provide necessary (but not sufficient) conditions for realizability. However, a careful investigation of results from related areas of Geometry has revealed the existence of complete and purely geometric tests to decide the correctness of whole families of drawings. The goal of this paper has been to further exploit this geometric side to obtain not only necessary, but also sufficient conditions for drawings of scenes with opaque polyhedra or, more generally, for drawings whose incidence structure is that of a polyhedral disk, possibly with one or more perforations.

These results can be extended in a number of ways. On the one hand, we note that the cross-section is just one among several known *reciprocal diagrams*. Other related diagrams have been found and used for the same purposes. For example, the dual diagram of the gradient-space approach is actually the same as Maxwell's reciprocal [22], and the known fact that Maxwell's reciprocal can be transformed to the cross-section through a plane polarity [45] indicates that all these diagrams are essentially the same, up to projective transformations. It would be helpful to clarify all the equivalences with a unifying aim in mind, since the readers usually have the impression that many authors use the same concepts under a different language, when following the classical literature [9, 15, 16, 20].

On the other hand, the main drawback of cross-sections is that no general method has been given yet to construct a cross-section or otherwise show that none exists, using a pencil and a straightedge alone. Although we have shown an incremental construction that works for trihedral drawings, a general tool still remains unknown. Its development, or a proof of its non-existence, are challenging open problems for further consideration.

Finally, it is worth mentioning that on perspective drawings of objects with several families of parallel edges, knowing the locations of the vanishing points of such edges may provide additional information to partially or fully constrain the position and orientation of all lines in a cross-section—and, by extension, limit the number of possible spatial interpretations. In fact, although working out the details is out of the scope of the present paper, it is not difficult to prove that the horizon lines of the face planes of the projected object in a drawing must be parallel to their corresponding cross-section lines, and that a projected edge between two faces f_i and f_j must not only be concurrent to the intersection of the cross-section lines L_{f_i} and L_{f_j} , but also to the point where the horizon lines of f_i and f_j meet. To fully unravel this connection seems a worthwhile task, given the availability of fast methods that detect the vanishing point positions (and hence the horizon lines) in perspective images [34]. This point currently concentrates part of our research efforts.

Acknowledgments

The authors learnt of the cross-section test during a stay of the first author with Walter Whiteley in Canada. They are very grateful to him, as many of the ideas behind the geometric approach presented here arose in the conversations with him. Also, the authors are very grateful to the reviewers of this paper, for their useful comments and all suggested improvements.

This work has been partially supported by the Spanish CICYT under contract TIC2000-0696, by a Ph.D. fellowship granted by the Spanish Ministry of Education with reference number FPI94-46634232, and by a Ramón y Cajal contract funding the first author, granted by the Spanish Ministry of Science and Technology.

Notes

1. "Homeomorphic" means "topologically equivalent".
2. Actually, Whiteley already conjectured the test to be valid in [44].

3. Throughout the paper, the term “disk” refers to a “closed disk”, i.e., the set of points (x, y) of the plane such that $x^2 + y^2 \leq 1$.

