International Journal of Computer Vision 43(2), 75–97, 2001

The Interpretation of Line Drawings with Contrast Failure and Shadows

MARTIN C. COOPER
IRIT, University of Toulouse III, 118, route de Narbonne, 31062 Toulouse, France
cooper@irit.fr

Abstract. In line drawings derived from real images, lines may be missing due to contrast failure and objects with curved surfaces may cast shadows from multiple light sources.

This paper shows that it is the presence of shadows, rather than contrast failure, that renders the line drawing labelling problem NP-complete. However, shadows are a valuable visual cue, since their presence is formally shown to reduce the average ambiguity of drawings. This is especially true when constraints concerning shadow formation are employed to differentiate shadow and non-shadow lines.

The extended junction constraint, concerning straight lines colinear with junctions, compensates the loss of information caused by contrast failure. In fact, we observe the contrast failure paradox: a drawing is sometimes less ambiguous when lines are partly missing due to contrast failure.

It is known that the coplanarity of sets of object vertices can be deduced from the presence of straight lines in the drawing. This paper shows that these coplanarity constraints are robust to the presence of contrast failure.

Keywords: line drawing labelling, contrast failure, shadows, coplanarity constraints, extended junction constraints, constraint satisfaction problem

1. Introduction

The interpretation of line drawings is a classic problem in artificial intelligence. The pioneering work of Huffman (1971), Clowes (1971) and Waltz (1975), in identifying constraints on the semantic labelling of junctions and the propagation of these constraints, made restrictive assumptions that limited the direct application of their work in computer vision systems. The seminal work of Sugihara (1982, 1984, 1986) in which necessary and sufficient conditions for the physical realisability of a labelled drawing were coded as a standard linear programming problem, retained the very restrictive assumptions that the drawing is a perfect projection (in the sense that there are no missing lines) of a polyhedral scene, thus limiting its application to the interpretation of human-entered drawings of polyhedral scenes. Shimshoni and Ponce's (1997) technique for dealing with inaccurate line drawings is also restricted to images of polyhedral scenes without contrast failure.

The constraints on the semantic labelling of junctions were generalised to curved objects with C^3 surfaces (Cooper, 1993, 1997a, 1999; Malik, 1987). However, the possible presence of phantom junctions (undetectable label transitions) on curved lines reduces the predictive power of a catalogue of semantic junction labellings, compared with the catalogue for polyhedral objects, since the two ends of a curved line do not necessarily have the same semantic label. Semantic junction labellings provide even less information when the drawing is an *imperfect* projection of the surface-normal discontinuity edges of a 3D scene.

However, we will show that the ambiguity caused by contrast failure can be greatly reduced by the introduction of a new constraint concerning lines colinear with junctions. Valuable information can also be obtained from constraints concerning shadows, in terms of the characteristic properties of shadow edges in the intensity image, in terms of the types of junctions that can be joined by shadow lines and in terms of the types of vertices which can cast distant shadows.

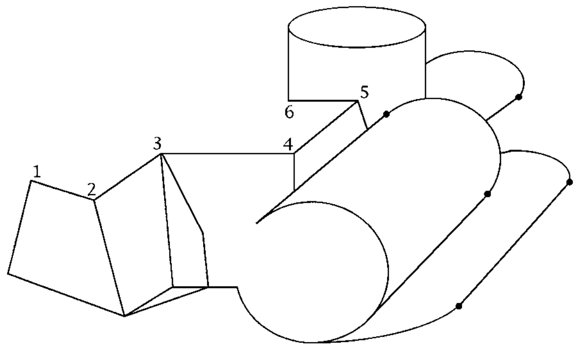

Figure 1. A drawing with shadow lines and contrast failure.

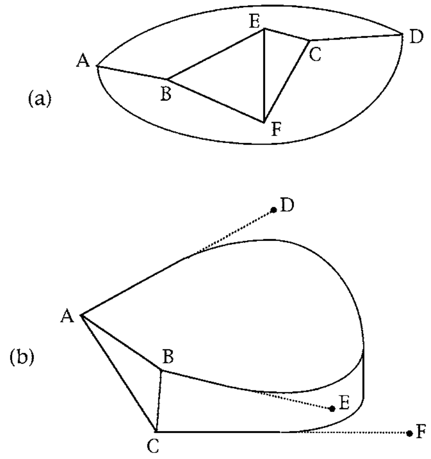

Figure 2. Drawings which are impossible due to coplanarity constraints.

An important source of information when interpreting a drawing of man-made objects is the presence of linear features such as straight lines, parallel lines, colinear points and coplanar points. Such linear features in the drawing can be translated into a standard linear programming problem involving a system of linear equations and linear inequalities on variables representing the depths of object vertices (Cooper, 2000). Computer vision systems must interpret line drawings derived from real images in which contrast failure means that certain lines will be missing and lighting effects mean that spurious lines will occur due to shadows. An example is given in Fig. 1. This paper shows that the constraints derived from linear features (such as coplanarity) are almost all robust to contrast failure and the presence of shadow lines. They allow us, among other things, to deduce that the points 1, 2, 3, 4, 5 and 6 in Fig. 1 are coplanar.

The strength of these coplanarity constraints is illustrated by the fact that they are able to identify the impossible drawings illustrated in Fig. 2. These drawings are physically unrealisable as 3D scenes under the assumption that straight lines are projections of straight edges formed by the intersection of surfaces which are planar in the vicinity of the edge. Two applications of the constraint to the drawing of Fig. 2(a) are sufficient to deduce that the points A, B, E, C, D are coplanar. Two further applications indicate that A, B, F, C, D are also coplanar. A contradiction arises from the fact that the points A, B, C, D should be colinear since they all lie on the intersection of these two planes.

The drawing in Fig. 2(b) shows that the coplanarity constraint can still be applied to straight edges in curved objects and even to hidden surfaces under a further assumption that all vertices are trihedral. The tangents AD and BE must be coplanar with the straight edge AB. Similarly the tangents BE and CF must be coplanar with the straight edge BC. Furthermore, AD and CF are both tangential to the same hidden surface which is planar in the vicinity of the straight edge AC. The extensions of lines AD, BE and CF should intersect in 3D space (at the point of intersection of the three planes ADBE, BECF and ADCF) which is clearly impossible since their projections in the drawing do not intersect. The drawing in Fig. 2(b) is therefore also unrealisable.

2. Problem Statement

The problem studied in this paper is the recovery of depth information from a line drawing of a 3D scene containing objects composed of smooth C^3 curved surfaces. Any curved line in the drawing could be the projection of an infinite family of 3D curves. Thus, given this overwhelming ambiguity in the shape of curved edges, we restrict ourselves to the determination of colinearity and coplanarity relationships between the visible object vertices. Semantic labellings (involving labels such as occluding, convex, concave or shadow for each line) will also be determined for each line junction in the drawing.

In order to obtain constraints on the semantic labellings of junctions and the depths of vertices, it is necessary to make assumptions concerning object shape and/or image formation. The following list includes all the assumptions that we will be led to make in order to obtain the constraints described in this paper. Note that some constraints are based on fewer assumptions than others.

General viewpoint assumption: A small perturbation in the position of the viewpoint does not change the configuration of the drawing (intersection of lines, presence of straight lines, colinear points, etc.).

General light source position assumption: A small change in the position of any light source does not change the configuration of the drawing.

General relative position assumption: A small change in the relative positions of objects does not change the configuration of the drawing.

Polyhedral vertices assumption: Every object vertex behaves locally as a polyhedral vertex: the surfaces $S_0, S_1, \ldots, S_{n-1}$ which intersect to form the vertex V each have unique tangent-planes $T_0, T_1, \ldots, T_{n-1}$ at V; the intersection of T_i and $T_{i+1(\mathrm{mod}\ n)}$ is tangential at V to the edge E_i which is the intersection of surfaces S_i and $S_{i+1(\mathrm{mod}\ n)}$.

Trihedral vertices assumption: Every object vertex behaves locally as a polyhedral vertex formed by the intersection of exactly three surfaces.

Non-tangential edges and surfaces assumption: None of the surfaces S_i and edges E_i which intersect to form a vertex V meet tangentially at V. Every object edge is formed by the non-tangential intersection of two surfaces.

Straight edge formation assumption: Every straight edge in 3D is formed by the intersection of locally planar surfaces.

Extended trihedral assumption: For a given object A, let P_A be the set of planes which are either tangential to a surface of A at a vertex or tangential to a surface of A along a straight edge. All objects A satisfy not only the trihedral vertices assumption but also the condition that no four planes in P_A meet at a point.

Note that the straight edge formation assumption is particularly powerful. It has been used to establish coplanarity constraints in the analysis of perfect projections of curved objects containing some straight edges (Cooper, 2000), but it can also be applied to deduce, for example, that a straight shadow line is the shadow cast by a straight edge onto a planar surface.

Under the assumptions of trihedral vertices where edges and surfaces meet non-tangentially, there are only five viewpoint-independent object vertex types which can project into junctions of degree less than four. These are illustrated in Fig. 3 and numbered 1 to 5. Note that a reflected version of vertex type 5 can also occur. Four more vertex types occur at points where the line of sight is tangential to an object surface (and are thus viewpoint-dependent vertices). These are the vertices numbered 6 to 9 in Fig. 3. In multi-object scenes, the multi-object vertices numbered 10 to 15 can also occur. To limit the number of multi-object vertices, we only incude vertices which are stable to small changes in the relative positions of objects. This implies, for example, that no three objects meet at a vertex, that no two object edges coincide and that no object vertex lies on the edge of another object.

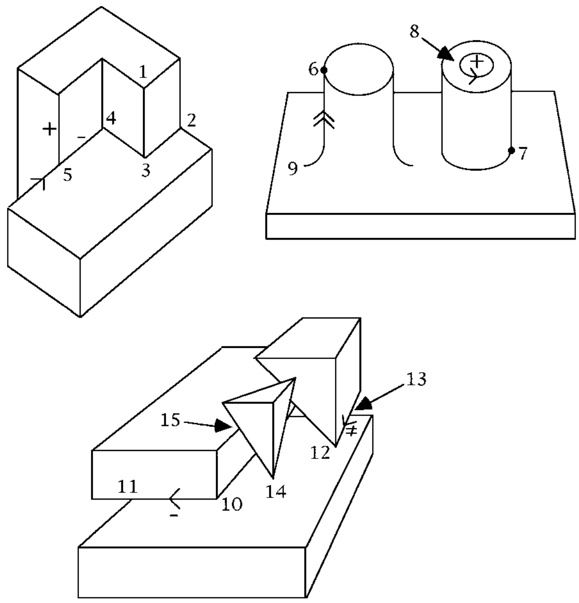

Figure 3. The basic set of 15 different vertices derived from the assumptions of general relative position, trihedral vertices and non-tangential edges and surfaces.

Edges can be classified as follows. *Occluding* edges are formed by the intersection of two surfaces, only one of which is visible, and are denoted by an arrow (with the occluding surface on the right as we follow the direction of the arrow). *Extremal* edges are formed by a single curved surface being tangential to the line of sight, and are denoted by a double-headed arrow. For example, the contour of a cube is made up of occluding edges, whereas the contour of a sphere is an extremal edge. A *convex* edge is formed by the intersection of two surfaces (both of which are visible) at an exterior angle of greater than π, and is denoted by "+". If the exterior angle is less than π, then the edge is *concave* and is denoted by "−". However, in multi-object scenes, concave edges can also be caused by two objects touching along a common surface (for example, the edge 10–11 in Fig. 3) or by two objects touching along the edge of one of the objects (for example, the edge 12–13 in Fig. 3). These two edges are denoted by the labels '$\leftarrow -$' and '$\leftarrow\neq$', respectively. Examples of all line labels can be found in Fig. 3. For the set of vertices studied in this paper (numbered 1 to 15 in Fig. 3), it is unnecessary to make the distinction between '$\leftarrow -$' '$\rightarrow -$' and '$-$' edges, since a pair of objects which touch along a common face can always be interpreted as a single object. In the terminology of the constraint satisfaction problem, the label '$-$' is neighbourhood substitutable (Cooper, 1997b, Freuder, 1991) for the labels '$\leftarrow -$' and '$\rightarrow -$' at every line-end.

