International Journal of Computer Vision 73(2), 195–212, 2007

DOI: 10.1007/s11263-006-9783-7

Constraints Between Distant Lines in the Labelling of Line Drawings of Polyhedral Scenes

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

Received October 18, 2005; Revised May 23, 2006; Accepted July 12, 2006

First online version published in September, 2006

Abstract. The machine interpretation of line drawings has applications both in vision and geometric modelling. This paper extends the classic technique of assigning semantic labels to lines subject to junction constraints, by introducing new constraints (often between distant lines). These include generic constraints between lines lying on a path in the drawing as well as preference constraints between the labellings of pairs of junctions lying on parallel lines. Such constraints are essential to avoid an exponential number of legal labellings of drawings of objects with non-trihedral vertices.

The strength of these constraints is demonstrated by their ability to identify the unique correct labelling of many drawings of polyhedral objects with tetrahedral vertices. These new constraints also allowed us to deduce a general polyhedral junction constraint for the case when there is no limit on the number of faces which can meet at a junction.

Keywords: line drawing labelling, parallel lines, polyhedral objects, non-trihedral vertices, valued constraint satisfaction problem, soft constraints

1. Introduction

The interpretation of line drawings is a classic problem in machine vision. The pioneers in this field (Clowes, 1971; Huffman, 1971) concentrated on perfect projections of opaque polyhedral objects and on the problem of labelling each line as concave, convex or occluding. With the further restriction to simple trihedral vertices (i.e. vertices formed by the intersection of three faces), they were able to write down the complete list of legal junction labellings, known as the Huffman-Clowes catalogue. Although determining whether a given line drawing has a global labelling consistent with the Huffman-Clowes catalogue is known to be an NP-complete problem (Kirousis and Papadimitriou, 1988), the median-case complexity has been empirically observed to be $O(n)$ (Parodi et al., 1998) and polynomial-time algorithms exist for certain special cases (Kirousis, 1990; Parodi and Torre, 1994). Furthermore, the Huffman-Clowes catalogue is often sufficient to reduce the theoretically exponential number of labellings to a manageable number of legal labellings when it is used in conjunction with the Outer Boundary constraint, which says that when the drawing represents an isolated object or group of objects, the outer boundary can be assigned a unique labelling corresponding to an occluding edge.

This initial success was tempered by the following points:

1. The catalogue of labelled junctions provides necessary but not sufficient conditions for physical realisability of the drawing.

2. Classifying lines as projections of concave, convex or occluding edges only provides partial information about the corresponding 3D scene.
3. The restriction to perfect projection of polyhedral objects with trihedral vertices is too unrealistic for most vision applications.

Sugihara (1984, 1986) resolved point (1) above by giving necessary and sufficient conditions for the physical realisability of a legally-labelled line drawing of a polyhedral scene, by expressing the problem as a linear programming problem whose solution represents the equations of visible faces of the objects in the scene. This technique theoretically also resolves point (2), but there is often a great deal of ambiguity in the result which is not present in a human interpretation of the same drawing. For example, a drawing of a cube is immediately interpreted by a human being as a cube, even though it could theoretically be the 2D projection of any of a large class of parallelipipeds.

Much progress has been made in recent years concerning point (3). Some systems have been specifically designed to allow for freehand sketching errors (Grimstead and Martin, 1996; Lipson and Shpitalni, 1996; Varley and Martin, 2000). Catalogues of labelled junctions have also notably been given for curved objects (Cooper, 1993; Malik, 1987), for curved objects with edges and surfaces which can meet tangentially (Cooper, 1997), for polyhedral objects with tetrahedral vertices (Varley and Martin, 2001), and for drawings of scenes with lighting effects such as shadows (Waltz, 1975) and contrast failure (Cooper, 2001). When the restriction to polyhedral objects with trihedral vertices is relaxed, theoretical analysis (Cooper, 1997, 2001) has shown that a catalogue of labelled junctions is insufficient to disambiguate line drawings. Furthermore, experimental trials indicated that time complexity to find the best labelling grows with the number of valid labellings which is now an exponential function of the size of the drawing (Varley, 2004).

Although junctions can be said to provide most of the information in a technical drawing, other features provide valuable clues. For example, if we know the vanishing points of all lines in a drawing of a polyhedral object with trihedral vertices, then the labelling problem is no longer NP-complete, but can be solved in polynomial time (Parodi and Torre, 1994). Other sources of 3D information include collinearity of a line with a junction or another line (Cooper, 2000, 2001), parallel lines (Cooper, 1999) and straight lines (in the case of curved objects with some straight edges) (Cooper, 2000).

When multiple interpretations are possible, many heuristics have been used to find the most plausible interpretation, such as maximising commonly occuring 3D features such as right angles, vertical edges, symmetries and parallel planes, while minimising the number of distinct objects, angles and edge-lengths in the reconstructed 3D scene (Lipson and Shpitalni, 1996).

In this paper we express the line drawing labelling problem as a valued constraint satisfaction problem (Cooper and Schiex, 2004). This allows us to mix hard constraints (which must imperatively be satisfied) with soft constraints expressing preferences between different combinations of labels. This framework not only allows us to combine several junction catalogues based on different assumptions (such as trihedral/tetrahedral vertices, polyhedra/curved objects) but also allows us to express preferences (for example for right angles or parallel faces).

The main contribution of this paper is the introduction of novel constraints between unconnected lines or junctions, based on parallel lines, cycles of lines or collinearity. These non-local constraints permit the propagation of information between unconnected components of the drawing. Among other things, these constraints formalise and generalise constraints between holes or bosses and the boundary of the planar surface on which it lies (Varley and Martin, 2003) or between lines separating the same two regions (Kirousis, 1990; Varley et al., 2004). These new constraints considerably strengthen the trihedral catalogue of Huffman (1971) and Clowes (1971) which is often insufficient to uniquely identify the correct labelling of a drawing. Our constraints thus provide an alternative to Sugihara's necessary and sufficient conditions for realisability (Sugihara, 1984). Although incomplete, they have the advantage of being applicable on small subsets of the drawing (in the same way as junction constraints) whereas Sugihara's test must be applied to each global labelling, of which there may be an exponential number. We report results of experimental trials in which we evaluate the strength of our new constraints in disambiguating drawings of polyhedral objects with tetrahedral vertices.

2. Line Drawing Labelling as Optimisation

To obtain labelling constraints, such as the Huffman-Clowes catalogue, we must make assumptions on the 3D scenes which may be represented and the projection operation which produces the drawing. Common assumptions include planar faces, trihedral vertices, general viewpoint, general object positions, perfect projection. We can measure the plausibility of a labelling by determining how many of these assumptions must be relaxed (and how many times). For example, consider a drawing and two labellings L_1 and L_2, such that L_1 requires that two vertices are tetrahedral and L_2 requires that one pair of parallel lines in the drawing are not actually parallel in 3D. Should we prefer L_1 to L_2 or vice-versa, or should we accept both as plausible interpretations?

Different applications may give rise to different answers to this and similar questions. For example if the line drawing has been derived from a human-entered drawing then the user may have explicitly specified that certain lines are projections of parallel lines in 3D (but this is not assumed in this paper) or knowledge of the application area may imply that tetrahedral vertices are quite common. In all cases, recent work on reduction operations on valued constraint satisfaction problems may be applied to mitigate the theoretical intractability of the resulting optimisation problem (Cooper, 2003; Cooper and Schiex, 2004; Cooper, 2004). Valued constraints are local cost functions which allow us to express preferences between legal labellings and also to mark other labellings as illegal by assigning them an infinite cost (Schiex et al., 1995).

In later sections we study non-local soft constraints based on a natural tendancy to minimise the number of distinct surface orientations in the interpretation of man-made objects. But first we begin by studying non-local hard constraints. It is well known that junction constraints alone are not sufficient to ensure realisability (Sugihara, 1984) and several workers have stated constraints on the labelling of sets of lines not meeting at a junction (Huffman, 1971; Kirousis, 1990). In the next section, we formalise constraints on the labelling of a set of lines intersected by a path from line L_A to L_B, where either $L_A = L_B$ or L_A, L_B are parallel. We emphasize that these constraints assume planar surfaces. They may also be applied to drawings of curved objects provided we have sufficient evidence that particular regions are the projection of planar faces. For example, it may be reasonable to assume that a surface which has a polygonal boundary is planar.

3. Parallel Lines Constraint

In this section we assume that the line drawing is an orthographic projection of the edges of opaque polyhedral objects from a general viewpoint. The general viewpoint assumption says that a small change in the position of the viewpoint does not alter the configuration of the drawing (including junction types and presence of parallel or collinear lines). The orthographic projection and general viewpoint conditions imply that parallel lines in the drawing are projections of parallel lines in the 3D scene. If, as will usually be the case, we mark lines as parallel if the angle between them is less than some threshold ϵ (to take into account, for example, rounding errors or a projection which is only approximately orthographic), then the constraints presented in this section become soft rather than hard constraints. The coding of such soft constraints is discussed in detail in Section 9.