References

1. M. B. Clowes, “On seeing things,” *Artificial Intelligence*, Vol. 2, No. 1, pp. 79–116, 1971.
2. M. C. Cooper, “Interpreting line drawings of curved objects with tangential edges and surfaces,” *Image and Vision Computing*, Vol. 15, pp. 263–276, 1997.
3. M. C. Cooper, “Linear-time algorithms for testing the realisability of line drawings of curved objects,” *Artificial Intelligence*, Vol. 108, pp. 31–67, 1999.
4. M. C. Cooper, “Linear constraints for the interpretation of line drawings of curved objects,” *Artificial Intelligence*, Vol. 119, pp. 235–258, 2000.
5. H. Crapo and W. Whiteley, “Statics of frameworks and motions of panel structures, a projective geometric introduction,” *Structural Topology*, Vol. 6, pp. 42–82, 1982.
6. H. Crapo and W. Whiteley, “Plane self stresses and projected polyhedra I: The basic pattern,” *Structural Topology*, Vol. 20, pp. 55–78, 1993.
7. D. Dori and K. Tombre, “From engineering drawings to 3D CAD models—Are we ready now?,” *Computer Aided Design*, Vol. 27, No. 4, pp. 243–254, 1995.
8. S. W. Draper, “The Penrose triangle and a family of related figures,” *Perception*, Vol. 7, pp. 283–296, 1978.
9. S. W. Draper, “The use of gradient and dual space in line-drawing interpretation,” *Artificial Intelligence*, Vol. 17, pp. 461–508, 1981.
10. B. Ernst, *Adventures with Impossible Figures*, Parkwest Publications, 1987.
11. I. J. Grimstead and R. R. Martin, “Incremental line labelling for sketch input of solid models,” *Computer Graphics Forum*, Vol. 15, No. 2, pp. 155–166, 1996.
12. E. R. Hancock and J. Kittler, “Edge labelling using dictionary-based relaxation,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 12, pp. 165–181, 1990.
13. A. Heyden, “On the consistency of line-drawings, obtained by projections of piecewise planar objects,” *Journal of Mathematical Imaging and Vision*, Vol. 6, No. 4, pp. 393–412, 1996.
14. D. A. Huffman, “Impossible objects as nonsense sentences,” *Machine Intelligence*, Vol. 6, pp. 295–323, 1971.
15. D. A. Huffman, “A duality concept for the analysis of polyhedral scenes,” *Machine Intelligence*, Vol. 8, pp. 475–492, 1977.
16. T. Kanade, “A theory of Origami World,” *Artificial Intelligence*, Vol. 13, pp. 279–311, 1980.
17. L. Kirousis and C. Papadimitriou, “The complexity of recognizing polyhedral scenes,” *Journal of Computer and System Sciences*, Vol. 37, No. 1, pp. 14–38, 1988.
18. F. Klein, *Elementary Mathematics from An Advanced Standpoint*. Dover Publications, 1939.
19. H. Lipson and M. Shpitalni, “Optimization-based reconstruction of a 3D object from a single freehand line drawing,” *Computer-Aided Design*, Vol. 28, No. 8, pp. 651–663, 1996.
20. A. K. Mackworth, “Interpreting pictures of polyhedral scenes,” *Artificial Intelligence*, Vol. 4, pp. 121–137, 1973.
21. J. Malik and D. Maydan, “Recovering three-dimensional shape from a single image of curved objects,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 11, No. 6, pp. 555–566, 1989.
22. J. C. Maxwell, “On reciprocal figures and diagrams of forces,” *Philosophical Magazine*, Vol. 27, No. 4, pp. 250–261, 1864.
23. R. Myers and E. R. Hancock, “Genetic algorithms for ambiguous labelling problems,” *Pattern Recognition*, Vol. 33, pp. 685–704, 2000.
24. P. Parodi, “The Complexity of Understanding Line-Drawings of Origami Scenes,” *International Journal of Computer Vision*, Vol. 18, No. 2, pp. 139–170, 1996.
25. P. Parodi, R. Lancewicki, A. Vijn, and J. K. Tsotsos, “Empirically-derived estimates of the complexity of labelling line drawings of polyhedral scenes,” *Artificial Intelligence*, Vol. 105, Nos. (1/2), pp. 47–75, 1998.
26. P. Parodi and V. Torre, “On the complexity of labelling perspective projections of polyhedral scenes,” *Artificial Intelligence*, Vol. 70, Nos. (1/2), pp. 239–276, 1994.
27. L. S. Penrose and R. Penrose, “Impossible objects: A special type of visual illusion,” *British Journal of Psychology*, Vol. 49, pp. 31–33, 1958.
28. L. Ros, “A kinematic-geometric approach to spatial interpretation of line drawings,” Ph.D. thesis, Technical University of Catalonia, 2000. Available at <http://www-iri.upc.es/people/ros>.
29. L. Ros, A. Sabater, and F. Thomas, “An ellipsoidal calculus based on propagation and fusion,” *IEEE Trans. on Systems, Man and Cybernetics (part B)*, Vol. 32, Nos. 4, pp. 430–442, 2002.
30. L. Ros, K. Sugihara, and F. Thomas, “Towards shape representation using trihedral mesh projections,” *The Visual Computer*, Vol. 19, Nos. (2/3), pp. 139–150, 2003.
31. L. Ros and F. Thomas, “Analysing spatial realizability of line drawings through edge-concurrency tests,” in *IEEE Int. Conf. on Robotics and Automation*, Vol. IV, Leuven, Belgium, pp. 3559–3566, 1998.
32. L. Ros and F. Thomas, “Overcoming superstrictness in line drawing interpretation,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 24, No. 4, pp. 456–466, 2002.
33. I. Shimshoni and J. Ponce, “Recovering the shape of polyhedra using line-drawing analysis and complex reflectance models,” *Computer Vision and Image Understanding*, Vol. 65, No. 2, pp. 296–310, 1997.
34. J. A. Shufelt, “Performance evaluation and analysis of vanishing point detection techniques,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 21, No. 3, pp. 282–288, 1999.
35. K. Sugihara, “Mathematical structures of line drawings of polyhedrons. towards man-machine communication by means of line drawings,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 4, pp. 458–469, 1982.
36. K. Sugihara, “An algebraic and combinatorial approach to the analysis of line drawings of polyhedra,” *Discrete Applied Mathematics*, Vol. 9, pp. 77–104, 1984a.
37. K. Sugihara, “An algebraic approach to shape-from-image problems,” *Artificial Intelligence*, Vol. 23, pp. 59–95, 1984b.
38. K. Sugihara, “A necessary and sufficient condition for a picture to represent a polyhedral scene,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 6, pp. 578–586, 1984c.
39. K. Sugihara, *Machine Interpretation of Line Drawings*. The MIT Press, 1986.
40. K. Sugihara, “Three-dimensional realization of anomalous pictures—An application of picture interpretation theory to toy