Under the assumption that straight lines are projections of straight edges formed by the intersection of locally planar surfaces, label transitions on straight lines in a drawing which is not subject to contrast failure are illegal. Nevertheless, when the surfaces which meet to form an edge may be curved, undetectable transitions from convex to occluding labels are possible and are known as C-junctions or phantom junctions. An example is the projection of vertex 8 in Fig. 3. Phantom junctions can also occur on straight lines, due to contrast failure at projections of vertices of type 5, 11 or 13 of Fig. 3; a T-junction is undetectable when the stem of the T-junction is missing due to contrast failure.

3. Constraints on the Form of Edges in the Intensity Image

If the drawing is derived by detecting edges in an intensity image, then this has two consequences:

- the existence of spurious or missing lines in the drawing;
- further information is available in the form of intensity values in the vicinity of edges which could be used to constrain their labelling as concave, convex, occluding, extremal, shadow, etc.

Edges can be classified as either step or ramp, depending on the form of their cross-section in the intensity image. Figure 4(a) shows a typical scene containing three objects and one light source, viewed from above. Figure 4(b) shows the cross-section of the intensity image of the larger of the three objects. Assuming Lambertian surfaces, step edges can occur due to discontinuities in surface albedo, surface orientation or illumination. For example, the step edge D in Fig. 4 is a shadow edge, caused by a discontinuity in illumination.

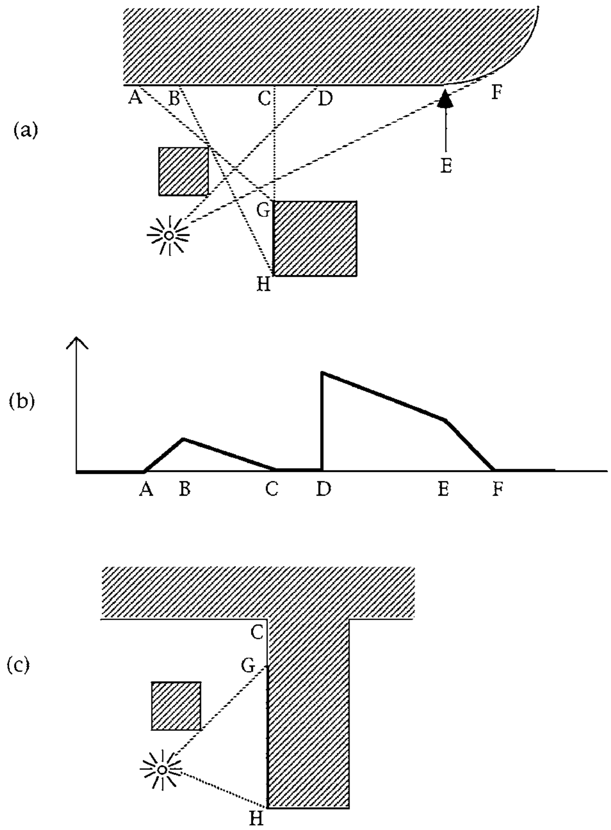

Figure 4. Example of the formation of different types of edge.

Objects do not always cast sharp shadows. Non-point light sources, such as windows, give rise to soft shadows consisting of an umbra and a penumbra. Even when no non-point light sources are apparently present, large reflective object surfaces provide secondary non-point light sources. In Fig. 4, the surface GH acts as a secondary light source, illuminating the viewed object between A and C, producing a penumbra on section AB. The penumbra is bounded by two ramp edges. The width of the penumbra depends on the width of the non-point light source. A shadow step edge such as D can be considered as a penumbra of zero width.

It can easily be shown that, when the secondary light source is planar, another ramp edge occurs at the intersection of this plane with the viewed object. Since GH is planar, a ramp edge thus occurs at C. More importantly, this is true even in the case when the planar surface GH actually intersects the viewed object surface at C, as shown in Fig. 4(c). In other words, a concave edge may be visible in the intensity image as a ramp edge, rather than a step edge, if it is not directly illuminated. Note that, on the other hand, for a directly-illuminated concave edge E, inter-reflection models (Shimshoni and Ponce, 1997) which take into account the fact that the surfaces meeting at E are secondary light sources, predict an increase in brightness as we approach E.

Ramp edges have two other possible causes, as illustrated by the points E and F in Fig. 4(a). The point E represents a discontinuity of surface curvature of the viewed object. In general, discontinuities of the nth derivative in the object surface are visible as discontinuities of the $(n - 1)$th derivative in the intensity image. Note that discontinuities of surface curvature are, in fact, disallowed by the assumption that object surfaces meet non-tangentially, but have been studied in previous papers (Cooper, 1993, 1997a). When rays emanating from a point light source strike an object surface tangentially, this gives rise to another ramp edge in the intensity image, as occurs at F in Fig. 4(a).

Ramp edges can be classified as valleys if the first derivative of intensity increases as we cross the edge, or as ridges if the first derivative decreases. It can be shown that edges A, C and F must always be valleys and that edge B must always be a ridge. However, the edge E would change from a valley to a ridge if it was illuminated from the right rather than from the left. When the light source is not a point light source, a characteristic feature of a shadow edge leaving a junction is that it starts as a step edge and gradually separates into two ramp edges, one ridge and one valley.

Shadow step edges cast by point light sources, such as edge D in Fig. 4(a), also have a characteristic property (Rubin and Richards, 1982) which we state as a constraint.

Shadow Edge Constraint

On the illuminated side of a shadow edge the viewed surface emits more light than on the shaded side, *and this is true at all points along the edge and for all wavelengths*.

For example, an edge separating a dark green region from a bright red region cannot be a shadow edge. Under the stronger assumption that all illumination falling on the viewed surface is from identical white light sources, the observed hue should, in fact, be identical on both sides of a shadow edge. However, colour saturation (proportion of white light) may be greater on the brighter side of the edge due to specular reflection.

We have seen that concave edges in 3D, without any direct illumination, can be visible as valley ramp edges in the intensity image. Under what circumstances can depth-discontinuity, concave or convex edges be invisible in the intensity image? This phenomenon is known as contrast failure. An occluding edge is not guaranteed to produce any kind of visible edge. A depth discontinuity between two parallel surfaces of identical surface characteristics and subject to the same illumination will necessarily be undetectable. Since, in man-made objects, pairs of parallel surfaces are common, a minimum requirement when analysing line drawings derived from intensity images is that certain occluding lines may be missing. It should also be noted that it is possible to construct depth discontinuities which are visible as ramp instead of step edges, if the two parallel surfaces have distinct curvatures.

Contrast failure can occur between any two surfaces of identical albedo if they receive identical illumination (for example, if the only light source happens to lie in the plane bisecting the two surfaces). Although this kind of contrast failure is viewpoint independent it is not independent of the light source positions. To obtain most of the constraints stated in this paper, we will assume that this kind of contrast failure does not occur, by making the general light source position assumption.

Two other causes of contrast failure must be mentioned, although they are assumed not to occur in this paper:

- the surfaces on the two sides of the edge emit no light, either due to zero albedo (i.e. the surface is completely black) or due to zero illumination (i.e. the surface is in complete darkness);
- the two surfaces are illuminated by a totally diffuse light source (i.e. equal illumination from all directions).

The first of these two cases can only occur in image regions which are particularly dark. It is possible to establish a catalogue of junction labellings which allows for this kind of contrast failure, and apply it solely at junctions including at least one sufficiently dark region.

It can be noted that in the case of totally diffuse lighting, not all edges are necessarily invisible. For example, when a cube is placed on a dark table and illuminated by a totally diffuse light source, its top surface receives approximately twice as much illumination as its sides, meaning that the boundary of the top surface is detectable in the intensity image. The upright edges of the cube would, however, be invisible. Although objects do not cast shadows under totally diffuse lighting, ramp "shadow" edges can still occur, for example at points such as C in Fig. 4(a), due to discontinuities in illumination gradient.

The following two assumptions, together with the general light source position assumption, limit contrast failure to occluding edges between two locally parallel surfaces.

Illumination assumption: every point of every visible object surface is illuminated by a light source which is not totally diffuse.

Albedo assumption: every point of every visible object surface has non-zero albedo.

The constraints presented in this paper only hold for idealised drawings in which contrast failure only occurs when surfaces are parallel. In practice, due the finite precision of imaging devices, contrast failure will occur for other reasons, such as a light source lying near to the bisecting plane of a 3D edge. It is important to understand the idealised case before progressing to more realistic constraints. These more realistic constraints can be expressed in a natural way as valued constraints (Cooper, not yet published). In the framework of the valued constraint satisfaction problem, all junction labellings that may occur due to unlikely, but not entirely impossible, scene configurations must be determined and assigned some low but non-zero value.

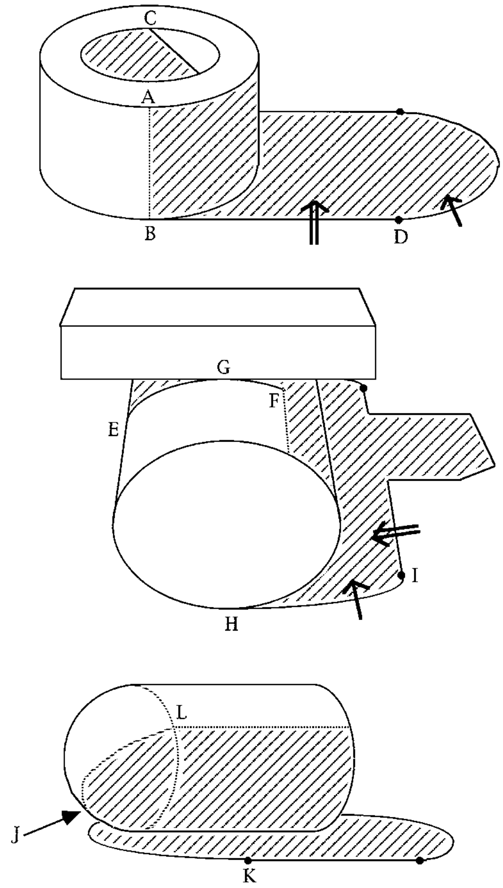

Figure 5. Junctions caused by shadows and curved surfaces.

4. Shadows and Curved Surfaces

Waltz (1975) studied in detail the ways in which shadows can be cast at polyhedral vertices. When object surfaces may be curved, new junctions can occur in the drawing or become visible due to the presence of shadows.

When rays emanating from a point light source touch an object surface tangentially they create a ramp shadow edge, such as AB in Fig. 5. We call such edges extremal illumination edges. They are always valley ramp edges in the intensity image. Extremal illumination edges terminate when they meet a surface-normal discontinuity edge E (such as at A and B in Fig. 5), and they intersect their cast shadow at this point (as at B) if and only if the edge E is concave. An L-junction occurs in an extremal illumination edge when it crosses a discontinuity in surface curvature of the underlying surface, as illustrated at junction L in Fig. 5 (Nalwa, 1988). Discontinuities of surface curvature, which also give rise to ramp edges in the intensity image, are discussed in detail in (Cooper, 1993), but are, in fact, disallowed in this paper by the assumption that surfaces meet non-tangentially. When a shadow edge S, representing a step edge in the intensity image, meets an extremal illumination edge produced by the same light source, S fades out as it approaches the junction. This is illustrated by the junction F in Fig. 5, where GF must fade out.

The shadow of point A in Fig. 5 is the point D. The black dot marks a discontinuity of curvature of the line in the drawing. Such junctions are known as curvature-L junctions. The shadow of the line BA arrives at D from the left. We introduce the label $\Uparrow$ for shadow edges cast by extremal illumination edges. The shadow of the outer rim of the cylinder arrives at D from the right. We use Waltz's (1975) label $\uparrow$ for shadows cast by surface-normal discontinuity edges. Both types of arrow point towards the shaded region. The line arriving from the right, when extended to the other side of D, must lie within the shadow. Imagine that ramp lines such as AB cannot be detected. Then a junction of type B would be indistinguishable from a junction of type H. Nevertheless, the first curvature-L junction encountered along the shadow line allows us to distinguish between these two types of junction. At I, it is the shadow line arriving from the left which, when extended, must lie within the shadow, meaning that HI must be labelled ↑ as shown.