Figure 1 shows combinations of labels that are physically impossible under these assumptions. In all figures we use a generic label $\bigtriangledown$ to represent any of the three labels $+$ (convex), $-$ (concave), $\rightarrow$ (occluding with the nearer object below the line). Similarly $\bigtriangleup$

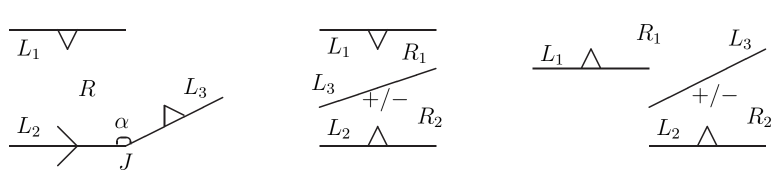

Figure 1. (a) The ParOcc constraint: This labelling is impossible if lines L_1, L_2, L_3 are adjacent to region R and lines L_1, L_2 are parallel; (b), (c) the ParCon constraint: this labelling is impossible if lines L_1, L_3 are adjacent to region R_1, lines L_2, L_3 are adjacent to region R_2, and lines L_1, L_2 are parallel but are not parallel to L_3.

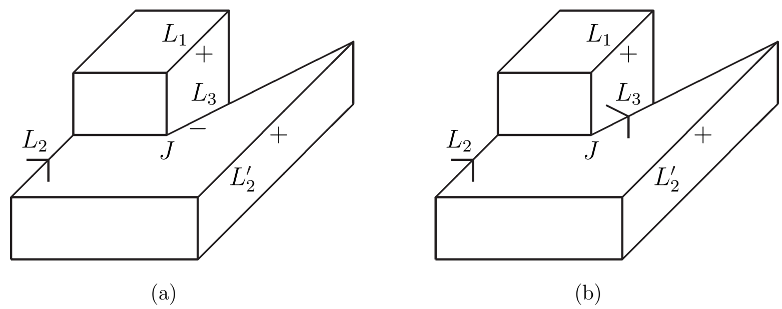

Figure 2. The trihedral labelling in (a) violates the ParCon constraint, whereas the tetrahedral labelling in (b) does not.

represents any of the labels $+, -, \leftarrow$. Thus, for example, the label $\triangledown$ for line L_1 in Fig. 1(a) implies that the edge projecting into L_1 lies on the surface projecting into region R. Note that in the ParOcc constraint (Fig. 1(a)), α can be any angle. A line L is said to be adjacent to a region R if L is part of the boundary of R when considered as a face in the planar graph representation of the drawing. For example, in the ParOcc constraint, L_1, L_2, L_3 are all adjacent to the region R as shown in Fig. 1(a). The junction J in Fig. 1(a) is any viewpoint-independent junction (such as a Y, W or L, but not a T-junction caused by a depth discontinuity). Let S_1, S_2, S_3 represent the 3D lines in the scene which project into L_1, L_2, L_3, respectively. The labelling in Fig. 1(a) implies that S_1 and S_2 are coplanar. Since S_2, S_3 intersect, they are also coplanar. Furthermore, since S_1, S_3 are parallel, they are also coplanar. It follows that S_1, S_2, S_3 all lie in the same plane, the face projecting into region R, but this contradicts the occluding label for L_2. A reflected version of the configuration in Fig. 1(a), with L_3 to the right of L_2, is also physically impossible.

In the ParCon constraint (Fig. 1(b)), the labellings shown are impossible provided L_1, L_2 are parallel but not parallel to L_3. The proof is omitted since we prove a more general result below. Note that the parallel lines L_1, L_2 do not necessarily face each other; for example, the configuration in Fig. 1(c) gives rise to the same constraint.

Figure 2 shows an example of the use of the ParCon constraint. Applying the trihedral catalogue (Clowes, 1971; Huffman, 1971) and the Outer Boundary constraint produces a single legal labelling, which includes the four labels given in Fig. 2(a). However, this labelling violates the ParCon constraint both on L_1, L_2, L_3 and on L_1, L'_2, L_3. Applying the tetrahedral catalogue (Varley and Martin, 2001) instead of the trihedral catalogue produces an alternative labelling which includes the four labels shown in Fig. 2(b). This labelling, in which junction J is the projection of a tetrahedral vertex, is consistent with the ParCon constraint.

We can generalise the ParOcc and ParCon constraints to a general constraint which can be applied to any pair of parallel lines. Consider a labelled drawing produced by orthographic projection. A *path* is a locus of points in the drawing from a point on a line L_1 to a point on a line L_2. It may intersect any number of intermediate lines between L_1 and L_2. A path Π is *parallel* if L_1, L_2 are parallel and the tangent to Π is parallel to L_1, L_2 at each intersection of Π with intermediate lines. In the following we assume (by rotation of the drawing if necessary) that L_1 and L_2 are horizontal. L_1 may be above or below L_2 in the drawing and Π may approach intermediate lines L from the left or the right (but always horizontally). Figure 3 shows two ways in which a path Π, shown as a broken line, may intersect an intermediate line. The direction of Π is indicated by an arrow.

An intersection between a parallel path Π and an intermediate line L is known as *strictly monotone* if L is labelled as in one of the two configurations given in Fig. 3. The justification for this term is as follows. Imagine an imaginary horizontal line segment H_P in the plane of the drawing such that H_P intersects Π at a

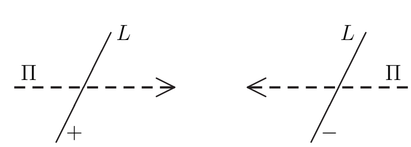

Figure 3. The catalogue of strictly monotone intersections.

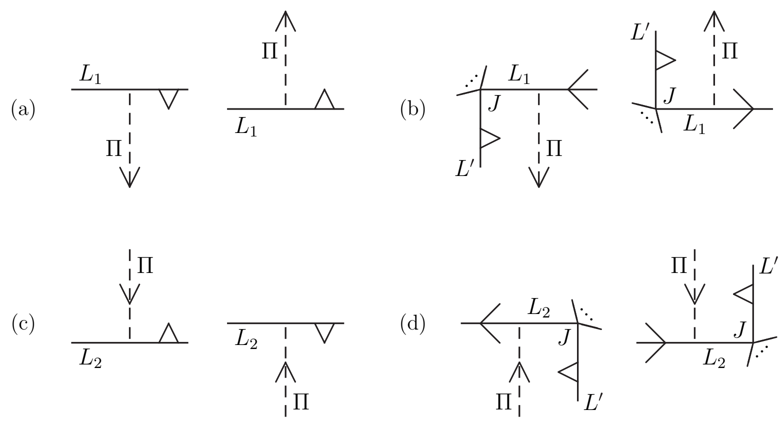

Figure 4. (a,b) How a strictly monotone path Π may begin; (c), (d) how a strictly monotone path Π may end.

point P. Let S_P be the 3D line segment which projects into H_P and which lies on the 3D planar surface visible at P. As P follows the path Π in the drawing and intersects an intermediate line L, S_P twists in 3D so that the depth z_r of its right hand end either increases or decreases. In both cases shown in Fig. 3, the depth z_r strictly increases. The angle of line L in Fig. 3 is arbitrary, but L must be at an angle greater than ϵ to the horizontal to avoid superstrictness problems (i.e. marking a legal labelling as illegal due to rounding errors in the positions of junctions). Note that the path Π is shown as a straight line but may be any continuous curve provided that the tangent to Π at its intersection with L is horizontal.

Figure 4 shows some ways in which a path may begin or end at the horizontal lines L_1, L_2. The generic label $\triangledown = \{+, -, \rightarrow\}$ is as in Fig. 1. J represents any viewpoint-independent junction (i.e. not a T-junction caused by a depth discontinuity). In Fig. 4(b), (d) any number of other lines can meet at J, as shown. The angle between L' and $L_1(L_2)$ is arbitrary. No other junction lies on $L_1(L_2)$ between J and the point P where Π intersects $L_1(L_2)$.

A parallel path is called *strictly monotone* if

1. it contains only strictly monotone intersections
2. it begins at a line L_1 in one of the configurations shown in Fig. 4(a) or (b)
3. it ends at a line L_2 in one of the configurations shown in Fig. 4(c) or (d)
4. *either* it contains at least one intersection *or* it begins with a configuration in Fig. 4(b) *or* it ends with a configuration in Fig. 4(d).

It is clear that if the two parallel lines L_1 and L_2 are projections of parallel lines in 3D, then there can be no strictly monotone path joining L_1 and L_2; otherwise the line projecting into L_2 would be twisted compared to the line projecting into L_1. The Parallel Lines constraint simply says that no strictly monotone parallel path exists in a labelled line drawing. It is easily seen that the ParOcc and ParCon constraints are just special cases of the Parallel Lines constraint.

