

Shape-from-Image via Cross-Sections

Lluís Ros and Federico Thomas
Institut de Robòtica i Informàtica Industrial (CSIC-UPC)
Gran Capità 2-4, 08034 Barcelona. Catalonia (Spain).
{llros,fthomas}@iri.upc.es

Abstract

Using structural geometry, 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. Our proof is constructive, showing how to derive all spatial interpretations, it relies on elementary synthetic geometric arguments, and, as a by-product, it yields a simpler and shorter proof of Whiteley’s result. Moreover, important properties of line drawings are visually derived as corollaries: realizability is independent of the adopted projection, it is an invariant projective property, and for trihedral drawings it can be checked with a pencil and an unmarked ruler alone.

1. Introduction

Emulating the human performance in interpreting drawings of polyhedra has been one of the goals of Computer Vision along the past three decades [2, 13, 12, 3, 1]. 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. See [12, 6] for a recopilation of results.

It is well known that the problem of deciding whether a line drawing is *realizable*—i.e. it actually represents the correct projection of a polyhedral scene—was solved by Sugihara in his series of papers [11, 10, 9], where he reduced it to an instance of linear programming. What is less known, however, is that for spherical polyhedra the correctness can be decided only checking the concurrence of groups of

three lines derived from the drawing itself. (By spherical we mean here that the polyhedron is homeomorphic to a sphere.) Although the authors independently proved this in [8], they later found that the result had already appeared in [14], due to Whiteley. Using elementary synthetic geometric arguments, we here see that the same test is valid for drawings of polyhedral surfaces homeomorphic to a disk. Our proof is constructive and it also yields a simpler and shorter proof of Whiteley’s result, where, contrary to [14], tools of structural geometry are not needed. Important properties visually emerge as corollaries. Namely, for a line drawing, (1) its spatial realizability is independent from the adopted projection; (2) if it is realizable, then a projective transformation applied to it yields another realizable drawing; and (3) if it is *trihedral* (i. e., all vertices have exactly three incident faces), its realizability can be checked with a pencil and an unmarked ruler alone, without resorting to Sugihara’s algebraic formulation.

2. Basic Background

A *line drawing* is a diagram with straight line segments, called *edges*, and points where two or more segments meet, called *vertices*. A *polyhedral disk*, or *polydisk* for short, is a piecewise linear and continuous surface homeomorphic to a disk, made up of a collection of planar polygonal faces, f_1, f_2, \dots, f_m , glued in pairs along the edges. For the sake of simplicity, we assume that the drawing, if correct, is produced by orthogonally projecting a single polydisk onto the XY plane, showing all its edges, even the hidden ones. Section 4 extends the results to other types of projection, and [6] extends them to drawings of more complicated scenes, without hidden edges, several objects and possible occlusions between them. We say that a drawing \mathcal{D} is *correct*, or *realizable*, if we can vertically lift the vertices of \mathcal{D} to construct a polydisk that projects onto \mathcal{D} , with distinct planes for adjacent faces. Such a polydisk is called an *interpretation* or a *lifting* of the drawing.

We assume that a drawing is given along with its *incidence structure* [12], that tells the sets of vertices of the drawing that form a planar face in any spatial in-

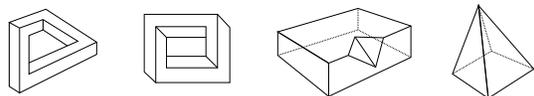


Figure 1. Line drawings.

PSfrag replacements

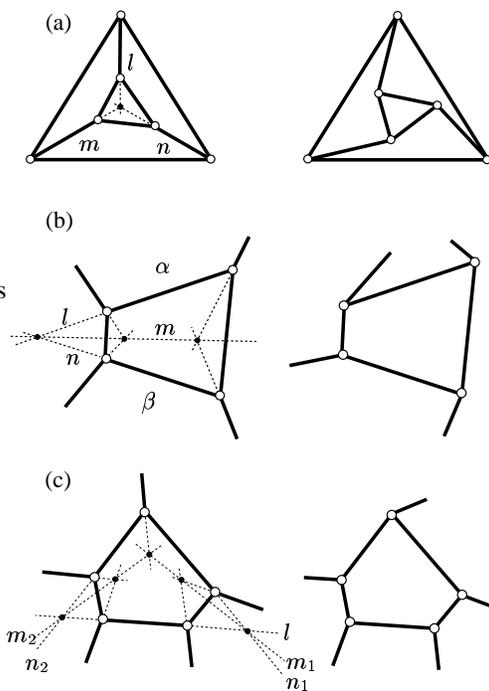


Figure 2. Correct (left) and incorrect drawings (right).

terpretation. The incidence structure can be computed by applying the method in [12, pag. 45], after a labelling of its edges has been obtained using techniques described in [2, 13, 5].