If we were to allow object surfaces to meet tangentially, then there would be two other causes of curvature-L junctions on a cast shadow line S: (1) a discontinuity of curvature in the edge casting the shadow S, or (2) a discontinuity of curvature in the surface casting the shadow, when S is the shadow cast by an extremal illumination edge. An example of case (2) is the junction K in Fig. 5, which is the shadow of junction L.

The junction C in Fig. 5 is a T junction. The bar of the T is a convex edge E and the stem of the T is a shadow edge S. In fact S is the shadow cast by E on one of the surfaces which intersect to form E. It can be shown that S actually fades out in the intensity image as it approaches the T-junction. However, since S meets the bar of the T-junction non-tangentially, this kind of junction provides no essentially new labelling since such a junction (without S fading out) is well known to be possible when a shadow falls on a straight convex edge (Waltz, 1975).

When a shadow line (representing either a step or ramp edge in the intensity image) meets an extremal line, such as the side of a cylinder or the contour of a sphere, it does so tangentially. Junctions E and J in Fig. 5 are examples. Junctions such as E are called 3-tangent junctions. Two of the lines meeting at the junction have continuous curvature.

The junctions G and H in Fig. 5 are essentially the same type of junction: the convex edge E of one object touches the surface of another object. Without shadows, this junction is invisible and has no effect on the labelling of the projection of E in the drawing. However, this junction may become visible due to the presence of shadows. Shadow lines may be visible on one side (e.g. junction H) or on both sides (e.g. junction G), and n light sources may produce up to n shadow lines on each side of the junction. Corresponding shadow lines (such as EG and GF) must have continuous curvature at the junction. Note that, although junctions such as G and H contradict the rule that an object edge should not meet an object surface tangentially, they are clearly too common to be ignored.

Another new junction is the intersection of two shadow lines (step or ramp), produced by two distinct light sources, in the form of an X. In fact, an X-junction is a valuable clue to line labelling, since the two lines crossing at the X-junction must be illumination or albedo discontinuity lines (such as shadows, reflections or surface markings).

5. A Constraint on Shadow Lines

We distinguish two types of junctions in a drawing: light-source dependent and light-source independent. When all light sources are subject to a small perturbation, the scene point projecting into a light-source-dependent junction changes, whereas light-source independent junctions are always projections of the same point in 3D. In Fig. 5, junctions A, B, C, D, E, F, I, J, K, L are all light-source dependent, whereas junctions G and H are light-source independent. Junctions which are projections of 3D vertices, such as vertices 1 to 15 in Fig. 3, are clearly light-source independent. In fact, junctions of type W, Y, K, Peak, Multi, 4-tangent and X_0 are always light-source independent. Definitions of the different possible junction types can be found in Appendix C (Fig. 24). A Ψ junction is also light-source independent if all but one of the lines meeting at the junction are known to be projections of surface-normal discontinuity edges. The following constraint concerns shadow lines colinear with or joining two light-source independent junctions. It is particularly useful for eliminating 'shadow' as a possible label for lines joining two light-source independent junctions in the drawing.

Straight Shadow Line Constraint

Under the assumptions of general viewpoint and general light source positions, a straight line colinear with two light-source independent junctions cannot be a shadow line.

This follows from the fact that the equation of a shadow line changes with a change in the position of the corresponding light source. On the other hand, since light-source independent junctions are projections of fixed 3D points, the line which joins them is invariant to changes in light source positions. Note that this constraint could be used to distinguish between ramp lines caused by illumination effects and ramp lines which are projections of surface-normal discontinuity edges.

The shadow line constraint described in this section is a powerful constraint which allows us to eliminate the

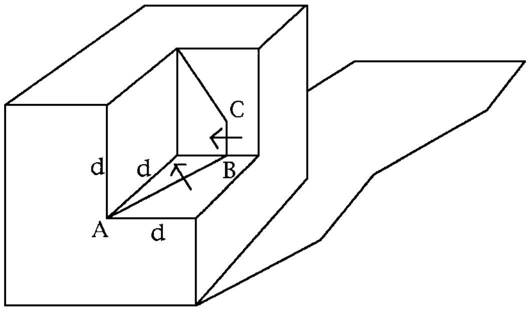

Figure 6. The straight shadow line constraint identifies the three lines labelled d as non-shadow lines, but also the line BC as a shadow line.

label 'shadow' for many lines. For example, the lines labelled d (for surface-normal discontinuity) in Fig. 6 cannot be shadow lines, since they are straight lines joining two light-source independent junctions. This, in turn, implies that line AB is a shadow line (under the assumption of trihedral vertices). Applying the straight shadow line constraint to line AB, which is now known to be a shadow, means that B cannot be a light-source independent junction. Thus the line BC must also be a shadow. This example shows that the shadow line constraint can not only be used to identify non-shadow lines, but can, in some cases, be used to identify shadow lines.

6. Extended Junction Constraints

Consider the junction A in Fig. 7(a). The fact that it is colinear with the straight line BC is a valuable clue which can be used to drastically reduce the number of possible legal labellings for A (from 43 to just 11). We use the term *extended junction* to denote a junction

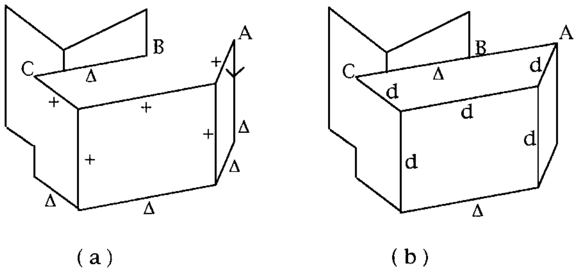

Figure 7. The contrast failure paradox: the drawing (a) is less ambiguous than the drawing (b) due to the fact that contrast failure has caused the line AB to be missing.

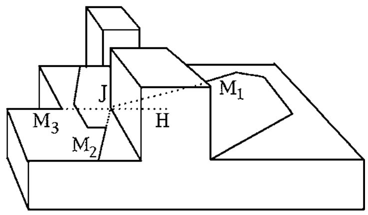

Figure 8. Different types of extended junctions J-M_1, J-M_2 and J-M_3, not caused by contrast failure between parallel surfaces.

together with a straight line which is colinear with it (such as A-BC in Fig. 7).

Consider an extended junction in which a line M is colinear with a junction J. By the general viewpoint assumption, we can deduce that M and the lines which meet at J are projections of lines which meet in 3D. The extension of M may be invisible at J because of contrast failure (as is clearly the case with the extension of line CB in Fig. 7(a)). It may also be invisible because of occlusion. For example, in Fig. 8, M_1 is the extension of a shadow line which is occluded at J. M_2 is a shadow line which is not visible at J because J is shaded by another object. Of the three edges which meet at the vertex which projects into J, only two are visible, the the third being hidden due to self-occlusion. M_3 is colinear with the projection H of this hidden edge. Note that we include the possibility that distinct 3D edges may be colinear, since this situation occurs too often in man-made objects for us to ignore it. When two visible lines in the drawing are colinear, this constrains the possible labellings for the two lines, as described in Cooper (2000).

Figure 9 lists all the possible forms of extended junctions involving a total of up to four lines. For example, the extended junction A-BC in Fig. 7(a) is of type L[K]. For each type of extended junction given in Fig. 9, Appendix A gives the list of legal labellings obtained by enumerating all possible cases. All the assumptions given in Sections 2 and 3 are required to establish this catalogue, including the very strong extended trihedral assumption. Note that reflected versions of all types of extended junctions exist (except for L[W] and L[Y] junctions), with slightly different lists of labellings.

Applying the L[K] constraint, defined by the list of labellings in Appendix A, to the extended junction A-BC in the drawing in Fig. 7(a), tells us that junction A has one of the four labellings $\rightarrow +$, $+\rightarrow$, $\leftarrow +$ or $+\leftarrow$ (given that BC cannot be a shadow line, since it is the bar of a T-junction). The global labelling

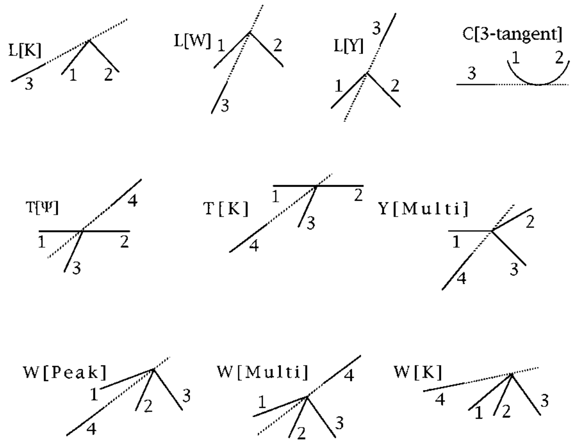

Figure 9. The types of extended junctions of total degree up to four.

which results after constraint propagation is shown in Fig. 7(a). We use the generic label Δ to represent any label from the set $\{-, \leftarrow, \rightarrow, \leftarrow\neq, \rightarrow\neq\}$. It is interesting to compare this global labelling with the result of constraint propagation on a drawing which is identical to the drawing in Fig. 7(a) except for the fact that the line AB is not subject to contrast failure. The global labelling obtained in this case is shown in Fig. 7(b), where the generic label d represents the set of labels $\Delta \cup \{+\}$. It turns out that the drawing in Fig. 7(b) is more ambiguous, due to the fact that junction A now has 12 legal labellings, instead of just 4. We call this phenomenon, in which a drawing is less ambiguous when more lines are missing, the *contrast failure paradox*.

A perfect projection is almost always more informative than an imperfect projection (due to the absence of phantom junctions on straight lines). However, if it is known that contrast failure can occur, the contrast failure paradox tells us that a line in the drawing which is partly missing due to contrast failure may provide more information than the corresponding line if it had been totally visible. In the example of Fig. 7, this is because we effectively add back the missing line with the extra information that it was missing due to contrast failure.

7. Constraint Between Vertices and the Shadows they Cast

Suppose that vertex V projects into junction J_V in the drawing and that the shadow cast by V projects into junction J_S. Such corresponding junction pairs (J_V, J_S)

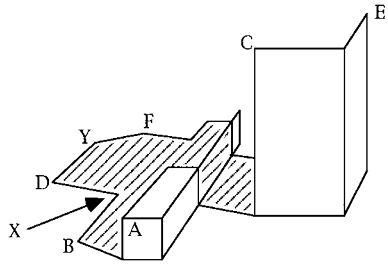

Figure 10. Examples of vertex-shadow correspondences: (A,B), (C,D) and (E,F); and examples of L-junctions, X and Y, that are not shadows of visible vertices.

may be detected, for example, from the fact that they are colinear with the projection of a known light source (under the general viewpoint and general light source position assumptions). It turns out that the number of legal labellings of the junction pair (J_V, J_S) is very limited. We call the resulting constraint the *vertex-shadow constraint*. For example, establishing that B, D, F in Fig. 10 are the shadows of A, C, E, respectively, reduces the number of legal labellings for A from 9 to just 1, for B, D, F from 31 to just 1 and for C, E from 34 to just 3. It should be noted that junctions composed uniquely of shadow lines are not necessarily shadows of vertices. They may be the intersection of the shadows of two distinct objects, as at junction X in Fig. 10. Furthermore, a shadow-casting vertex V may be invisible in the drawing due to occlusion or contrast failure (which is the case for the vertex casting the shadow at junction Y in Fig. 10).

Trihedral vertices necessarily cast shadows J_S in the form of L-junctions. Analysis of the intensity image gives us the direction of the shadow lines ($\downarrow$ or $\uparrow$) meeting at J_S. This then provides a strong constraint on the possible labellings of J_V. We distinguish two cases, depending on whether J_S is a convex L-junction (i.e. the shaded area subtends an angle α less than π) or concave L-junction ($\alpha > \pi$). If J_S is a convex L-junction, then V must be a vertex of type 1 in Fig. 3 (such as the corner of a cube), which gives rise to the following constraint on the possible labellings of J_V.