Figure 5 shows an example of the use of the Parallel Lines constraint. Even assuming trihedral vertices and using the Outer Boundary constraint, this drawing still has four legal labellings. All but one of these labellings can be eliminated by applying the Parallel Lines constraint to the lines AB, EF with a path Π passing by the intermediate lines BC, ED. The labelling in Fig. 5(a) is an example of a labelling that violates this constraint. Figure 5(b) is the only legal labelling consistent with the Parallel Lines constraint.

A strictly monotone path may begin and end at the same line, since a line is necessarily parallel to itself. In this case, we do not need the assumption of an orthographic projection. A strictly monotone path may contain no intermediate lines provided it begins or ends in one of the configurations in Fig. 4 (b), (d). Figure 6 shows a constraint on the labelling of three consecutive lines on the boundary of a region R. Any number of lines can meet at junctions A and B, as shown, and the angles α, β are arbitrary. The labelling shown in Fig. 6 is illegal if junctions A, B are viewpoint-independent (i.e. not depth-discontinuity T-junctions), there are no other junctions between A and B, and the surface projecting into region R is planar. This constraint is easily deduced from the Parallel Lines constraint with a strictly monotone path leaving the line AB towards the right and immediately returning to AB.

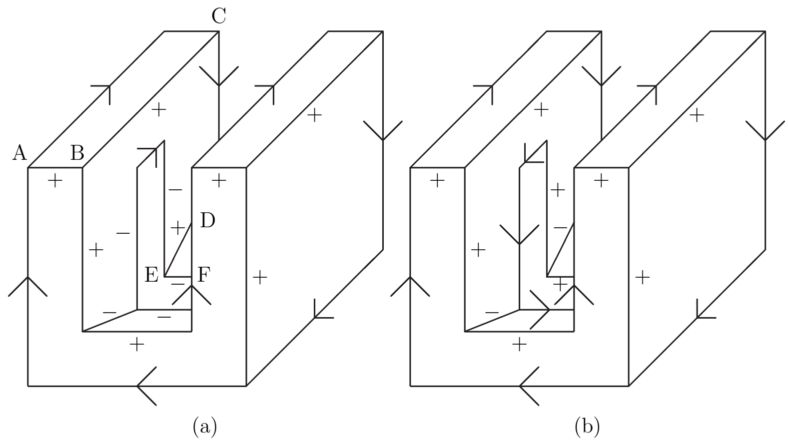

Figure 5. (a) A legal labelling according to the trihedral catalogue and the Outer Boundary constraint; (b) the unique and correct labelling found by also applying the Parallel Lines constraint and propagating.

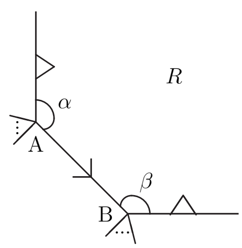

Figure 6. This labelling is impossible by the Parallel Lines constraint.

It is clear that when applying the Parallel Lines constraint it is not necessary to specify the actual locus of points traced by the path Π: we identify Π with the sequence of lines and regions it intersects. However, if the length t of a path is defined as the number of lines it contains (including the beginning and end lines), then the number of strictly monotone paths of length t is $\Theta(n^t)$ in the worst case where n is the number of lines in the drawing. Therefore, in practice, we recommend limiting the use of the Parallel Lines constraint to parallel paths with, for example, at most two intermediate lines. When a line drawing contains many parallel lines, we can often identify parallel planes. For example, a human being sees at a glance that the object depicted in Fig. 2 has three pairs of parallel planar surfaces. If we know (from their labelling) that lines L_1, L_2 lie on the planar surface projecting into region R, lines L'_1, L'_2 lie on the planar surface projecting into region R', L_1, L'_1 are parallel, L_2, L'_2 are parallel, but L_1, L_2 are not, then we can deduce that the surfaces projecting into R, R' are parallel planes. In this case, we can extend the Parallel Lines constraint as follows. A strictly monotone path can jump from any point in the interior of R to any point in the interior of R', without affecting the validity of the Parallel Lines constraint. We call the resulting constraint the Extended Parallel Lines constraint.

As a trial of the usefulness of the Parallel Lines constraint, we examined the labelled line drawings given in Varley and Martin's paper to illustrate their catalogue of trihedral and tetrahedral junction labellings (Figs. 4–9 and 24–166 of Varley and Martin (2001)). Eliminating mirror-image versions of the same drawing produced a total of 92 drawings. Nine of these drawings are such that if taken out of the context of the paper (where several drawings from different viewpoints are given of each different object) it would not necessarily be interpreted by a human being as indicated by Varley and Martin (2001). These drawings were ignored, which left us with 83 line drawings, containing an average of 14.8 lines each. Out of these 83 drawings, 54 of them were correctly labelled by applying the Outer Boundary constraint (which simply imposes an occluding label on the external boundary of the drawing) and prefering trihedral to tetrahedral labellings. For all the remaining 29 drawings, applying the Parallel Lines constraint or its extended version was sufficient to make the correct labelling the unique optimal labelling. In only three cases did we require the Extended Parallel Lines constraint. In all 29 cases, the number of intermediate lines in the Parallel Lines constraint was never greater than two. In the experimental trials described in this section, we applied the Parallel Lines constraint (coded as a hard constraint) to all parallel paths in the drawing involving up to two intermediate lines.

In a second trial, we studied the same sample of 83 drawings but this time without applying the Outer Boundary Constraint. If for each drawing, we imagine the same object this time depicted resting on a large flat surface, such as a table-top, the human interpretation of the drawing does not change, even though the Outer Boundary Constraint no longer applies. The resulting drawing is only ambiguous for a human being in that we cannot determine whether the object is actually touching the surface or not, and if so along which edges. Surprisingly, the Parallel Lines constraint, together with our preference for trihedral junction labellings, was sufficient to identify as optimal exactly the interpretations corresponding to human intuition in all but one of the 83 drawings. Note that the almost 100% success rate in imitating human intuition as to which is the most likely interpretation of these drawings is no doubt due to the fact that Varley and Martin used parallel lines as the main visual cue to avoid ambiguity in their drawings. Other sets of drawings may require reasoning about, for example, collinearity or symmetry.

In a previous paper we studied the interpretation of line drawings in which lines may be missing due to contrast failure (Cooper, 2001). If contrast failure is assumed to only occur between parallel surfaces, then the Parallel Lines constraint is still valid for strictly monotone paths Π beginning and ending at the configurations in Fig. 4(a) and (c). This follows from the fact that, even if the path Π intersects a missing occluding line, the 3D orientation of the occluding and occluded surfaces are identical since these surfaces are assumed to be parallel (Cooper, 2001).

In the following section we give a theoretical application, by showing how the Parallel Lines constraint can be used in the construction of a junction catalogue for polyhedral vertices.

4. A Universal Constraint for Simple Polyhedral Junctions

In this section we derive a constraint on the labelling of a large class of junctions when there is no limit on the number of surfaces which can meet at a vertex.

Definition 4.1. A *simple junction* J is any intersection of lines terminating at J such that no two of these lines are collinear.

Figure 7 shows a simple junction J. For any line L terminating at J, we can apply the Parallel Lines constraint to the path Π shown in Fig. 7, since L is parallel to itself. Note that, in this special case, we no longer require the assumption of orthographic projection to apply the Parallel Lines constraint. In this section, therefore, we study the labelling of projections of polyhedral vertices under only a general viewpoint assumption.

It is a direct consequence of the Parallel Lines constraint that the labellings shown in Fig. 8 are illegal. Note that there may be any number $m \geq 0$ of lines arriving at J from above and any number $n \geq 0$ of lines arriving at J from below. The special case when L is the only line terminating at J is clearly also illegal, although this is not a consequence of the Parallel Lines constraint.

As an example of the use of the Polyhedral Junction constraint, consider the two labellings of a Multi junction in Fig. 9. The Polyhedral Junction constraint tells us that the labelling of Fig. 9(a) is illegal, whereas the labelling of Fig. 9 is legal (as will be proved below). The difference between the two configurations is that deleting the concave line leaves a W junction in Fig. 9(a) and a Y junction in Fig. 9(b). A $(+, +, +)$ labelling is legal for Y junctions but not for W junctions. Thus Fig. 9 provides a refinement to the list of labellings for

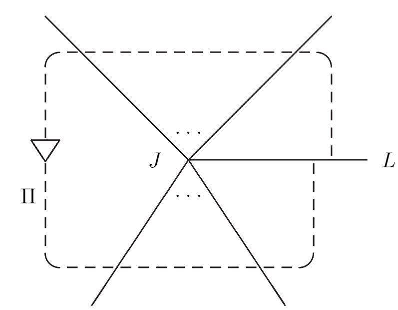

Figure 7. A simple junction J.

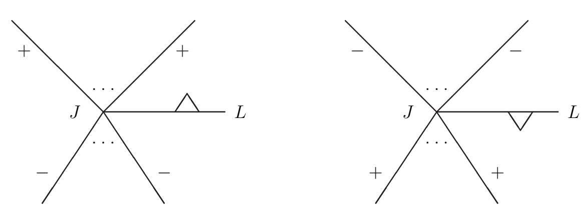

Figure 8. The Polyhedral Junction constraint: the labellings shown are illegal.