For a polydisk \mathcal{P} in 3-space, we say that a face f of \mathcal{P} is *sequentially adjacent* to faces $f_{i_1}, f_{i_2}, \dots, f_{i_m}$ of \mathcal{P} if the edges between f and these faces are *sequentially linked*, meaning that if (f, f_{i_k}) denotes the edge between faces f and f_{i_k} , then in the sequence $s = \{(f, f_{i_1}), (f, f_{i_2}), \dots, (f, f_{i_m})\}$ every edge shares a vertex with the previous and the next one, except for the first and last edges in s , which only share a vertex with the next and the previous edge, respectively.

3. The cross-section test

Some examples point out that a drawing's correctness is checkable solely using concurrence conditions. The truncated tetrahedron in fig. 2a is only correct when its three edges l, m, n meet at a common point. The 4-calotte in fig. 2b, 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 at the line m of intersection of the planes α and β . Likewise, the 5-calotte in fig. 2c 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. The challenge was to characterize a set of also sufficient concurrence conditions for a drawing to be correct. We next present one.

Consider a spherical polyhedron in 3-space such as, for example, the tetrahedron in fig. 3, top. Now, ob-

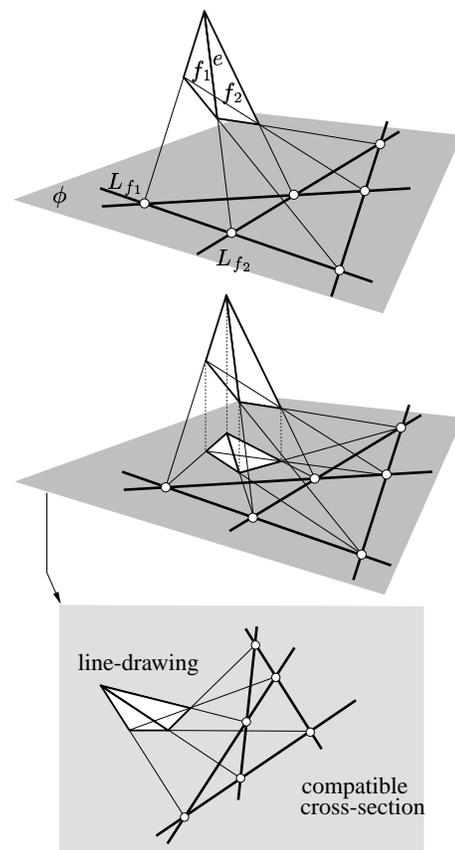


Figure 3. The cross-section test.

tain the intersections of the planes of its faces with an external plane ϕ in general position. The resulting arrangement of lines is called a *cross-section* of the polyhedron. It is clear that the edge line e between two faces, say f_1 and f_2 , must be concurrent to the point of intersection of the lines L_{f_1} and L_{f_2} of intersection of the planes of f_1 and f_2 with ϕ . These trivial concurrence conditions 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. 3, 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. 3, bottom). Whiteley's theorem states that the converse is also true. We will also prove the following.

Theorem 1. *A line drawing of a polydisk is realizable if and only if it has a compatible cross-section, with non-coincident cross-section lines.*

By "compatible" we mean here that all non-boundary edges of the polydisk are concurrent to their corresponding point in the cross-section, leaving the boundary ones unconstrained.

Proof. The "only if" part is proven by the above arguments. We will prove the "if" part using the drawing and its compatible cross-section to explicitly construct

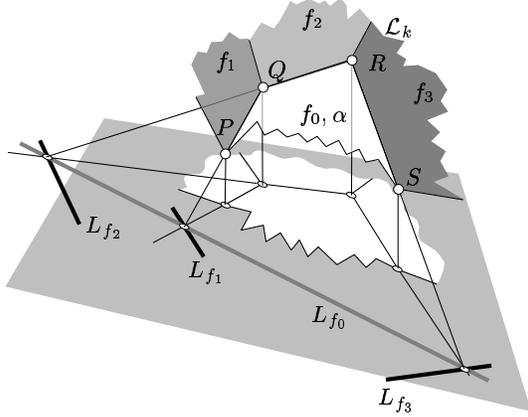


Figure 4. Lifting an intermediate face.

a polydisk, one whose face-planes generate the given cross-section when intersected with the plane of the drawing.