Vertex-Shadow Constraint when the Shadow of the Vertex is Convex

If J_V is a Y-junction then it has the labelling $+ + +$. If J_V is a W-junction then it has the labelling $\rightarrow + \rightarrow$. If J_V is an L-junction then it has one of the labellings $+ \rightarrow$, $\rightarrow +$ or $\rightarrow\rightarrow$ (except that the labelling $\rightarrow\rightarrow$ is

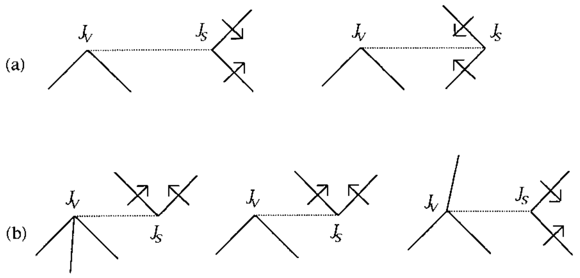

Figure 11. (a) In these two cases, the labelling →→ for J_V is illegal, since no edge could cast the upper shadow edge at J_S; (b) in these three cases, no legal labelling exists for J_V. (N.B. Only the solid lines are actually present in the drawing.)

illegal for the two configurations shown in Fig. 11(a)). Furthermore, the configurations in Fig. 11(b) are physically impossible.

For example, knowing that B is the shadow of A in Fig. 10 uniquely determines the labelling $\rightarrow + \rightarrow$ for A. Knowing that D is the shadow of C narrows down the labelling of C to three possibilities (depending on which of the three lines is missing due to occlusion or contrast failure).

If, on the other hand, J_S is a concave L-junction, then V must be a vertex of type 2, 5, 11, 13 or 15 in Fig. 3. In this case, J_V may include shadow lines, since these vertices can cast local as well as distant shadows. Given the technical complexity of the resulting constraint on the labellings for J_V, we relegate the details to Appendix B.

The vertex-shadow constraint is a very strong constraint. For example, if J_S is convex and J_V is either a Y or W junction, then the labelling of J_V is uniquely determined. Even more striking is the fact that, whatever the form of J_V, the labelling of J_S is always uniquely determined. However, we should be wary of applying the vertex-shadow constraint unconditionally to all pairs of junctions (J_V, J_S) which are colinear with a light source, to within an angle of error ε. The number of spurious correspondences thus found increases as $(\varepsilon/\pi)n^2$.

8. Tractability of the Line Labelling Problem

The classic approach to the interpretation of line drawings is to assign semantic labels to the lines of the drawing in accordance with a catalogue of legal junction labellings. The catalogue, giving the list of legal labellings for each type of junction which may occur in the drawing (L, Y, T, etc.), is derived from assumptions on object shape and image formation. In certain cases, the existence of a legal global labelling of the drawing is not only a necessary condition but also a sufficient condition for the drawing to be realisable as the projection of a physically possible 3D scene (Cooper, 1999).

For a given catalogue, we use the term LDLP (line drawing labelling problem) to denote the computational problem of determining whether a drawing given as input has at least one global labelling consistent with the catalogue. It is known that the LDLP is NP-complete for perfect projections of polyhedral scenes (Kirousis and Papadimitriou, 1988) but that the LDLP is solvable in linear time for perfect projections of curved objects (Cooper, 1999). When the LDLP is intractable, the propagation of constraints (Tsang, 1993; Waltz, 1975) has nevertheless proved to be a powerful tool for reducing the ambiguity in the labelling problem by filtering out labels which cannot be part of a global legal labelling.

A catalogue of labelled junctions can thus be judged according to two important criteria: (1) is the associated LDLP tractable, and (2) what is the predictive power of the catalogue. The predictive power measures the number of bits of information per line-end provided by the catalogue (Cooper, 1997a) and can easily by used to estimate the probability that a random labelling of a drawing will actually be a legal labelling. This section is exclusively concerned with the tractability of different versions of the LDLP in the presence of contrast failure and/or shadows. The corresponding catalogues of junction labellings are given in Appendix C.

This section presents two main results. Firstly, the possibility of missing lines due to contrast failure between parallel surfaces renders the LDLP solvable in linear time, even for polyhedral scenes, provided there are no shadows. Secondly, in all cases, in the presence of shadows the LDLP is NP-complete.

Theorem 1. *The LDLP is solvable in linear time for drawings subject to contrast failure, but without shadows.*

Proof: This result follows from the possible existence of undetectable phantom junctions on any line in the drawing, meaning that any label in the set $\Delta = \{-, \leftarrow, \rightarrow, \leftarrow\neq, \rightarrow\neq\}$ can be transformed into any other label in this set. Such label transitions, caused by T-junctions with missing stems due to contrast failure, can occur on straight as well as curved lines. The original LDLP is, hence, equivalent to a LDLP in which

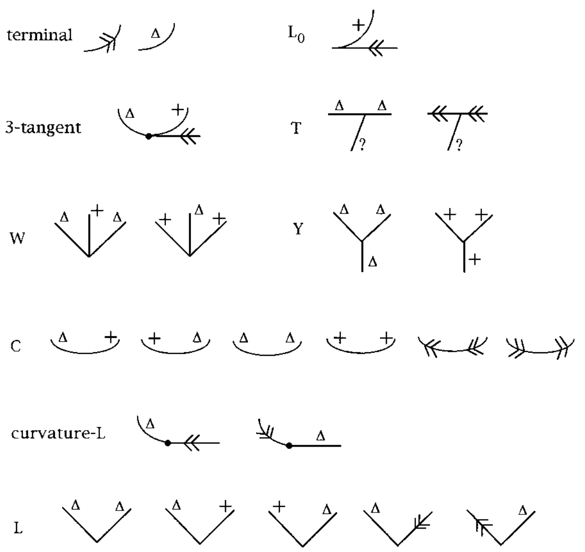

Figure 12. The catalogue of labelled junctions for line drawings with missing lines due to contrast failure between parallel surfaces, but no shadows.

all the labels in the set $\{-, \leftarrow, \rightarrow, \leftarrow\neq, \rightarrow\neq\}$ are replaced by a unique label Δ. The resulting catalogue of labelled junctions is given in Fig. 12. A question mark denotes any of the four labels: Δ, $+$, $\twoheadrightarrow$, $\twoheadleftarrow$. Reflected versions of curvature-L and 3-tangent junctions also exist. □

The constraints at L_0 and 3-tangent junctions can be replaced by a combination of unary constraints. The constraints at T, W and Y junctions can be replaced by a combination of binary constraints. For example, a W junction joining lines with labels x, y, z can be replaced by

$$(x = \text{`}\Delta\text{'} \ \vee \ y = \text{`}\Delta\text{'}) \wedge \ (x = \text{`}+\text{'} \vee \ y = \text{`}+\text{'}) \\ \wedge \ (z = \text{`}\Delta\text{'} \vee y = \text{`}\Delta\text{'}) \\ \wedge \ (z = \text{`}+\text{'} \vee y = \text{`}+\text{'})$$

On curved lines, transitions between $+$ and Δ can occur, meaning that these labels can be replaced by the even more generic label $\mathrm{d} = \Delta \cup \{+\}$. On lines joining two L junctions, the label Δ is neighbourhood substitutable for all other labels (Cooper, 1997b; Freuder, 1991) and hence all such lines can immediately be labelled Δ without altering the solvability of the problem. In the resulting reduced problem, each line has at most two possible labels, and all constraints are unary or binary. The LDLP has thus been reduced to 2SAT, for which there is a known linear time algorithm (Melhorn, 1974).

Theorem 2. *The LDLP is solvable in linear time for drawings subject to contrast failure, and containing shadow lines which have all been identified as such.*

Proof: Consider the line drawing which results when all shadow lines have been erased. The remaining junctions are exactly those listed in Fig. 12. The presence of shadows may restrict the possible labellings for a given junction. For example, the presence of a shadow line leaving a Y junction prohibits the labelling $(+, +, +)$. Nevertheless, it is easy to verify that the reduction to 2SAT, described above, still holds. It is noteworthy that junctions which are only detectable due to the presence of shadows, such as junctions B and G in Fig. 5, provide no essentially new type of constraint, once shadows have been erased. □

Although we omit the details, it is easily verified that the reduction to 2SAT still holds when extended junction constraints are applied to line drawings without shadows or with shadows identified as such. A similar remark holds for the vertex-shadow constraints, which, of course, can only be applied to line drawings with shadows identified as such.

Theorem 3. *The LDLP for line drawings with shadows is NP-complete, whether or not contrast failure between parallel surfaces can occur.*

Proof: NP-completeness follows immediately from the following reduction from PLANAR 3SAT (Lichtenstein, 1982) in which an arbitrary instance of PLANAR 3SAT is represented by a drawing to be labelled. It is sufficient to demonstrate constructions for

1) producing multiple copies of the same variable v;
2) negating a variable $\neg v$;
3) the disjunction of three variables $v_1 \vee v_2 \vee v_3$.

We associate the value *false* with the label 'shadow'. (Note that the orientation of the shadow label is implicit in the difference in intensity between the two sides of the corresponding edge in the intensity image, and so does not need to be specified). We associate the value *true* with the generic label $\mathrm{d} = \{+, -, \leftarrow, \rightarrow, \leftarrow\neq, \rightarrow\neq\}$. All lines in the drawing under construction are curved, meaning that any two labels in the set d can be transformed one into the other by a sequence of phantom junctions on any line.

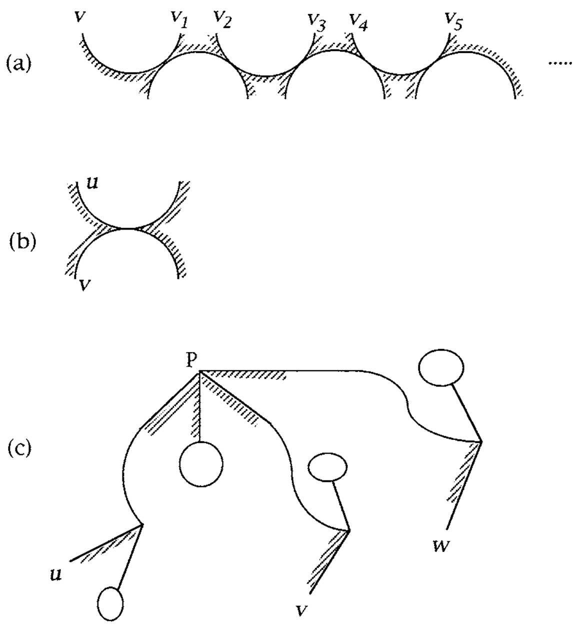

Figure 13. Constructions required for the reduction of PLANAR 3SAT to the LDLP with shadows: (a) $v = v_1 = v_2 = \ldots$ (b) $u = \neg v$ (c) $v_1 \vee v_2 \vee v_3$.

Figure 13 gives the three constructions mentioned above. In Fig. 13(a), all the variables v_i must have the same value (d or 'shadow') as v. This follows from the set of legal labellings for X_0 junctions, given in Appendix C. The construction in Fig. 13(b) represents negation, since u = d iff $v =$ 'shadow'. In Fig. 13(c), it is easy, though tedious, to verify that variables $u,\ v,\ w$ can simultaneously take on all combinations of values except (shadow, shadow, shadow). Such an assignment of values would imply that the first, third and fourth lines meeting at the Peak junction P would have to be simultaneously labelled d, which is not a legal labelling. In each construction in Fig. 13, the shading next to a line indicates which side of the corresponding edge in the intensity image is darker.

The reduction from PLANAR 3SAT is now complete. Note that the constructions work whether contrast failure between parallel surfaces can occur or not.

The tractability results proved in this section (Theorems 1, 2, 3) are robust to the changes in the catalogue of junction labellings which follow when discontinuities of surface curvature are allowed (Cooper, 1993) or when object edges and surfaces may meet tangentially (Cooper, 1997a).

Since the assumption that contrast failure can only occur between parallel surfaces may, in practice, be too optimistic, we should also consider the LDLP in which contrast failure can occur at any non-shadow edges. We call this phenomenon generalised contrast failure. In the case of line drawings without shadows, the constraints under the assumption of generalised contrast failure are even weaker than those given in Fig. 12, and the LDLP is still solvable in linear time.