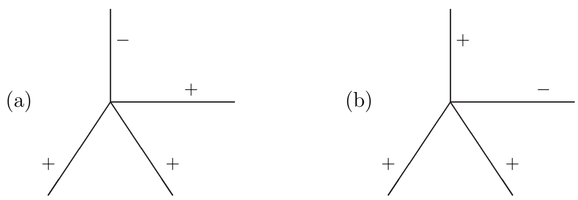

Figure 9. The labelling (a) is illegal, but the labelling (b) is legal.

Multi junctions in the tetrahedral catalogue (Varley and Martin, 2001).

In fact, the rest of this section is devoted to showing that the constraint of Fig. 8 is the tightest possible constraint on labellings of simple junctions.

Definition 4.2. A labelled junction J_1 is a *sublabelling* of a labelled junction J_2 if J_1 can be obtained from J_2 by deleting some number (possibly zero) of lines from J_2.

For example, a W-junction labelled $+ - +$ has 3 L-junction labellings, namely $+-$,$-+$, $++$, as sublabellings.

Figure 10 shows some legal polyhedral junction labellings of L, Y and W junctions.

Lemma 4.3. *All labellings of a simple junction J which are not illegal by the Polyhedral Junction constraint (Fig. 8) contain one of the labelled junctions in Fig. 10 as a sublabelling.*

Proof: Let J be a labelled junction which has none of the labelled junctions in Fig. 10 as sublabellings. We will show that J is illegal by the Polyhedral Junction constraint.

J cannot have two or more occluding labels, since this would imply that it necessarily had as a sublabelling an L-junction with two occluding labels and all such labellings of L-junctions are listed in Fig. 10. Suppose that J has just one line A with an occluding label,

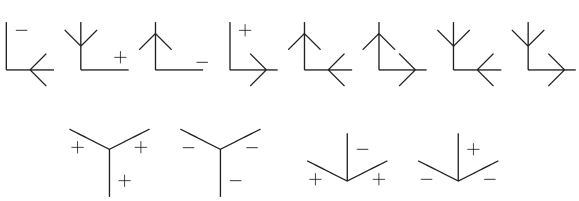

Figure 10. Minimal legal polyhedral junction labellings.

and suppose that A is labelled with an arrow pointing towards the junction (the proof for other case being entirely similar). Since J has neither of the first two L-junction labellings given in Fig. 10 as sublabellings, it must be of the form shown on the left of Fig. 8 and hence must be illegal.

If each line meeting at J is labelled $+$, then since J has no Y-junction, labelled $+ + +$, as a sublabelling, there must be an angle $\alpha > \pi$ between two adjacent lines. But then J is again illegal by the Polyhedral Junction constraint. A similar argument applies if all the lines meeting at J are labelled $-$.

The only case left to consider is when J has no occluding labels and at least one $+$ label and at least one $-$ label. Let A, B be two lines meeting at J such that A, B have different labels and such that the angle β between A and B is minimal. Without loss of generality, suppose A is labelled $+$ and B is labelled $-$. This situation is illustrated in Fig. 11. The extensions of lines A, B divide the plane into four quadrants as shown in Fig. 11. In Q_{opp} there cannot be two lines with different labels; if not this would contradict the minimality of β. In Q_A there can be no line labelled $-$, otherwise J would have the W-junction labelling $- + -$ as a sublabelling. Similarly, in Q_B there can be no line labelled $+$, otherwise J would have the W-junction labelling $+ - +$ as a sublabelling. The labelling of J is then illegal by the Polyhedral Junction constraint (with $L = A$ if Q_{opp} contains only $+$ labels and with $L = B$ if Q_{opp} contains only $-$ labels). □

Lemma 4.4. *Let J be a labelled simple junction. If J has a legal polyhedral junction labelling J_0 as a sublabelling, then J is also a legal polyhedral junction labelling.*

Proof: Let J_0 be a legal polyhedral labelling which is a sublabelling of J. Let V_0 be a vertex which projects

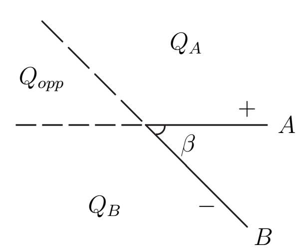

Figure 11. Q_{opp} contains either only $+$ lines or only $-$ lines, Q_A contains only $+$ lines and Q_B only $-$ lines.

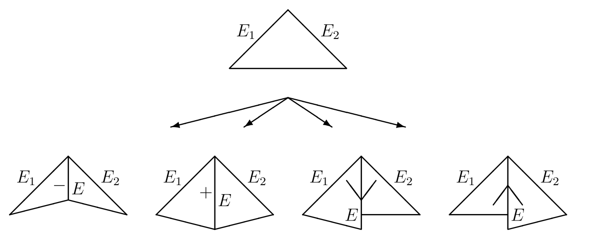

Figure 12. Introduction of a new edge E of any type within a surface bounded by non-collinear edges E_1, E_2.

into J_0. Suppose that L is a line of J which does not occur in J_0. L must lie in some region R between adjacent lines L_1, L_2 meeting at J_0. Let S be the surface which projects into R.

Consider first the case in which S is a one of the surfaces meeting at the vertex V_0 and that S is bounded by 3D edges E_1, E_2 projecting into lines L_1, L_2. Partition S into two distinct faces S_1, S_2 separated by a 3D line E which projects into line L. Since E_1, E_2 are not collinear, by our assumption that J contains no collinear lines, it is possible to rotate S_1 about E_1 and S_2 about E_2, by small angles ϵ_1, ϵ_2, so that E becomes either a concave or convex edge. By also creating a hidden edge, it is possible, in a similar manner, to construct an occluding edge. Figure 12 shows how it is possible to introduce a concave, convex or occluding edge in a vertical surface. For sufficiently small ϵ_1, ϵ_2, the labelling of the lines L_1, L_2 of J_0 remains unchanged. The cases in which one or both of E_1, E_2 occlude the surface S can be dealt with by entirely similar constructions.

By repeating this operation for each line L of J not in J_0, we can clearly construct a vertex V which projects into J. □

The following theorem now follows immediately from Lemma 4.3 and Lemma 4.4.

Theorem 4.5. *Let J be a labelled simple junction. J is a legal polyhedral junction labelling if and only if it satisfies the Polyhedral Junction constraint (Fig. 8).*

Huffman (1972) characterised all legal labellings of simple polyhedral junctions in terms of the possibility of finding in dual space a closed trace corresponding to the junction. Theorem 4.5 provides a more explicit characterisation.

The table in Fig. 13 gives the number of labellings which are realisable as projections of trihedral, tetrahedral or polyhedral vertices, for each junction type of arity up to four. The figures for trihedral vertices refer to the basic Huffman-Clowes catalogue (Clowes, 1971; Huffman, 1971). The figures for tetrahedral vertices refer to Varley and Martin's catalogue (Varley and Martin, 2001). The lists of legal polyhedral labellings for L, Y, W, Multi and Peak junctions are easily obtained from Theorem 4.5. The list of legal polyhedral T-junction labellings can be found in Huffman (1972) and the lists of legal polyhedral labellings for K, Y and X junctions were obtained by exhaustive search making ample use of the Parallel Lines constraint (which incidently was not quite sufficient on its own to eliminate all illegal labellings since two illegal labellings for the ψ-junction also satisfy the Parallel Lines constraint).

Junction type	Trihedral	Tetrahedral	Polyhedral	Combinatorial
L	6	8	8	16
Y	5	32	52	64
W	3	28	52	64
T	4	20	24	64
Multi	-	9	240	256
Peak	-	11	240	256
K	-	8	139	256
ψ	-	-	94	256
X	-	-	48	256

Figure 13. The number of labellings for junctions of arity up to four.

The figures in the rightmost column are the number of combinatorially possible labellings. It is clear that the general polyhedral junction constraints are far too weak to obtain a unique labelling of a line drawing. Consider a catalogue containing junction types belonging to a set C (e.g. $C = \{L, Y, W, T\}$ in the trihedral catalogue). Let N be the total number of line-ends in the catalogue (e.g. $N = 2 + 3 + 3 + 3$ for the trihedral catalogue). Assuming a uniform distribution of junction types, a drawing with n lines contains on average $\frac{2n}{N}$ junctions of each type in C. For each junction type $t \in C$, let p_t be the ratio of the number of legal labellings to the number of combinatorially possible labellings. Since each line can have one of four labels, the average number of global legal labellings of a drawing with n lines is thus

$$4^n \prod_{t \in C} p_t^{\frac{2n}{N}}$$

We can therefore calculate that the average number of global legal labellings of a drawing with n lines is $\Omega(1.49^n)$ for the tetrahedral catalogue and $\Omega(2.8^n)$

for the polyhedral catalogue. This can be compared with the $O(0.73^n)$ average number of legal labellings using the trihedral catalogue. These figures highlight the need for new constraints involving small subsets of lines to reduce the exponential number of global legal labellings. On the drawings tested, the Outer Boundary constraint and the Parallel Lines constraint turned out to be the most powerful such constraints. In the following sections we establish constraints on the labellings of sets of lines intersected by a another form of path, based on depths rather than orientations of surfaces.