The proof proceeds using induction. First we show how to lift a first face of the drawing keeping all its vertices coplanar. Then we assume that a polydisk \mathcal{L}_k with k faces has already been lifted correctly and prove that any other face that is sequentially adjacent to some faces in \mathcal{L}_k can be properly lifted too, to form a lifted polydisk \mathcal{L}_{k+1} with $k + 1$ faces.

To start with, take any of the faces of the drawing 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 the face, and all edges and vertices of the face are then lifted vertically to lie on this plane.

As induction hypothesis, assume that we have a polydisk \mathcal{L}_k already lifted correctly, where for every edge between two faces, its line meets the point where the cross-section lines of these two faces intersect. This applies to all edges of \mathcal{L}_k , even those at its boundary. Note that this hypothesis is true if \mathcal{L}_k only contains the first lifted face.

Now, we can lift an additional face f_0 , not in \mathcal{L}_k , that is sequentially adjacent to m faces f_1, \dots, f_m of \mathcal{L}_k (and to no other face of \mathcal{L}_k) through m edges PQ, QR, RS, \dots (see fig. 4, where we represent the case $m = 3$). We will prove that these edges and the cross-section line L_{f_0} of f_0 are all coplanar, and define a plane α that is different from all planes assigned to the faces f_1, \dots, f_m .

To see this, note that, as the edge line PQ is incident with L_{f_0} (by induction hypothesis), PQ and L_{f_0} are coplanar. Let us call α the plane they define. Moreover, the edge line QR is also coplanar with α , as it contains two points of this plane: the point where it intersects with L_{f_0} and the point Q of line PQ . The same applies to line RS as it is incident with R and L_{f_0} . Clearly, the argument can be iterated to prove that all other edges between f_0 and \mathcal{L}_k are coplanar with α .

With the plane for f_0 already fixed, all other edges

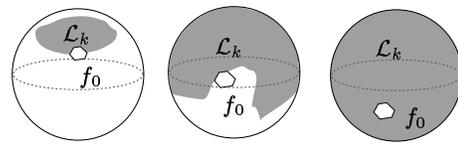


Figure 5. Evolution of \mathcal{L}_k .

between f_0 and faces not in \mathcal{L}_k can be fixed too by lifting them vertically to lie in this plane. For every such edge, say between face f_0 and face 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_0} meet. Clearly, this line 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_0} , 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_1, \dots, f_m . But this is trivially true, as the only way for α to coincide with one of such planes would be that the corresponding cross-section lines are identical, which is not the case, by the premises of the theorem. \square

4. Derived results

Corollary 1 (Whiteley's theorem). *A line drawing of a spherical polyhedron is realizable if and only if it has a compatible cross-section, with non-coincident cross-section lines.*

Proof. The spherical polyhedron can be constructed by generating a sequence of lifted polydisks $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3, \dots$, adding one face at a time. At the end, a last face will close the polyhedron, but all its edges will be coplanar, as the face is sequentially adjacent to the previous polydisk (fig. 5). \square

However, the test cannot be extended to deal with projections of polyhedral objects with holes, as they cannot be lifted by subsequently adding faces that are sequentially adjacent to a previous polydisk. See this, e. g., for a topologic disk with a hole in fig. 6a. As a counterexample, consider three pairwise adjacent faces with a hole (fig. 6b). They always have a compatible cross-section, but the drawing is not realizable unless the three non-boundary edges, l, m and n , are concurrent.

Corollary 2. *Realizability does not depend on the type of projection assumed.*

Proof. If we observe fig. 4 we see that the same proof works for central projection, the only difference being that the vertices P, Q, R, \dots , are lifted along lines that meet at the center of projection. Similarly, the proof is valid for oblique parallel projection. \square