When generalised contrast failure can occur in drawings of scenes with shadows, tractability depends on whether there is just one or many point light sources. With multiple point light sources the LDLP becomes almost trivial: any line in the drawing, whose corresponding edge in the intensity image does not prohibit the label 'shadow', can be labelled 'shadow' and subtracted out of the drawing without any effect on the solvability of the LDLP. The drawing that remains when all such lines have been subtracted out is an instance of the LDLP without shadows. We have just seen that this problem is solvable in linear time. On the other hand, if we know that the scene is illuminated by at most one point light source (and any number of non-point light sources), then the LDLP becomes NP-complete again. The only change required in the proof of NP-completeness is that the construction for $v_1 \vee v_2 \vee v_3$ in Fig. 13(c) must be replaced by a simpler construction consisting of a single W-junction. Indeed, under the assumptions of generalised contrast failure and a single point light source, all combinations of the labels d and 'shadow' are legal for a W-junction, except ('shadow', 'shadow', 'shadow').

We conclude that the tractability results proved in Theorems 1, 2, 3 are not merely consequences of arbitrary assumptions about object shape or the kind of contrast failure that can occur, but instead are quite general results. Contrast failure makes the LDLP solvable in linear time, whereas the problem becomes NP-complete for drawings with shadows.

9. Predictive Power of the Junction Labelling Constraints

Another important criterion for judging a set of constraints is their strength. The stronger the constraints, the less the expected number of spurious interpretations of the drawing. We can quantify the notion of strength of a catalogue of junction labellings by calculating its predictive power, which is defined as the number of bits of information that the catalogue provides per line-end. A necessary (though not sufficient) condition for the expected number of spurious global labellings of a

Table 1. A comparison of the predictive power of different catalogues.

	pp	$pp_{\max}/2$
No CF, no shadows	3.22	3.00
No CF, shadows	4.28	3.59
CF, no shadows	2.44	3.00
CF, shadows	3.57	3.59

drawing to tend to zero, as the number of lines in the drawing tends to infinity, is that $pp > pp_{\max}/2$, where pp is the predictive power and $pp_{\max}$ is the theoretical upper bound on pp (Cooper, 1997a). Table 1 presents the values of pp and $pp_{\max}/2$ for the catalogues given in Appendix C under four different assumptions: drawings with or without contrast failure (CF), drawings with or without shadows.

In calculating the predictive power of catalogues involving shadows, we assume that analysis of the intensity image allows us to determine, for each line-end at a junction, which side of the line is darker. For example, analysis of the intensity image at an X-junction reduces the number of legal labellings for the junction from four to just one. On the other hand we are pessimistic in that we assume that all lines in the drawing are curved, which implies that any number of $+\leftarrow$ transitions can occur an any line.

The results of Table 1 confirm something that is well known to artists and photographers, that a scene with shadows is less ambiguous than the same scene without shadows. In general, a line drawing with shadows has a smaller expected number of spurious interpretations than the same drawing without shadows. Unfortunately, when contrast failure can occur, the value of pp falls below the critical value $pp_{\max}/2$, meaning that the catalogue of junction labellings does not provide sufficient information to uniquely label line drawings subject to contrast failure. The dramatic drop in the value of pp, when contrast failure is allowed, is due to the fact that transitions such as $\rightarrow -$ are legal on any line in the drawing. Such transitions are caused by contrast failure at vertices such as number 5 in Fig. 3.

The possibility of missing lines due to contrast failure thus increases the importance of other constraints, such as the identification of non-shadow lines by analysis of the intensity image (Section 3), the straight shadow line constraint (Section 5), the extended junction constraints (Section 6), the vertex-shadow constraints (Section 7), the absence of $+\leftarrow$ transitions on straight lines, the colinear lines labelling constraint (Cooper, 2000), and the identification of extremal lines and $+\leftarrow$ transitions by analysis of the intensity image (Cooper, 1997; Malik and Maydan, 1989).

10. Coplanarity Constraints from Straight Edges

The coplanarity of four 3D points can be deduced from the presence of a straight line in the drawing, under the straight edge formation assumption (see Section 2). The case when the drawing is not subject to contrast failure, which occurs for example when the drawing is human-entered, has been treated in detail in a previous paper (Cooper, 2000). The constraints derived in this section cover the generalisation of these coplanarity constraints to the case of possibly missing lines due to contrast failure. Although coplanarity constraints can be used to establish linear equations between the depths of object vertices (Cooper, 2000), they can also be used to simply group together 3D edges lying in the same plane. This makes explicit information which could be useful for later tasks such as object recognition, for example.

If we assume that all vertices are trihedral, then coplanarity constraints can still be applied. For example, the existence of three lines, identified as non-shadow lines, meeting at each junction implies that no contrast failure has occurred. This constraint is given in Fig. 14. A straight edge joins the vertices j and k. If the edge leaving j in direction of i is also a straight edge, then we can take i to be the next vertex encountered on this edge, otherwise we can let i be an arbitrary point lying on the tangent to the edge. For example, in Fig. 15, applying the constraint to the straight edge 2–3, allows us to deduce that the 3D points 1, 2, 3, 4 are coplanar. The common plane may, in fact, be a partially or totally hidden surface. For example, the constraint applied to the straight edge 4–5 in Fig. 15 tells us that points 3, 4, 5, 6 are coplanar, even though the common plane is hidden at points 4, 5, 6.

Note that any number of shadow edges may also be present at j and k, even though they are not shown in

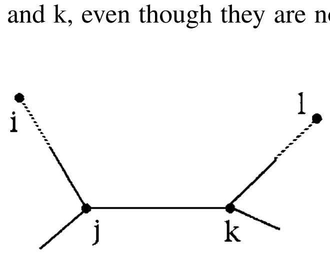

Figure 14. The coplanarity of i, j, k, l follows from the fact that jk is a straight edge and that all three edges are visible at each junction.

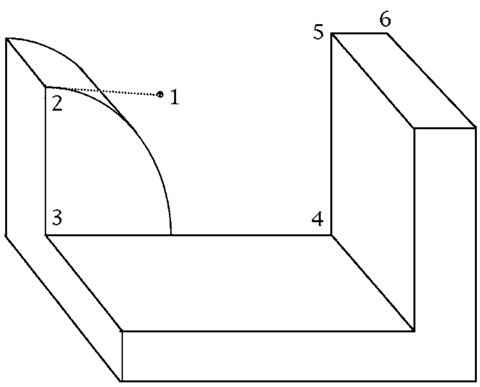

Figure 15. The coplanarity of points 1, 2, 3, 4, 5, 6 can be deduced from three applications of the coplanarity constraint of Fig. 14.

Fig. 14. The three non-shadow lines meeting at j or k form Y or W junctions, but not T-junctions. However, any number of T-junctions may occur on the straight line jk without affecting the validity of the constraint. For example, the constraint applied to the straight edge 3–4 in Fig. 15, tells us that points 2, 3, 4, 5 are coplanar.

This coplanarity constraint can be applied to show that certain impossible drawings are unrealisable as 3D objects. In Fig. 16, we obtain the labels shown after applying the extended junction constraint at junctions A and C, together with the straight shadow line constraint to lines BD, DF, BF, FE. These line labels indicate that ABCD, CDEF and EFAB are three sets of coplanar points. However, this is impossible since the projections of lines AB, CD and EF (when extended to infinity in both directions) do not intersect at a point in the drawing.

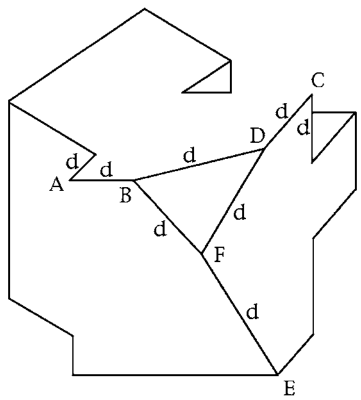

Figure 16. This drawings is impossible due to the coplanarity constraint.

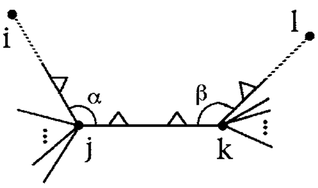

Figure 17. The coplanarity of i, j, k, l follows from the fact that jk is a straight edge, whenever we have sufficient evidence to deduce that i, j, k, l all lie on the same surface.

Coplanarity can still be deduced even in some cases when less than three non-shadow lines have been identified at each junction. This is illustrated in Fig. 17. As in Fig. 14, any number of lines identified as shadows can be added without affecting the validity of the constraint. Apart from the lines ij, jk and kl, any number of other lines of any type may occur, as shown, always lying on our right hand side as we follow the path ijkl. Unlike the constraint of Fig. 14, in this constraint the junctions j and k may be T-junctions. In fact, they may be any junctions satisfying $\alpha \neq 0 \wedge \alpha \neq \pi$ (and $\beta \neq 0 \wedge \beta \neq \pi$). This constraint does not even require the trihedral vertices assumption.

However, the coplanarity constraint of Fig. 17 is only valid for certain labellings of the lines ij, jk and kl in the vicinity of the junctions. The generic label $\triangleright$ on a vertical line represents any type of edge which, if formed by the intersection of two surfaces, is such that the surface which is visible just to the right of the line is one of these surfaces. Equivalently, we could replace $\triangleright$ in the statement of this constraint, by the set of possible labels $\{+, -, \uparrow, \uparrow\neq$, 'shadow'$\}$. This constraint is only valid if it is known that the pair of edges ij and jk (and the pair of edges jk and kl) meet in 3D. We can deduce that the edges ij and jk meet in 3D if at least one shadow line is present at j, or if $\alpha < \pi$ (given the restrictions on the labelling for j). If three non-shadow lines have been identified at j, and meet at a Y or a W junction, then no restrictions are necessary, neither on the angle α nor on the labels for ij and jk at j.

The constraint of Fig. 17 was deduced from examining all junctions under the assumption that contrast failure can only occur between parallel surfaces.

Note that in the coplanarity constraints of Figs. 14 and 17, if the point i (or point l) is a T or L junction, then it may, in fact, be the projection of two distinct 3D points. We can avoid any ambiguity by saying that it is the 3D lines ij, jk, kl which are coplanar, rather than the points i, j, k, l.

11. Results on some Sample Drawings

This section demonstrates the information which can be deduced from a line drawing subject to contrast failure. The result of applying the constraints described in this paper to a drawing consists in certain restrictions on the possible labelling of the drawing together with a list of sets of coplanar points.

Many other constraints could also be applied, concerning sets of colinear points, sets of straight lines which intersect at a point (Cooper, 2000), or sets of parallel lines whose projections meet at a vanishing point (Parodi and Torre, 1994). If vanishing points are detected (Straforini et al., 1993; Tai et al., 1993), however, it should be noted that sets of shadow lines $S_1, \ldots, S_r$ may converge to a point which is not a vanishing point. This occurs, for example, if $S_1, \ldots, S_r$ are shadows of parallel edges. Unless restrictions are placed on the shape of object surfaces, extremal lines may also converge to a common point in the drawing which is not a vanishing point. The extended trihedral assumption is necessary to identify the intersection of four (or more) lines, each labelled with the generic label d, as a vanishing point.

Figure 18 shows a drawing involving shadows and contrast failure. The labels shown for lines or line-ends are the result of applying the catalogue given in Appendix C and then filtering by constraint propagation (Tsang, 1993). The generic labels d and Δ are shown when it is not possible to uniquely determine a label. The absence of a label on a line indicates that we have not been able to obtain sufficient restictions on the line label to limit it to one of the generic labels d or Δ. To avoid cluttering up the figure unnecessarily, we have used generic labels for lines (or line-ends) rather than the actual lists of possible junction labellings. In fact, many junctions have a much smaller number of possible labellings than indicated. Junctions 1, 5 and 8, for example, each have no more than three possible labellings.

The coplanarity constraints also allow us to deduce the following facts:

(1) points 1, 2, 8, 7 are coplanar;
(2) points 2, 3, 9, 8, are coplanar;
(3) points 3, 4, 13, 9, 10, 11, 12 are coplanar;
(4) points 1, 2, 3, 4, 5, 6 are coplanar;
(5) points 7, 8, 9, 10, 11 are coplanar;
(6) points 13, 4, 5, 14 are coplanar.