5. Lines Sharing the Same Two Regions

In this section we do not need the assumption of orthographic projection, but we assume that the drawing is a projection of a polyhedral scene from a general viewpoint. In order to introduce a generic constraint on lines intersected by a path in the drawing, we first study the special case of a pair of lines adjacent to the same two regions. In a line drawing of a polyhedral scene, each line separates two distinct regions of the drawing. The 2Reg constraint applies when two non-collinear lines are adjacent to the same two regions. For example, in both drawings in Fig. 14, the region above the line AB is the same as the region below the line *CD* and the region below the line AB is the same as the region above the line *CD*. We represent this generic situation by the diagram in Fig. 15(a) where the straight lines L_1, L_2 labelled l_1, l_2 are lines present in the drawing and the other lines are construction lines which simply indicate which side of the lines share the same region in the drawing. This is, in fact, a very common situation since it occurs at every L-junction. There is a distinct and less common case illustrated in Fig. 15(b).

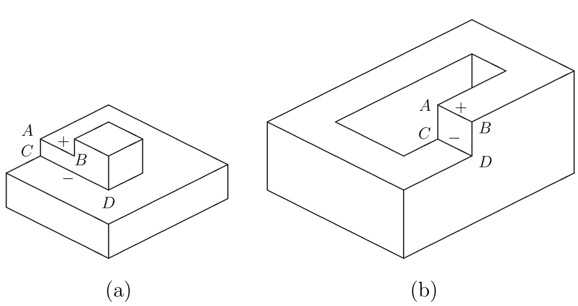

Figure 14. Two impossible pictures of polyhedral objects, both detected by the 2Reg constraint ((b) is known as Sugihara's box (Sugihara, 1986)).

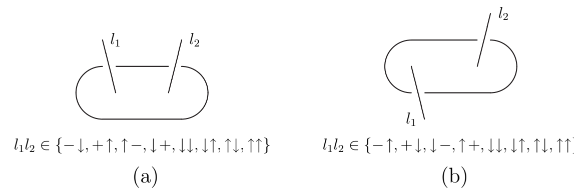

Figure 15. The 2Reg constraint.

Figure 15 gives the list of legal labellings of the two lines in both cases. Note that the angles between the two lines are arbitrary. For example, if the two lines are parallel this does not provide a stronger constraint. This 2Reg constraint follows from simple reasoning about the depth of scene points. As an example, consider the case in which $l_1 = +$. It is clear that $l_1 = +$ (convex) implies that a scene point just to the right of the line L_2 is nearer to the viewer than a point just to the left of L_2. This, in turn, implies that $l_2 = \uparrow$. An important point is that we do not need an assumption of trihedral vertices to obtain this constraint. This is, therefore, a very general constraint, from which can be deduced, for example, the L-junction constraint in the general polyhedral catalogue. Indeed the list of eight labellings in Fig. 15(a) coincides with the list of eight legal L-junction labellings given in Fig. 10.

An illustration of the strength of the 2Reg constraint is that it can be used to detect some classic examples of drawings of impossible polyhedral objects. For example, both of the drawings in Fig. 14 have legal labellings according to the trihedral catalogue, but none of these labellings is consistent with the 2Reg constraint. In all trihedral labellings, lines *AB* and *CD* have labels $+$ and $-$, respectively. But the $+-$ labelling is illegal according to the 2Reg constraint.

A special case of the 2Reg constraint that needs to be considered separately occurs when the two lines are collinear. In this case, the lists of legal labellings are given in Fig. 16. These lists are completely different to those for the 2Reg constraint (Fig. 15). This new constraint, which we call Col2Reg, was derived using the general viewpoint assumption which implies that collinear lines in the drawing are projections of collinear lines in the 3D scene.

When two lines share only one common region then without any other information we can deduce nothing about the possible labellings of the lines. However, in the case that the two lines are collinear, this gives rise to a constraint which we call the Col1Reg constraint. The lists of legal labellings are given in Fig. 17.

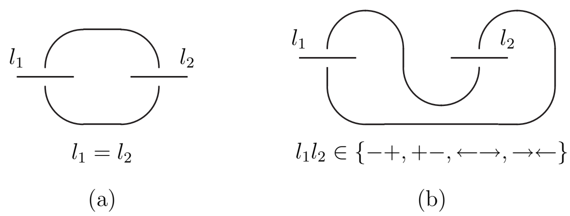

Figure 16. The Col2Reg constraint.

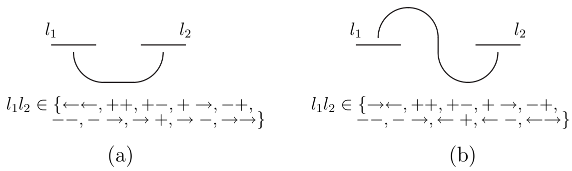

Figure 17. The Col1Reg constraint.

To illustrate the utility of the Col2Reg and Col1Reg constraints, consider the example line drawing in Fig. 18. After applying the trihedral catalogue and the Outer Boundary constraint, seven of the line labels are still ambiguous. However, by applying the Col2Reg constraint to the pair of collinear lines (AB, CD) and the Col1Reg constraint to the pairs of collinear lines (EF, GH) and (IJ, HK), a unique labelling is determined after propagation. For example, the fact that the line AB is labelled as convex (+) implies immediately by Col2Reg (case (b) of Fig. 16) that line CD is concave (−) rather than occluding (→). Similarly, the concave label (−) for line EF implies by Col1Reg (case (a) of Fig. 17) that line GH cannot be labelled as ←.

Kirousis (1990) gave a constraint for L-chains (sequences of lines connected by L-junctions) in drawings of objects with trihedral vertices. Stated succinctly, the L-chain constraint says that a + or − label for a line on an L-chain C uniquely determines the labelling of all lines on C and that an occluding label for a line on C reduces the number of possible labels for all other lines to at most three. This constraint is a direct consequence of the 2Reg and Col2Reg constraints. It therefore follows that the L-chain constraint also holds for drawings of objects with non-trihedral vertices.

6. Cyclic Path Constraint

In Fig. 15(a), (b) it was necessary to make explicit a cyclic path intersecting the two lines, since the constraint is not identical for the two distinct cases shown in Fig. 15. Furthermore, a completely different constraint is obtained when the lines are collinear (Fig. 16). The 2Reg and Col2Reg constraints can be generalised to the case of a path passing through $n > 2$ lines. Since, even for $n = 3$ we identified 95 distinct cases, we prefer to give a generic constraint valid for all cases and all values of n. The resulting Cyclic Path constraint, given below, was inspired by but also strictly generalises Huffman's cut-set rule based on reasoning in dual space (Huffman, 1972). Our rule allows for paths of arbitrary shape which do not necessarily begin and end at the same point and we correct Huffman's treatment of intersections with occluding lines. It can be considered as an application of Draper's sidedness reasoning (Draper, 1981) along a path in the drawing.

A *cyclic path* is a path in a drawing which begins at a point P_1 on line L_0 and ends at a point P_2 on L_0. It is *anchored* at a point Q collinear with L_0. Any two or three of the points P_1, P_2 and Q may, and often do, coincide. Suppose that a cyclic path passes through regions $R_1, \ldots, R_t$ which are projections of planar surfaces $S_1, \ldots, S_t$ and let $z_1, \ldots, z_t$ denote the 3D depths of these surfaces at a point which projects into Q in the drawing. Then a labelling is clearly illegal if it implies $z_1 \leq z_2 \leq \cdots \leq z_t \leq z_1$ with at least one of the inequalities being strict.

We define the intersection of a path Π with a line L, separating regions R_i and R_{i+1}, to be *strictly positive* if the labelling of L implies that $z_i < z_{i+1}$. This is the case for the intersections in Fig. 19(a) and (b), in which Q lies on the same side of L as R_i (respectively R_{i+1}) if L is labelled + (−). An intersection is said to be *null* if the labelling of L implies $z_i = z_{i+1}$. This is the case for the intersection shown in Fig. 19(c), where Q is collinear with L and L is labelled either convex or concave. An cyclic path Π is *strictly positive* if

1. Π contains only strictly positive or null intersections
2. Π begins as shown in Fig. 20(a)
3. Π ends as shown in Fig. 20(b–d)
4. *either* Π ends as in Fig. 20(b) and Q does not coincide with a junction *or* Π contains at least one strictly positive intersection.