- design," *Pattern Recognition*, Vol. 30, No. 7, pp. 1061–1067, 1997.
41. D. Trytten and M. Tuceryan, "The construction of labeled line-drawings from intensity images," *Pattern Recognition*, Vol. 28, No. 2, pp. 171–198, 1995.
 42. D. Waltz, "Understanding line drawings of scenes with shadows." In: *The Psychology of Computer Vision*, P.H. Winston (Ed.) Mc. Graw Hill, 1975, pp. 19–91.
 43. W. Whiteley, "Motions and stresses of projected polyhedra," *Structural Topology*, Vol. 7, pp. 13–38, 1982.
 44. W. Whiteley, "From a line drawing to a polyhedron," *Journal of Mathematical Psychology*, Vol. 31, pp. 441–448, 1987.
 45. W. Whiteley, "Weavings, sections and projections of spherical polyhedra," *Discrete Applied Mathematics*, Vol. 32, pp. 275–294, 1991.



Lluís Ros was born in Vilafranca del Penedès (Catalonia), on June 30, 1968. He is currently a research scientist at the Institut de Robòtica i Informàtica Industrial of the Spanish High Council for Scientific Research, where he holds a Ramón y Cajal position since February 2003. He received a M.Sc. degree in mechanical engineering in 1992, and a Ph.D. degree (with honors) in advanced robotics in 2000, both from the Polytechnic University of Catalonia. His research interests are in geometry, with applications to machine vision, robot kinematics and computer graphics.



Federico Thomas was born in Barcelona, on October 5, 1961. He received the B.S. degree in telecommunications engineering in 1984, and the Ph.D. degree (with honors) in computer science in 1988, both from the Polytechnic University of Catalonia. Since March 1990, he has been a research scientist in the Institut de Robòtica i Informàtica Industrial of the Spanish High Council for Scientific Research. His research interests are in geometry and kinematics, with applications to robotics, computer graphics and machine vision.