The rule implicitly used for merging sets of coplanar points S_1 and S_2 is that S_1 and S_2 can be merged if $S_1 \cap S_2$ contains at least three non- colinear points. Since both T and L junctions can be caused by one edge occluding another distant edge, they may in fact be the projection of two distinct points in 3D. Thus, care must be taken when applying those coplanarity constraints above involving the junctions 1, 6, 7, 11, 13 and 14. For example, the point 13 which is coplanar with points 4, 5 and 14 is the point projecting into the T-junction and lying on the line 4–13.

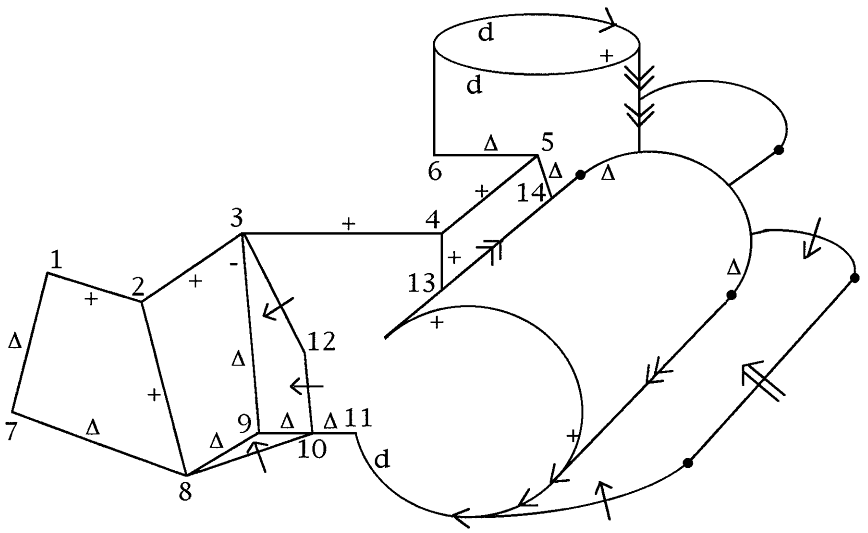

Figure 18. A drawing with shadow lines and contrast failure between parallel surfaces.

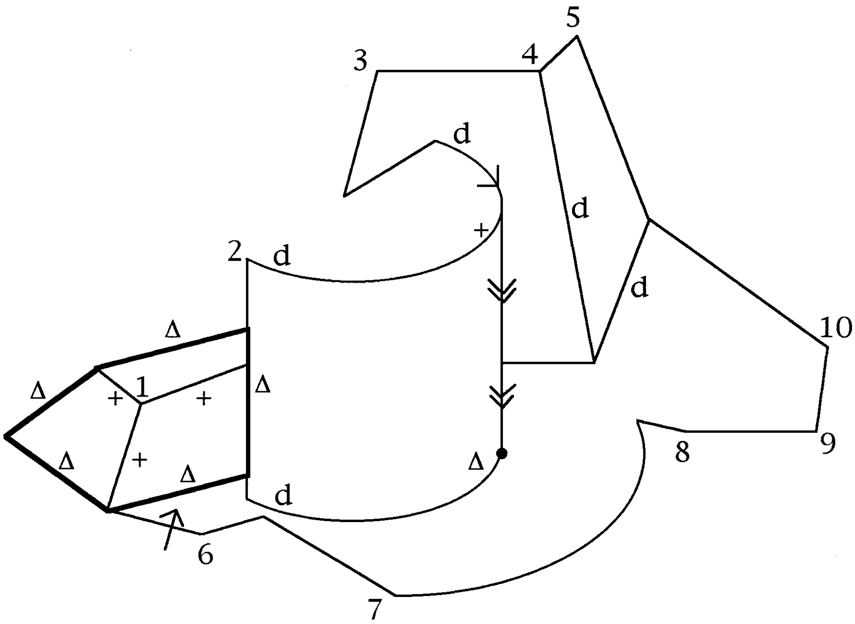

Figure 19. A drawing involving shadows and contrast failure between parallel surfaces.

Figure 19 shows another example drawing involving shadows and contrast failure. The thicker lines are assumed to have been identified as non-shadow lines by the shadow edge constraint (see Section 3). For example, this would be the case if, at some point along the line, there were a simultaneous increase in the red component and decrease in the green component as we cross the line. The identification of these lines as non-shadow allows us, by propagating junction labelling constraints, to identify the convex edges of the left-most object. The planarity of the three visible surfaces of this object (or rather the edges bounding these surfaces) can then be deduced from several applications of the coplanarity constraints. The middle object, the half cylinder, is successfully labelled as shown, with little ambiguity, due to the presence of 3-tangent, curvature-L and T-junctions.

However, practically no information is obtained concerning the right-most object or the cast shadow lines in Fig. 19. An obvious idea is to apply a heuristic rule that prefers labellings involving as few missing lines as possible. This would allow us to find the correct labelling of the cast shadow lines, but would give a completely erroneous labelling of the right-most object. One such erroneous labelling is shown in Fig. 20. Another idea is to try to determine the position of the light source by generating half-lines passing through pairs of junctions

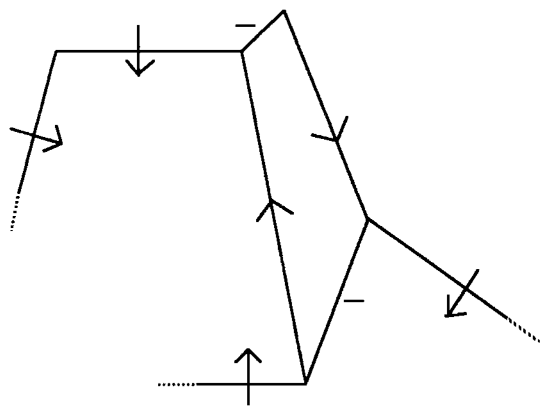

Figure 20. An unlikely, but legal labelling, for the right-most object in Fig. 19, which does not require lines to be missing due to contrast failure.

which could be projections of a vertex and its shadow, and looking for points where several of these lines intersect. If we suppose that the position of the light source in the drawing of Fig. 19 has been determined in this way or is known a priori, then we can apply the vertex-shadow constraint to the pairs of junction (1, 6), (2, 7), (3, 8), (4, 9), (5, 10). The global labelling which results after constraint propagation is shown in Fig. 21. A consequence of the tighter constraints on junction

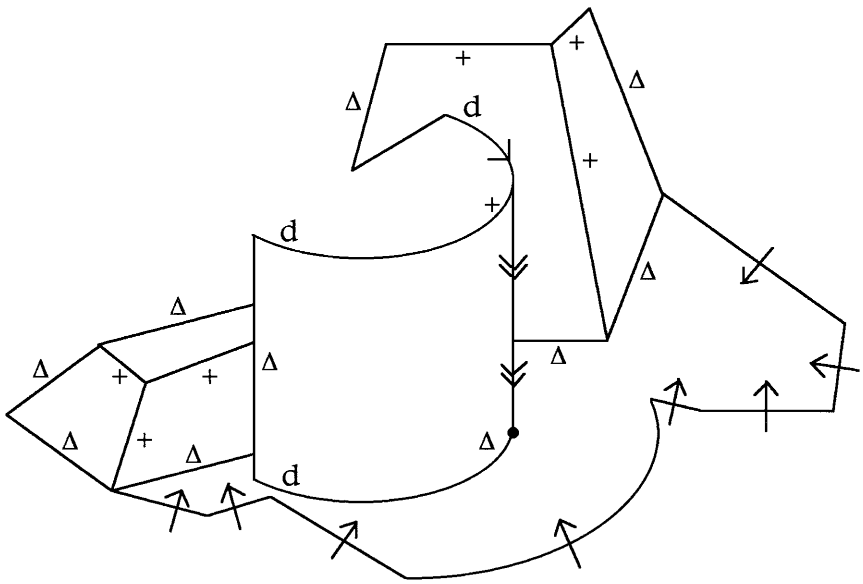

Figure 21. The global labelling of the drawing of Fig. 19 after application of the vertex-shadow constraint.

labellings thus determined, is that the planarity of the two visible faces of the right-most object can now be deduced.

12. Sufficient Conditions for Realisability

An interesting theoretical question is what are the necessary and sufficient conditions for a line drawing to be realisable as the projection of a physically possible 3D scene. This is called the realisability problem and has been solved for polyhedral scenes (Sugihara, 1984) and scenes composed of various classes of objects with curved surfaces (Cooper, 1997a; Cooper, 1999; Cooper, 2000). The constraints on line drawings with contrast failure and shadows stated in this paper are clearly necessary but not sufficient conditions for realisability. Extending the set of constraints so that they become a sufficient condition for realisability would require adding inequality constraints on the 3D positions of vertices (of the form "point A is in front of the plane BCD") to encode the three dimensional meaning of the labels $+$, $-$, $\rightarrow$, $\leftarrow$, $\rightarrow\neq$, $\leftarrow\neq$ (Cooper, 2000; Sugihara, 1984). For example, if three straight lines AB, AC, AD meet at the Y-junction A labelled $+ + +$, then point D must be behind the plane ABC. Many of these inequality constraints still hold even when contrast failure may occur. We have not stated them because they are not as informative as coplanarity constraints and they do not propagate as easily as semantic labelling constraints.

Another important constraint which has not been applied in this paper is the correspondence between edges and the shadow-edges they cast. We would have to consider correspondences between edges and the shadows they cast at shadow-producing junctions, at extended junctions and at vertex-shadow correspondences. An elegant unified approach is to assign to each shadow line-end labelled $\downarrow$ or $\uparrow$ an extra shadow-casting-edge label representing the edge E which casts it. This shadow-casting-edge label is "CF" if E is missing in the drawing due to contrast failure and is "OC" if E is missing due to occlusion. The junction labelling constraints, the extended junction constraints and the vertex-shadow constraints can all be modified to include shadow-casting-edge labels with each $\downarrow$ or $\uparrow$ label. Although a detailed discussion of shadow-casting-edge labels is beyond the scope of this paper, it is worth noting that such labels can provide extra information concerning semantic line labels. This follows directly from the fact that concave edges cannot cast shadows. Thus, for example, if the straight edge AB is known to cast the shadow CD, with C the shadow of vertex A and D the shadow of vertex B, then no semantic label transitions (such as $- \rightarrow$) are possible between A and B.

Determining correspondences between edges and the shadow edges they cast would also allow us to apply extra coplanarity constraints. For example, knowing that edge 12–10 is the shadow of edge 2–8 in Fig. 18, we can deduce that points 2, 8, 10, 12 are coplanar. This follows from the fact that the shadow of a straight edge E lies in the plane passing through E and the light source. Under an assumption of C^3 surfaces, we could also deduce that points 3, 4, 9, 12 are coplanar in Fig. 18, since the shadow line 3–12 must lie in the surface on which it is cast.

The presence of point light sources casting shadows also implies certain inequality constraints. We have seen that the presence of a shadow at a vertex may imply certain coplanarity constraints. A less obvious constraint is that the *absence* of a visible shadow at a vertex implies restrictions on the relative positions of the light source(s) and the tangents to the surfaces meeting at the vertex.

In conclusion, the elaboration of necessary and sufficient conditions for the realisability of line drawings with contrast failure and shadows remains an open problem. Nevertheless, the constraints stated in this paper amply demonstrate the feasibility of the analysis of line drawings with contrast failure and shadows.

13. Conclusion

When analysing line drawings of curved objects in which lines may be missing due to contrast failure, the classic technique of propagating junction labelling constraints does not provide sufficient information to adequately interpret the drawing. This is mainly due to the possibility of undetectable label transitions on lines. However, the extra information required to adequately interpret the drawing can often be obtained from various constraints concerning shadow formation and extended junctions (a straight line colinear with a junction).

Although the presence of shadows in a drawing is what causes the labelling problem to become NP-complete, shadows are a valuable source of information since the drawing is, on average, less ambiguous.

Further information, concerning the shape of objects in the drawing, can be obtained from coplanarity constraints between sets of object vertices, derived from the presence of straight lines in the drawing. These constraints have been shown to be robust to the presence of contrast failure.

Appendix A: Legal Labellings for Extended Junctions

In the following sets of legal labellings, the label d can be replaced by any of the labels +, −, →, ←, →≠, ←≠. Horizontal arrows represent occluding edges whereas vertical arrows represent shadows.