In the configuration of Fig. 20(b), if the point Q does not coincide with a junction, then the occluding label implies a strict depth inequality between the two surfaces at Q ($z_1 < z_t$). In the configurations illustrated in Fig. 20(a) and (c) Q may lie either to the left or the right of P_1 or P_2. In Fig. 20(d), the path Π may arrive

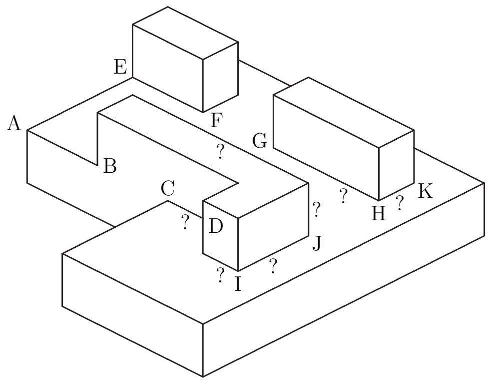

Figure 18. A line drawing which has seven ambiguous labels (marked with a question mark) after applying the trihedral catalogue and the Outer Boundary constraint, but which has a unique labelling after applying the Col1Reg and Col2Reg constraints.

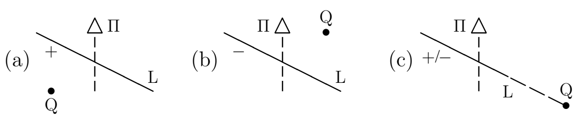

Figure 19. (a), (b) the two types of strictly positive intersections; (c) a null intersection.

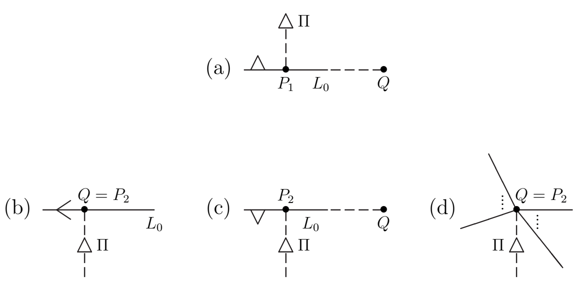

Figure 20. (a) How strictly positive cyclic paths begin; (b–d) how strictly positive cyclic paths end.

from any angle. Whatever the angle, the surface S_t projecting into the region R_t through which Π arrives at $P_2 = Q$ either intersects or passes behind the vertex projecting into Q, implying $z_1 \leq z_t$.

The Cyclic Path constraint simply says that a strictly positive cyclic path is illegal, since a net increase in depth as we traverse a cycle of surfaces is physically impossible. To avoid superstrictness problems, we could classify an intersection as strictly positive only if the perpendicular distance between the anchor point Q and the extension of L exceeds some minimum value δ. Furthermore, we could choose to only allow the cases in Fig. 19(c) or Fig. 20(a) and (c) in which the anchor point Q actually lies on the line L or L_0 (respectively), to avoid superstrictness problems due to accidental collinearity of Q with these lines. The 2Reg and Col1Reg constraints are just special cases obtained by studying cyclic paths containing either one or zero intermediate lines. The Col2Reg constraint can be obtained from the Cyclic Path constraint and the Parallel Lines constraint (applied to an imaginary line lying anywhere on a surface but not parallel to the lines in the drawing).

A path is not specified by an actual locus of points, but rather by the sequence of regions and lines it intersects. For any given cyclic path intersecting lines $L_0, \ldots, L_{t-1}$, we do not need to test the Cyclic Path constraint for every point Q collinear with L_0. Imagine the extension of L_0 divided into t segments by its intersection with the extensions of the lines $L_1, \ldots, L_{t-1}$. It suffices to test the constraint for $2t - 1$ points Q, one per segment together with each of the $t - 1$ intersection points. As with the Parallel Lines constraint, we suggest applying the Cyclic Path constraint only to cyclic paths involving only a small number of lines, since in the worst case there are $\Theta(n^t)$ cyclic paths of length t in a drawing containing n lines.

It is important to note that the Cyclic Path constraint does not require an orthographic projection, and is hence more generally applicable than the Parallel Lines constraint. Applied to our sample of drawings from Varley and Martin (2001), we found that the Cyclic Path constraint invalidated a subset of the labellings invalidated by the Parallel Lines constraint (where both constraints were applied to all paths involving up to two intermediate lines). As a concrete example, the Polyhedral Junction constraint can easily be deduced from the Cyclic Path constraint, instead of the Parallel Lines constraint. Nevertheless the Cyclic Path constraint and Parallel Lines constraint often complement each other. As an illustration of the strength of the Cyclic Path constraint, consider the drawing in Fig. 21. The labels shown are consistent with the trihedral catalogue, the Outer Boundary constraint and the Parallel Lines constraint. However, a cyclic path passing through regions R_1, R_2, R_0 and intersecting lines AB, CD and EF invalidates the labelling $(-, +, -)$ for (AB, CD, EF) if we choose the anchor point Q as shown. Similarly a cyclic path passing through regions R_3, R_4, R_0 and intersecting lines GH, IJ, KL invalidates the labelling $(\rightarrow, +, +)$ for (GH,IJ,KL) if the anchor point is chosen to be any point on GH. A more subtle example is the labelling

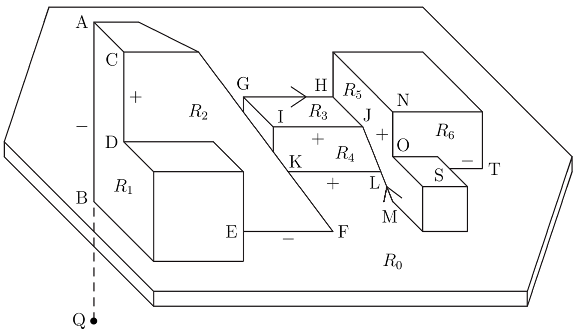

Figure 21. Three examples of applications of the Cyclic Path constraint.

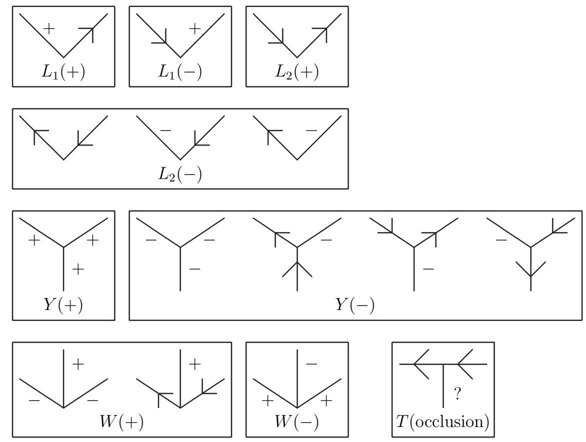

Figure 22. The catalogue of labelled junctions which are projections of trihedral vertices. A question mark represents any label.

(←, +, −) for (LM, NO, ST) invalidated by a cyclic path passing through regions R_5, R_6, R_0, intersecting lines LM, NO, ST and anchored at any point of LM.

7. Parallel Junctions on Distinct Faces

The constraints described in this section express a preference for a small number of distinct orientations of 3D edges in the scene reconstructed from the drawing. We do not assume planar faces and these preference constraints can thus be applied even when surfaces are curved. In this case, the straight lines depicted in the figures in this section represent tangents to the curved lines which meet at a junction. All the constraints on junction pairs given in this section are only valid if the total number of distinct orientations of 3D edges meeting at the corresponding pair of vertices is equal to three. Thus these constraints certainly hold if we can assume that all faces in the scene are parallel to one of only three planes. In general, this is too strong an assumption and hence these constraints simply express a preference for 3D interpretations involving pairs of trihedral vertices V_1, V_2 such that each of the planes meeting at V_1 is parallel to one of the planes meeting at V_2. Such interpretations are more likely in the case of man-made objects, which tend to have many parallel planes, by a simple application of Bayes's theorem.

Before giving our constraints on parallel junction pairs (pairs of junctions involving at least one pair of parallel lines), we reproduce in Fig. 22 the Huffman-Clowes (trihedral) catalogue of labelled junctions. We follow (Parodi and Torre, 1994) in dividing the set of labellings for Y and W junctions into subcategories Y(+), Y(−) and W(+), W(−). For example, a Y(+) junction is the projection of a convex vertex whereas a Y(−) is the projection of a concave vertex or a saddle point. Knowledge of the positions of the vanishing points of all lines (under perspective projection) is sufficient to classify Y and W junctions as + or − (Parodi and Torre, 1994). For notational convenience, we also divide the six labellings of L-junctions into four subcategories (called $L_1(+)$, $L_1(-)$, $L_2(+)$, $L_2(-)$) as shown in Fig. 22. Let V be a trihedral vertex projecting into an L-junction J. One of the edges meeting at V is not visible in the drawing. Let H denote the projection in the drawing of this hidden edge. When extended, the two visible lines meeting at J divide the plane of the drawing into four quadrants. For each of the four subcategories $L_1(+)$, $L_1(-)$, $L_2(+)$, $L_2(-)$, the hidden line H lies in a different quadrant. In the Huffman-Clowes catalogue, the set of legal labellings for an L-junction is simply the union of the sets of legal labellings for $L_1(+)$, $L_1(-)$, $L_2(+)$, $L_2(-)$ junctions.