Corollary 3. *Realizability is a projectively invariant property.*

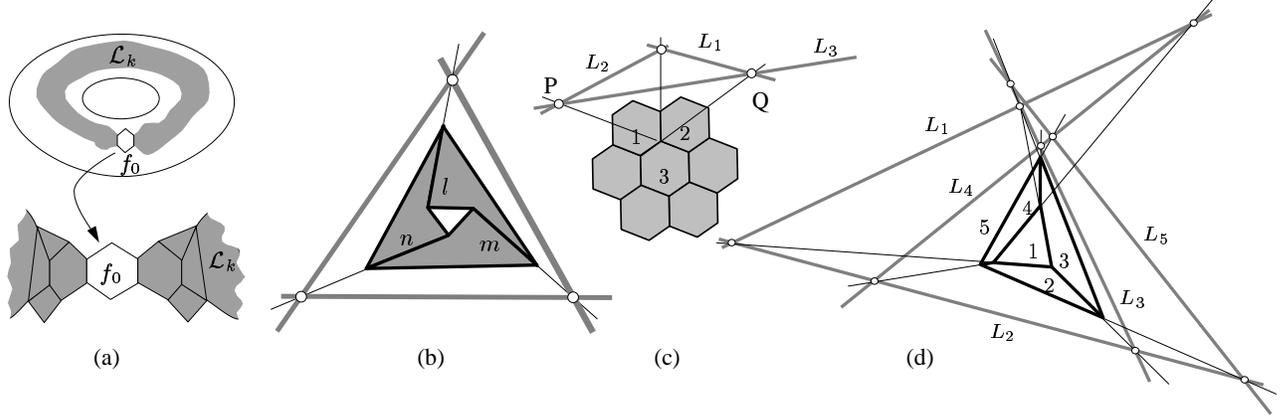


Figure 6. (a) and (b): polydisks with holes. (c) and (d): test of a trihedral drawing.

Proof. 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 [4]. Hence, if a drawing has a compatible cross-section, the transformed drawing will also have one. \square

Corollary 4. *Realizability of trihedral drawings can be checked with a pencil and an unmarked ruler alone.*

Proof. Note that when lifting the vertices of a correct drawing, one can always choose the heights of four vertices independently, namely, those of the two vertices of an edge, and the height of one vertex in each of the faces of this edge. So, these two faces can receive arbitrary planes. 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 concurrent to a same point of their common edge-line. For trihedral drawings these two initial lines completely determine the rest of the cross-section. We see this in fig. 6c, where, after fixing lines L_1 and L_2 for faces 1 and 2, L_3 can be automatically deduced as it must contain P and Q , the points where the edge lines between faces 3 and 1, and 3 and 2 meet with the cross-section lines of faces 1 and 2, respectively. We can obtain lines for the other faces iterating this process, and construct the cross-section. If at some point a concurrence condition does not hold, we conclude that the drawing is incorrect. If we complete the construction, by corollary 1 it is correct. Fig. 6d shows the process on a correct truncated tetrahedron. \square

5. Conclusions

Whiteley’s proof in [14] is very valuable from the standpoint of structural geometry, as it clearly shows the connections with self-stressed grillages. Although ours loses this connection, it is highly visual, intuitive and simple.

Unlike trihedral drawings, general ones do not have determined cross-sections. Hence to make them prac-

tical for drawing interpretation, one needs a way to automatically find all compatible cross-sections or show that none exists. In [6] we provide one, together with a tool for correcting incorrect drawings (also in [7]).

References

- [1] H. Barrow and J. Tenenbaum. Retrospective on “interpreting line drawings as three-dimensional surfaces”. *Artificial Intelligence*, 59:71–80, 1993.
- [2] D. Huffman. Impossible objects as nonsense sentences. *Machine Intelligence*, 6:295–323, 1971.
- [3] T. Kanade. Recovery of the three-dimensional shape of an object from a single view. *Artificial Intelligence*, 17:409–460, 1981.
- [4] F. Klein. *Elementary Mathematics from An Advanced Standpoint*. Dover Publications, 1939.
- [5] P. Parodi and V. Torre. On the complexity of labeling perspective projections of polyhedral scenes. *Artificial Intelligence*, 70(1-2):239–276, October 1994.
- [6] L. Ros. *A Structural Geometric Approach to Spatial Interpretation of Line Drawings*. PhD thesis, Polytechnic University of Catalonia, To appear.
- [7] L. Ros and F. Thomas. Overcoming superstrictness in line drawing interpretation. Submitted for publication.
- [8] L. Ros and F. Thomas. Analysing spatial realizability of line drawings through edge-concurrence tests. In *IEEE Int. Conf. on Robotics and Automation*, May 1998.
- [9] K. Sugihara. An algebraic and combinatorial approach to the analysis of line drawings of polyhedra. *Discrete Applied Mathematics*, 9:77–104, 1984.
- [10] K. Sugihara. An algebraic approach to shape-from-image problems. *Artificial Intelligence*, 23:59–95, 1984.
- [11] K. Sugihara. A necessary and sufficient condition for a picture to represent a polyhedral scene. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-6:578–586, September 1984.
- [12] K. Sugihara. *Machine Interpretation of Line Drawings*. The MIT Press, 1986.
- [13] D. Waltz. Understanding line drawings of scenes with shadows. In P. H. Winston, editor, *The Psychology of Computer Vision*, pages 19–91. Mc. Graw Hill, 1975.
- [14] W. Whiteley. Weavings, sections and projections of spherical polyhedra. *Discrete Applied Mathematics*, 32, pages 275–294, 1991.