L[K]: → +d, + → d, ← +d, + ← d, ← + ↑,
→ − ↓, − →↓, → →≠↑, →≠→ ↓, → → ↓,
←≠→ ↓, ←→ ↓ .

L[W]: → −d, − → d, − → ↑, → − ↓, → → d,
→ → ↑, → → ↓, → →≠d, →≠→ d,
→≠→ ↑, → →≠↓, + ← ↓, ← + ↑,
←→ ↓, ←→ ↑, → ← ↓, → ← ↑ .

L[Y]: → −d, − → d, → → d, → →≠d,
→ ≠→ d, − → ↑, → − ↓, → → ↑, → → ↓,
→ ≠→ ↑, → →≠↓, → ← ↓, → ← ↑,
→ ←≠↓, ←→ ↓, ←→ ↑, ←≠→ ↑ .

C[3-tangent]: −− ⇑ .

T[Ψ]: − ← + ↑, − ← ← d, ← − +d, ←≠← + ↑,
← ← + ↑, ← ←→ ↑, ← ← ← ↑, − ← ← ↑,
←≠← ← ↑, ← − + ↓, ← ←≠+ ↓,
← − → ↓, ← ←≠→ ↓, ← ← ↓↓, ← ← ↑↓,
← ← ↓↑, → → ↓↓, → → ↑↓, → → ↓↑ .

T[K]: − ← ← d, ← − +d, ← − + ↑, ← ←≠ + ↑,
← ← ← ↑, ← ← ← ↓, ← ←≠ + ↓, → →≠↓↓,
→ − ↓↓, →≠→ ↑↑, − → ↑↑, → → ↑↑,
→ → ↓↑, → → ↓↓ .

Y[Multi]: ← − ↑d, − ↑→ ↑, → ↓ −d, → ↓ − ↓,
→≠↑→ d, →≠↑→ ↑, → ↓→≠ d,
→ ↓→≠↓, → ↑→ d, → ↓→ d,
→ ↑→ ↑, → ↑→ ↓, → ↓→ ↑, → ↓→ ↓,
→ → − ↓, − ←→ ↑, ↓→ ←≠↓,
← ←≠↓↑, ↓→≠← ↑, ←≠← ↓↑,
← ← ↓↓, ← ← ↑↑, ← ← ↓↑,

←←↑↓, ↓→←↓, ↑→←↓,
− ←↓↑, ← −↑↓, ↓→ −↓,
→≠←↓↑, ←→≠↑↓, ↓←→≠↓,
→←↑↓, →←↑↑, →←↓↑, →↓←↓,
→↑←↓, →↓←↑, →←↑↓, ←→↑↑,
←→↓↑, ↑→→↓, ↓→→↓,
←↓→↓, ←↑→↓, ←↓→↑ .

W[Peak]: − ↑→ d, → − ↓d, →↓−d, − →↓d,
− →↓↑, →− ↓↓, →+ →↓, →+ →↑,
→→↓d, →→↓↓, →+ →≠↓,
→≠ + →↑, →≠→↓d, →≠→↓↑,
→→≠↓d, →→≠↓↓, + − + ↓,
→ − →↓, ←↑ + ↑, +↓←↓, ↓→ − ↓,
↓− →↓, ↓→≠→↓, ↓→→≠↓, ↓→→↓,
→↓− ↓, − ↑→↑, ↓←≠→↓,
→↑←↑, →↓←↑, →↓←↓, →→↓↑,
←↑→↑, ←↓→↑, ←↓→↓,
↓←→↓ .

W[Multi]: −↑→d, → − ↓d, →↓−d, − →↓d,
→→↓d, →≠→↓ d, →→≠↓ d,
↓→ −↓, ↓− →↓, →≠+ →↑,
→≠→↓↑, ↓→≠→↓, ↓→→≠↓,
→←≠↓↓, → + →≠↓,
↓→→↓, →←↓↓, →←↓↑,
→−↓↓, →↓−↓, −↑→↑,
→→≠↓↓, →→↓↑, →→↓↓,
→↓←↓, →↓←↑, →↑←↓, ←↓→↓,
←↓→↑, ←↑→↓ .

W[K]: ←↑+d, + ↓←d, ←↑+↑, − + − ↓,
↓→ −↓, → −↓↓, ↓− →↓,
− ←↓↓, →≠ + →↓, ↓→≠→↓,
→≠←↓↓, ↓→→≠↓, →←≠↓↓,
→ + →≠↓, → + →↓, ↓→→↓,
→←↓↓, − ↑→↓, ↓←≠→↓,
←↓→↓, ←↑→↓, ↓←→↓ .

Appendix B: Vertex-Shadow Constraint when the Shadow of the Vertex is Concave

It is assumed that junction J_V is the projection of a vertex V whose shadow is visible in the drawing as the concave L-junction J_S. The constraints given below concern only the *non-shadow* lines meeting at J_V. Thus, applying the vertex-shadow constraint to $(J_V,\ J_S)$ means eliminating from the list of labellings for J_V all labellings whose restriction to non-shadow lines does not occur in the appropriate list below.

Very rarely, the presence of a shadow line in a labelling for J_V may invalidate the correspondence $(J_V,\ J_S)$. In fact, this occurs in only one case. Suppose that the labelling for J_V is one of the labellings given in Fig. 22 (together with any number of extra shadow lines added to the W-junctions shown). Then this identifies V as a type 3 vertex (according to the numbering scheme of Fig. 3), but such vertices cannot cast distant shadows. Thus such labellings should also be eliminated as illegal.

Figure 23 shows the six types of junctions which can occur at J_V, and the three types of junction which can occur at J_S. The dotted line represents the line joining J_V and J_S. For all physically possible junction pairs $(J_V,\ J_S)$, we give below the list of legal labellings for J_V. Reflected versions of most of these constraints also exist.

((1), (a)): + →, → +, →→, ←→, →← .
((2), (a)): → −, →→≠, →→, ←← .
((3), (a)): →→ −.
((5), (a)): →→ −, →→→≠, →→→, + → −,
+ →→≠, + →→, ←→→ .
((6), (a)): ←←→, ←← +, ←←← .

((1), (b)): ← −, − ←, ←→≠, ←→, →≠←,
→← .
((2), (b)): → −, − ←, ← +, + →, ←→, →→,
←← .

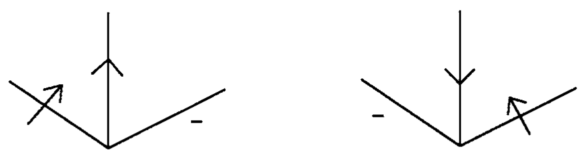

Figure 22. Illegal labellings for J_V.

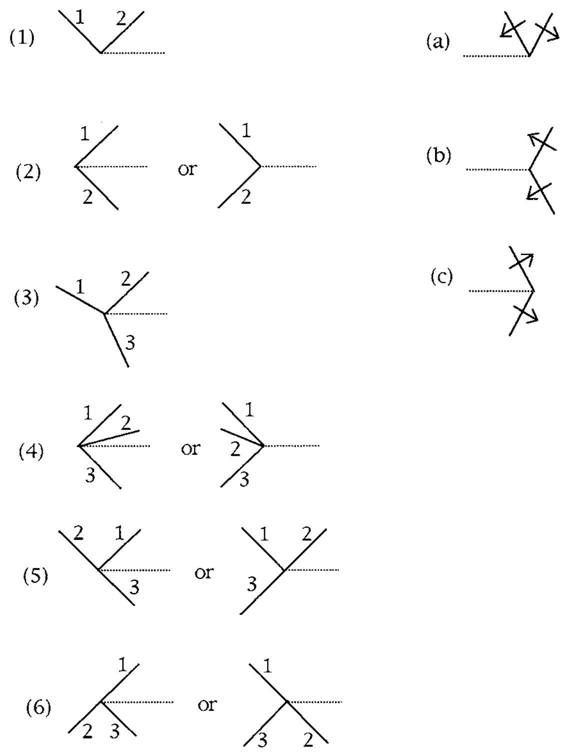

Figure 23. (1)–(6) the 6 possible types of J_V ; (a), (b), (c) the 3 types of J_S.

((3), (b)): $\leftarrow - \leftarrow, - \leftarrow\rightarrow$.

((4), (b)): $+ - +$.

((5), (b)): $\leftarrow - \rightarrow, \leftarrow\rightarrow\neq\rightarrow, \leftarrow\rightarrow\rightarrow, +- \rightarrow, + \rightarrow\neq\rightarrow, + \rightarrow\rightarrow, \rightarrow\rightarrow\rightarrow$.

((6), (b)): $\leftarrow\leftarrow\leftarrow, \leftarrow\leftarrow\rightarrow, \leftarrow - \rightarrow, \leftarrow\leftarrow\neq\rightarrow, \leftarrow -+, \leftarrow\leftarrow\neq +, \leftarrow\leftarrow +$.

((1), (c)): $\leftarrow -, - \leftarrow, \leftarrow\rightarrow\neq, \leftarrow\rightarrow, \rightarrow\neq\leftarrow, \rightarrow\leftarrow$.

((2), (c)): $\rightarrow -, - \leftarrow, \leftarrow\rightarrow, \rightarrow\rightarrow, \leftarrow\leftarrow$.

((3), (c)): $\leftarrow - \leftarrow, - \leftarrow\rightarrow$.

((5), (c)): $\leftarrow - \rightarrow, \leftarrow\rightarrow\neq\rightarrow, \leftarrow\rightarrow\rightarrow, \rightarrow\rightarrow\rightarrow$.

((6), (c)): $\leftarrow\leftarrow\leftarrow, \leftarrow\leftarrow\rightarrow, \leftarrow - \rightarrow, \leftarrow\leftarrow\neq\rightarrow$.

Appendix C: Catalogue of Junction Labellings

We give the catalogue of junction labellings in two stages: the first catalogue is the list of junction labellings in the case that there are no missing lines; the second catalogue is the list of all *extra* labellings which have to be added to the first catalogue in the case that contrast failure between parallel surfaces may cause lines to be missing. The catalogues have been derived under the assumptions of general light source positions and general viewpoint. We assume that the only vertices that may occur in the scene are the 15 vertices shown in Fig. 3. Junctions may also be caused by occlusion or by the interaction of shadows and curved surfaces (as described in Section 4). We suppose that ramp lines have not been detected in the intensity image.

It is worth noting the differences between our catalogue and that of Waltz (1975). Although Waltz allowed for accidental alignment and cracks, which we have not considered, our approach is more general in that we allow curved objects, contrast failure and multiple light sources. Because of the possible presence of multiple light sources, there is no theoretical limit to the number of lines which can meet at junctions in the drawing. We have thus been obliged to place an arbitrary limit on the degree of junctions to be included in the published catalogue. Figure 24 shows all junctions of degree less than or equal to 4. At a 3-tangent junction, the lines numbered 2 and 3 have continuous curvature, as do the lines numbered 3 and 4 at a 4-tangent junction. At a Multi junction the lines 1 and 2, if extended through the junction, would both lie within the sector bounded by lines 3 and 4.

In the comma-separated lists of junction labellings, each labelling is given in the form $L_1L_2..L_m$ where m is the degree of the junction and each L_i is the label of the line numbered i in Fig. 24. Although some of the lists are long, it should be remembered that many labellings will immediately be eliminated by analysis of the intensity image. For example, the shadow label $\uparrow$ and $\Uparrow$ are incompatible with a line saparating a lighter region above the line from a darker region below.

Catalogue of Junction Labellings without Contrast Failure

terminal: $\twoheadrightarrow, \uparrow, \downarrow, \Uparrow, \Downarrow$.

C: $+ \rightarrow, \rightarrow +, + \leftarrow, \leftarrow +, ++, --, \rightarrow\rightarrow, \leftarrow\leftarrow, \leftarrow\neq\leftarrow\neq, \rightarrow\neq\rightarrow\neq, \twoheadrightarrow\twoheadrightarrow, \twoheadleftarrow\twoheadleftarrow, \uparrow\uparrow, \downarrow\downarrow, \Uparrow\Uparrow, \Downarrow\Downarrow$.