Consider any pair of junctions J_1, J_2 in the drawing. Suppose that the two vertices V_1, V_2 projecting into these junctions are both trihedral and, furthermore, that each of the three edges meeting at V_1 is parallel to one of the edges meeting at V_2. Then the list of possible labellings of J_1, J_2 is given by the table of possible junction-types in Fig. 23 (if J_1, J_2 are Y or W junctions), Fig. 24 (if J_1 is a Y or W junction and J_2 is an L junction) or Fig. 25 (if J_1, J_2 are both L junctions). We call the corresponding constraint the Par3-3, Par3-2 or Par2-2 constraint, according to the number of lines meeting at junctions J_1 and J_2. The Par3-3 constraint has already been stated in a previous paper (Cooper, 1999), but is repeated here for completeness.

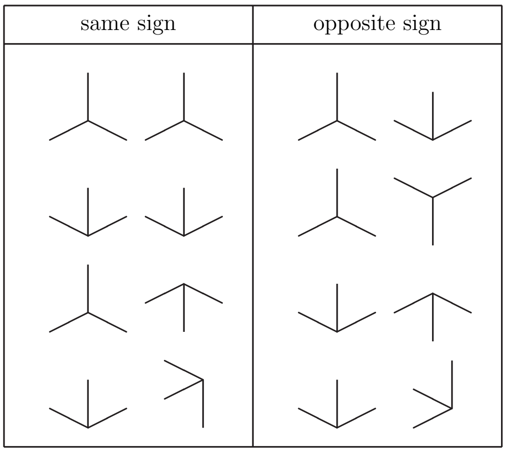

Figure 23. The Par3-3 constraint: pairs of junctions in the left hand column are of the same sign; pairs of junctions in the right hand column are of opposite sign

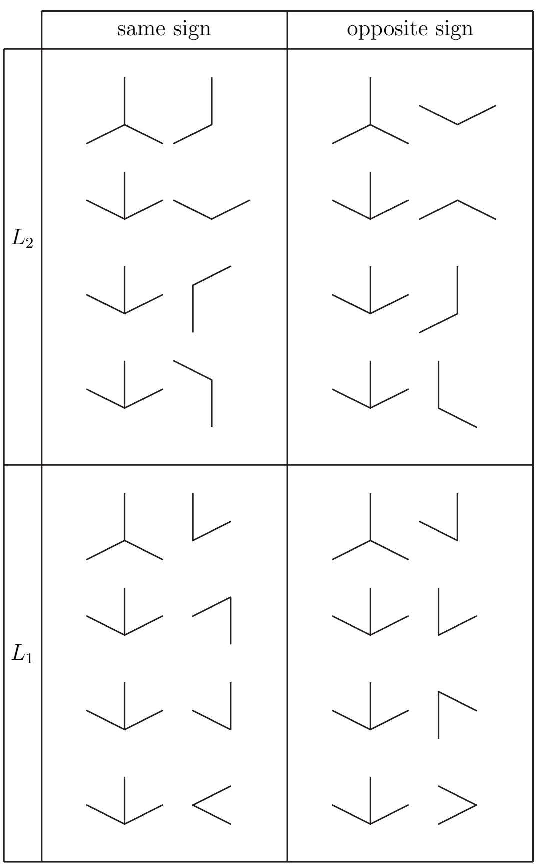

Figure 24. The Par3-2 constraint: pairs of junctions in the left hand column are of the same sign; pairs of junctions in the right hand column are of opposite sign; L-junctions in the top half of the figure are L_2 junctions; L-junctions in the bottom half of the figure are L_1 junctions.

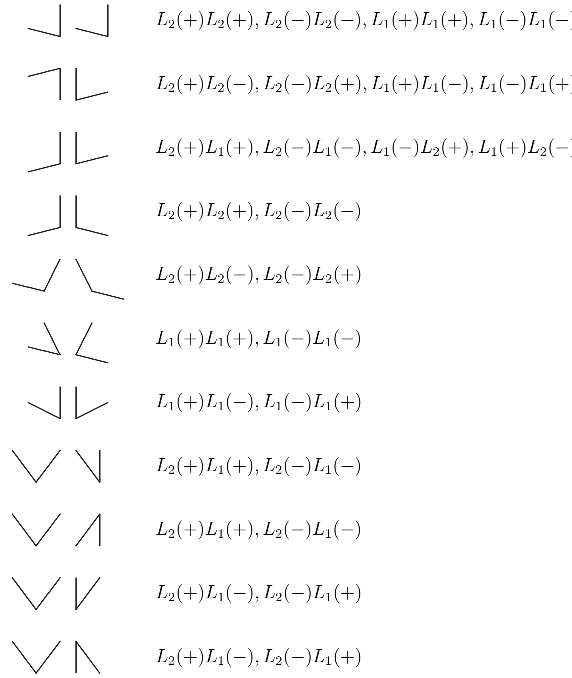

Figure 25. The Par2-2 constraint: for each pair of L-junctions shown on the left, the list of possible junction types are as shown on the right.

For example, assuming trihedral vertices and that parallel lines are projections of parallel 3D lines, the pair of Y-junctions at the top of the left hand column of Fig. 23 must be the same sign (i.e. Y(+)Y(+) or Y(−)Y(−)). As another example, the sixth configuration of pairs of L-junctions given in Fig. 25 has only two legal labellings which involve only three distinct face orientations. The constraints in Figs. 23–25 were derived by exhaustive search over all possible types of vertex-pairs formed by parallel planes, making use of the catalogue of labelled trihedral junctions in wireframe projections (Cooper, 2005).

Consider the drawing in Fig. 26(a). The trihedral catalogue and the Outer Boundary constraint imply the unique labelling of some lines, but the triangle *ABC* remains very ambiguous. Indeed, the triangle *ABC* has eight legal labellings. Note that ParOcc eliminates three of these eight labellings, but this still leaves five physically possible labellings. Par3-2 applied to the pair of junctions (D, A) indicates that A is an $L_1(+)$ junction which implies a unique labelling for the whole drawing, corresponding to the most natural interpretation. Note that applying Par3-2 to any of the junction pairs (B, E), (C, F), (C, I) or Par2-2 to either of the junction pairs (G, A), (B, H) also

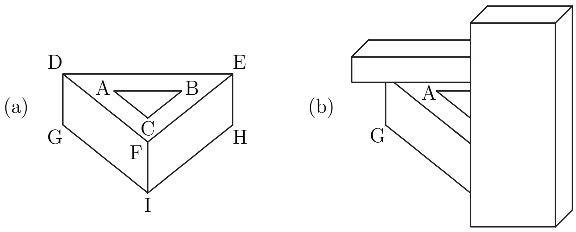

Figure 26. (a) A line drawing which has eight legal labellings according to the trihedral catalogue and the Outer Boundary constraint, but has only one labelling satisfying the Par3-2 constraint; (b) a line drawing which can be uniquely labelled using the Par2-2 constraint.

implies the same unique and correct interpretation of the drawing. Therefore, Par3-2 and Par2-2 would allow us to interpret correctly the triangle ABC even if all but one of the junctions D, E, F, G, H, I were occluded in the drawing. Figure 26(b) shows such an example.

The encoding of the Par3-3, Par3-2 and Par2-2 constraints as valued constraints is discussed in detail in Section 9.

8. Parallel or Collinear Lines Sharing a Region

Among the legal labellings for pairs of collinear lines given by the Col2Reg constraint, some require more distinct 3D face orientations than others and, hence, may be considered less likely. In the list of legal labellings given in Fig. 16(b), the labellings $+-$, $+-$ are more likely than the labellings $\leftarrow\rightarrow$, $\rightarrow\leftarrow$. Note that all legal labellings given in Fig. 16(a) are equally likely.

Among the legal labellings given in Fig. 17(a) as part of the Col1Reg constraint, the labellings $+-$, $-+$, $+\rightarrow$, $\rightarrow +$ all require more 3D face orientations than the other legal labellings, and are hence less likely. By the same reasoning, among the labellings given in Fig. 17(b), the labellings $++$, $--$, $\leftarrow\rightarrow$, $\leftarrow -$, $-\rightarrow$ are less likely than the others.

All legal labellings given in the 2Reg constraint are equally likely, since they all require three distinct face orientations, except in the special case when the lines labelled l_1, l_2 in Fig. 15 are parallel. In this case, the labellings $+\uparrow$, $\downarrow +$ are more likely than the other labellings given in Fig. 15(a) since they only require two distinct face orientations. Similarly, the labellings $-\uparrow$, $\downarrow -$, $\uparrow\uparrow$, $\downarrow\downarrow$, $\downarrow\uparrow$ are more likely than the other labellings given in Fig. 15(b).