(The first four labellings are illegal on straight lines.)

curvature-L: $\Uparrow\uparrow, \twoheadleftarrow\leftarrow, \twoheadleftarrow -, \twoheadleftarrow\leftarrow\neq, \rightarrow\twoheadrightarrow, - \twoheadrightarrow, \rightarrow\neq\twoheadrightarrow, \downarrow\Downarrow$.

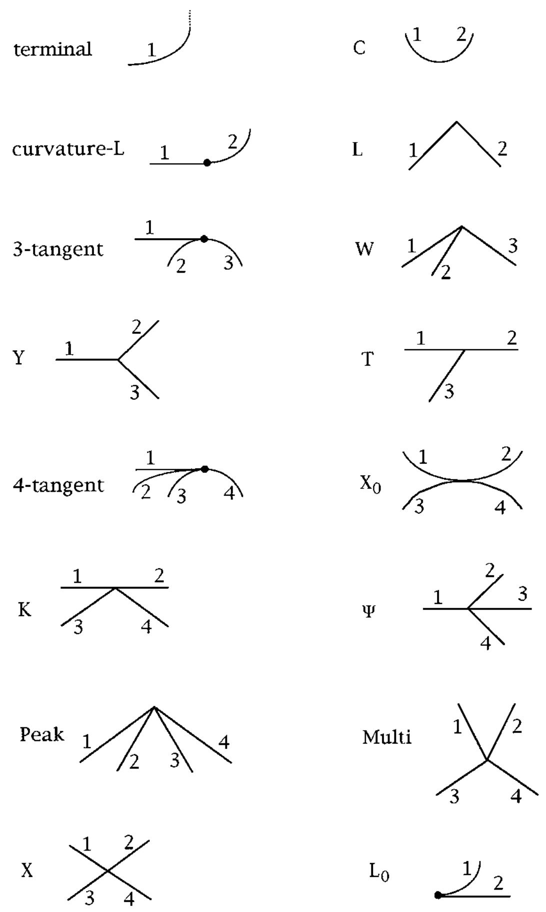

Figure 24. List of junctions of degree less than or equal to 4.

(If line 1 is straight, then only the first four labellings are legal.)

L: →→, ←←, − →, → −, ← +, + ←,
→≠→, →→≠, ↓↓, ↑↑, ⇑↑, ↑⇑, ⇑⇑ .

3-tangent: ↠ + →, ⇓ −−, ⇓→≠→≠, ↑↞↞,
↓↞↞, ↓→→, ⇑↞↞, ⇓↞↞.

(If line 1 is straight, then only the first three labellings are legal; if lines 2 and 3 are straight, then only the last five labellings are legal.)

W: − + −, → + →, → + →≠, →≠ + →,
+ −+, + ↓←, ←↑ +, ↓→ −, → − ↓, − ←↓,
↓ − →, ↓→→≠, →←≠↓, →≠←↓, ↓→≠→,
→←↓, ↓→→ .

Y: − − −, + + +, →→ −, ← − ←, − ←→,
− ↑→, ← − ↑, ↓→ −, ↑ − ←, →↓ −,
− ←↓, ↑→←, ↓→←, →↑→, →↓→,
←←↑, ←←↓, →≠↑→, ←←≠↑,
↓→←≠, ↑→≠←, →↓→≠,
←≠←↓ .

T: ←←?, ↞↞?, ← − →, ← −+, ←←≠→,
←←≠ +, − ←←, − ← +, ←≠←←,
←≠← +, ++ ↑, ++ ↓, →→↑, →→↓,
+ + ⇑, ++ ⇓, →→⇑, →→⇓ .

4-tangent: ↓↓→→ .

X_0: ←←↑↑, ←←↓↓, ↓↓→→, ↑↑→→ .

(If lines 1 and 2 are straight, then only the first two labellings are legal.)

K: − ←+↑, ← −↑ +, ←≠←+↑, ←←≠↑+,
←←↑→, ←←↓→, ←←→↑, ←←→↓,
←≠← + ↓, ←←≠↓ +.

Ψ: +↑+↓, +↓+↑, −↑−↓, −↓−↑,
+⇑+⇓, +⇓+⇑, −⇑−⇓, −⇓−⇑,
→ +−↑, → + →≠↑, →→ −↑,
→→→≠↑, −↑←→, ←≠↑←→, ←↓−+,
←↓←≠ +, − ←→↓, →≠←→↓,
←↓ − ←, ←↓←≠← .

Peak: +↓ −+, +−↓ +, →↓ − →, → −↓→,
←↑↓ +, + ↓↑←, ↓− + −, − + − ↓,
↓↓← −, ↓→ −↓, → − ↓↓, ↓↓− →,
↓ − →↓, − ←↓↓, ↓→≠ + →,
→≠ +→↓, ↓↓←≠→, ↓→≠→↓,
→≠←↓↓, ↓↓←→≠, ↓→→≠↓,
→→≠↓↓, ↓→ + →≠, → +→≠↓,
↓→ + →, →+→↓, ↓↓←→,
↓→→↓, →←↓↓ .

Multi: ↓→ − →, −↑←←, ↑− ←←, ←↓→ −,
− ←↑↑, ← −↑↑, ←+ ↓←≠, + ←≠←↑,
↑→↓←≠, ←←≠↑↑, →≠↑←↑,

↑↓→≠→, ←≠←↑↑, →↑←≠↓,
↑→≠↑←, ↑↓→→≠, + ←←≠↓,
←≠ + ↑←, ↑↑→→, ↓↓→→,
↑↓→→, ↓↑→→, ←←↑↑,
↑→↑←, ↑→↓←, →↑←↑,
→↑←↓ .

X: ↑↑↑↑, ↑↓↓↑, ↓↑↑↓, ↓↓↓↓, ⇑↑↑⇑, ⇑↓↓⇑,
⇓↑↑⇓, ⇓↓↓⇓, ↑⇑⇑↑, ↑⇓⇓↑, ↓⇑⇑↓, ↓⇓⇓↓,
⇑⇑⇑⇑, ⇑⇓⇓⇑, ⇓⇑⇑⇓, ⇓⇓⇓⇓,

Catalogue of Extra Junction Labellings due to Contrast Failure

terminal: ←, → .

C: −→, → −, − ←, ← −, →≠→, →→≠,
←≠←, ←←≠ .

curvature-L: ∅

L: ← −, ←→, ←↠, →←, − ←, ↠←, ←←≠,
←→≠, →≠←, ←≠←, →←≠, ←≠→,
+ →, → +, ←↑, ←↓, ↑←, ↓←,
→↑, →↓, ↑→, ↓→, ←⇑, ←⇓,
⇑←, ⇓←, →⇑, →⇓, ⇑→, ⇓→ .

3-tangent: ∅

W: →↓ −, − ↑→, ↓←≠→, →→≠↓, ←↓→,
←↑→, →→↓, ↓←→, →↓←, →↑← .

Y: ↑←←, →↓←, →←↓, ↓←←, →↑←,
→←↑, ↑←≠←, →↓←≠, →≠←↓,
↑→→, ←↓→, ←→↓, ↓→→, ←↑→,
←→↑, ↓→→≠, ←≠↑→, ←→≠↑ .

T: →−↓, →→≠↓, −→↑, →≠→↑ .

4-tangent: ∅

X_0: ∅

K: →→≠↓↑, → − ↓↑, →≠→↑↓, − →↑↓,
→→↑↓, →→↑↑, →→↓↑ .

Ψ: ←↓←↑, ←↓←↓, ←↑←↑, →↓→↑,
→↑→↑, →↓→↓ .

Peak: →↓−↓, →↓↑−, ↓ −↓→, −↑↓→,
→→≠↓↓, ↓↓→≠→, →↑↓←,
→↑↑←, →↓↑←, →↑←↓, →↓←↓,
→→↓↓, ←↑↓→, ←↑↑→, ←↓↑→,
↓←↑→, ↓←↓→, ↓↓→→ .

Multi: ↑↓→ −, ↑ − ↑←, ↓↑ − →, −↑←↑,
− ↑↑←, ↓← − ↓, ↑− ←↑, ←↓↓ −,
↑↓→←≠, ↑←≠↑←, ↓↑←≠→,
←≠↑←↑, ↓↑→←, ↓↓→←, ↑↓→←,
↑←↓←, ↑←↑←, →↑↑←, →↑↓←,
→↓↑←, ↑←→↑, ↑←→↓, ↓←→↑,
↓↑←→, ↓↓←→, ↑↓←→, ←↑←↓,
←↑←↑, ↑→←↑, ↑→←↓, ↓→←↑,
←↑↑→, ←↑↓→, ←↓↑→ .

X: ∅

L_0: + ↞.

References

Clowes, M.B. 1971. On seeing things. *Artificial Intelligence*, 2:79–116.

Cooper, M.C. 1993. Interpretation of line drawings of complex objects. *Image and Vision Computing*, 11(2):82–90.

Cooper, M.C. 1997a. Interpreting line drawings of curved objects with tangential edges and surfaces. *Image and Vision Computing*, 15:263–276.

Cooper, M.C. 1997b. Fundamental properties of neighbourhood substitution in constraint satisfaction problems. *Artificial Intelligence* 90:1–24.

Cooper, M.C. 1999. Linear-time algorithms for testing the realisability of line drawings of curved objects. *Artificial Intelligence*, 108:31–67.

Cooper, M.C. 2000. Linear constraints for the interpretation of line drawings of curved objects. *Artificial Intelligence*, 119:235–258.

Cooper M.C. Not yet published. Reduction operations in fuzzy and valued constraint satisfaction, *Fuzzy Sets and Systems*.

Faugeras, O. 1993. *Three-Dimensional Computer Vision*. MIT Press: Cambridge, MA, USA.

Freuder, E.C. 1991. Eliminating interchangeable values in constraint satisfaction problems. In *Proceedings AAAI-91*, Anaheim, CA, pp. 227–233.

Huffman, D.A. 1971. Impossible objects as nonsense sentences. In *Machine Intelligence*, Vol. 6, B. Meltzer and D. Mitchie (Eds.). Edinburgh University Press: UK, pp. 295–323.

Kirousis, L.M. and Papadimitriou, C.H. 1988. The complexity of recognizing polyhedral scenes. *Journal of Computer and System Sciences*, 37(1):14–38.

Lichtenstein, D. 1982. Planar formulae and their uses. *SIAM J. Computing*, 11:329–343.

Malik, J. 1987. Interpreting line drawings of curved objects. *International Journal of Computer Vision*, 1:73–103.

Malik, J. and Maydan, D. 1989. Recovering three-dimensional shape from a single image of curved objects. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(6):555–566.

Melhorn, K. 1974. *Graph Algorithms and NP-Completeness, EACTS*. Springer-Verlag: Berlin.

Nalwa, V.S. 1988. Line drawing interpretation: A mathematical framework. *International Journal of Computer Vision*, 2(2):103–124.

Parodi, P. and Torre, V. 1994. On the complexity of labeling perspective projections of polyhedral scenes. *Artificial Intelligence*, 70:239–276.

Rubin, J.M. and Richards, W.A. 1982. Color vision and image intensities: When are changes material. *Biological Cybernetics*, 45:215–226.

Shimshoni, I. and Ponce, J. 1997. Recovering the shape of polyhedra using line-drawing analysis and complex reflectance models. *Computer Vision and Image Understanding*, 65(2):296–310.

Straforini, M., Coelha, C., and Campani, M. 1993. Extraction of vanishing points from images of indoor and outdoor scenes. *Image and Vision Computing*, 11(2):91–99.

Sugihara, K. 1982. Mathematical structures of line drawings of polyhedrons—Towards man-machine communication by means of line drawings. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 4(5):458–469.

Sugihara, K. 1984. A necessary and sufficient condition for a picture to represent a polyhedral scene. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6(5):578–586.

Sugihara, K. 1986. *Machine Interpretation of Line Drawings*. MIT Press: Cambridge, MA.

Tai, A., Kittler, J., Petrou, M., and Windeatt, T. 1993. Vanishing point detection. *Image and Vision Computing*, 11(4):240–245.

Tsang, E. 1993. *Foundations of Constraint Satisfaction*. Academic Press: San Diego.

Waltz, D. 1975. Understanding line drawings of scenes with shadows. In *Psychology of Computer Vision*, P.H. Winston (Ed.). McGraw-Hill, New York, pp. 19–91.