A soft constraint also exists when two parallel lines share a single common region. This situation does not give rise to a hard constraint, since all 16 combinations of labels are theoretically possible. However, certain labellings require three distinct 3D face orientations and hence are less likely than those requiring only two. These unlikely labellings are given in Fig. 27 for the three distinct configurations of two parallel lines adjacent to a common region. Note that in Fig. 27(c), the line labelled l_1 may lie to the left or to the right of the line labelled l_2.

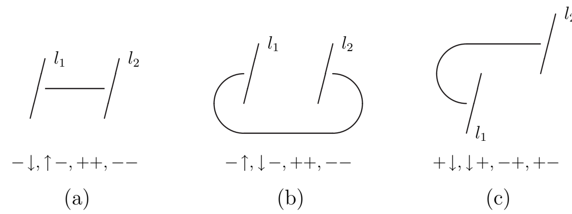

Figure 27. The Par1Reg soft constraint: the labellings shown are *unlikely* labellings $l_1 l_2$.

9. Encoding of Soft Constraints

In a Valued Constraint Satisfaction Problem (VCSP), a valued constraint is represented by a local cost function and the aim is to minimise the sum of these cost functions (Cooper and Schiex, 2004; Schiex et al., 1995). State-of-the-art VCSP solvers (http://carlit.toulouse.inra.fr/cgi-bin/awki.cgi/ToolBarIntro) maintain a form of soft arc consistency (such as FDAC) during branch and bound search and use appropriate variable/value ordering heuristics (Larrosa and Schiex, 2003). FDAC propagates hard constraints (represented by infinite costs) in all directions and propagates finite costs towards earlier variables in the instantiation order so as to produce a better lower bound with which to prune the search tree. This paper is concerned uniquely with new hard and soft constraints and their encoding as a VCSP; algorithmic aspects of VCSPs are treated in independent papers (Cooper, 2003; Cooper and Schiex, 2004).

For a given optimisation problem, many different encodings are possible as a VCSP. We say that the encoding of a line drawing labelling problem is *faithful* if the set of optimal solutions to the VCSP coincide with the most likely interpretations. A simple, but naive, encoding associates a fixed cost to the violation of each soft constraint. Consider a pair of parallel lines, involved in a large number m of Parallel Lines or Parallel Junctions (Par3-3, Par3-2 or Par2-2) constraints. It is possible, by a single violation of the general viewpoint assumption, that these two lines are in fact projections of 3D edges with two distinct 3D orientations, meaning that these m constraints are invalid. The naive encoding is therefore not faithful, since the likelihood that the general viewpoint assumption is violated is independent of m.

A more faithful encoding is obtained by the introduction of auxiliary variables, which of course has the drawback of increasing the size of the search space. For example, we can introduce a variable par_S for each set S of parallel lines in the drawing (par_S being true iff all lines in S are projections of parallel 3D edges). Then each Parallel Lines constraint along a parallel path Π beginning and ending at lines in S can be encoded as a hard conditional constraint of the form

$$par_S \Rightarrow \text{(the Parallel Lines constraint is satisfied by the lines on } \Pi\text{)}$$

An assignment $par_S = \texttt{false}$ incurs a fixed finite penalty corresponding to the violation of the general viewpoint assumption.

In the case of imperfect line drawings, such as those derived from freehand sketches, we can obtain an optimal partitioning of 2D lines into sets $S_1, \ldots, S_r$ of near-parallel lines, by using, for example, a linear-time optimal segmentation algorithm applied to the sorted array of their angles (Cooper, 1998). An auxiliary boolean variable par_{S_i} would be required for each S_i containing at least two distinct lines involved in at least one Parallel Lines or Parallel Junctions constraint.

A Par3-3 constraint between J, J' can be simply encoded as a hard conditional constraint of the form

$$par_{S_p} \wedge par_{S_q} \wedge par_{S_r} \Rightarrow \text{(the Par3-3 constraint is satisfied on } J, J'\text{)}$$

where the lines meeting at the junctions J, J' belong to the sets of parallel lines S_p, S_q, S_r. In the case of Par3-2 constraints, the corresponding conditional constraint is a soft constraint, since even if the two visible lines are projections of parallel 3D edges, we are merely expressing a preference for vertices in which the hidden third edge is parallel to the visible third edge. In other words, the constraint

$$par_{S_p} \wedge par_{S_q} \wedge par_{S_r} \Rightarrow \text{(the Par3-2 constraint is satisfied on } J, L\text{)}$$

can be violated with a fixed finite cost chosen as a function of the likelihood that the hidden third edge is parallel to the visible third edge. (As above, the lines meeting at J belong to the sets S_p, S_q, S_r.) The encoding of the Par2-2 constraint is entirely similar.

Alternative, more faithful, encodings exist at the cost of the introduction of more auxiliary variables. For example, we could introduce an auxiliary variable V_L with domain $\{0, \ldots, r\}$, for each L junction L, where $V_L = i$ if the hidden third line is parallel to the lines in S_i and $V_L = 0$ if this hidden line is parallel to no other visible lines in the drawing. In other words, we attempt to explicitly reconstruct the directions of hidden lines.

10. Discussion

In this paper we have presented constraints for the labelling of line drawings of polyhedral scenes. The resulting constraints are necessary but not sufficient conditions for physical realisability. The aim of this research is to identify low-arity hard or soft constraints which can be applied before or during the search for an optimal interpretation. We consider such constraints to be essential for any practical line drawing interpretation system which is not restricted to objects with trihedral vertices. Indeed, based on junction constraints alone, we have seen that the number of legal labellings would be an exponential function of the size of the drawing.

We have chosen to restrict our attention to constraints on lines intersected by a path as well as constraints derived from the presence of parallel or collinear lines. Many other constraints exist. Examples include constraints derived from the presence of lines collinear with junctions (Cooper, 2001), collinear lines (Lipson and Shpitalni, 1996), cubic corners (Conesa Pastor

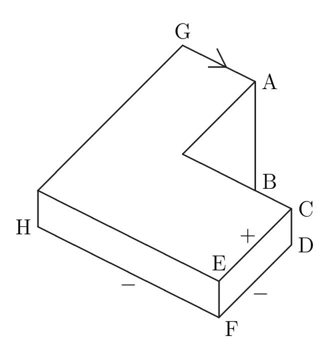

Figure 28. An example of a physically unrealisable labelled drawing not detected by our constraints.

et al., 1999; Varley and Martin, 2002), skew symmetry (Piquer Vicent et al., 2003) or vanishing points (Parodi and Torre, 1994). Furthermore, we have only studied low-arity constraints, and thus our constraints are necessarily incomplete even in the treatment of parallel and collinear lines. For example, the labelled drawing in Fig. 28 is physically unrealisable although it satisfies all the constraints given in this paper. Physically unrealisability is a consequence of the following argument. From the labelling and the presence of parallel lines, we can deduce that GACE and DFH lie on two parallel planes. But then, since AB is parallel to CD, it follows that B should not be visible since it must lie behind the plane DFH. We can thus deduce an arity-4 constraint based on the presence of three pairs of parallel lines. It is, of course, debatable whether constraints based on the presence of several pairs of parallel lines will be sufficiently robust to be worth applying in practical applications.

Acknowledgments

This paper was considerably improved thanks to discussions with Ralph Martin and Peter Varley concerning earlier drafts of this paper.

References

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

Conesa Pastor, J., Company Calleja, P., Gomis Marti, J.M. 1999. Initial modelling strategies for geometrical reconstruction - optimisation-based approaches. In *Proc. 11th Int. Conf. on Design Tools and Methods in Industrial Engineering*, pp. 161–171.

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

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

Cooper, M.C. 1998. The tractability of segmentation and scene analysis. *Int. J. Computer Vision*, 30(1):27–42.

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. 2001. The interpretation of line drawings with contrast failure and shadows. *International J. Computer Vision*, 43:75–97.

Cooper, M.C. 2003. Reduction operations in fuzzy or valued constraint satisfaction. *Fuzzy Sets and Systems*, 134:311–342.

Cooper, M.C. 2005. Wireframe projections: physical realisability of curved objects and unambiguous reconstruction of simple polyhedra. *International Journal of Computer Vision*, 64(1):69–88.

Cooper, M.C. and Schiex, T. 2004. Arc consistency for soft constraints. *Artificial Intelligence*, 154(1–2):199–227.

Cooper, M.C. 2004. Cyclic consistency: A local reduction operation for binary valued constraints. *Artificial Intelligence*, 155(1–2):69–92.

Draper, S.W. 1981. The use of gradient and dual space in line-drawing interpretation. *Artificial Intelligence*, 17:461–508.

Grimstead, I.J. and Martin, R.R. 1996. Incremental line labelling for sketch input of solid models. *Computer Graphics Forum*, 15(2):155–166.

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

Huffman, D.A. 1977. Realizable configurations of lines in pictures of polyhedra. In: E.W. Elcock and D. Michie (Eds.), *Machine Intelligence 8*, Ellis-Horwood, Chichester, UK, pp. 493–509.

Kirousis, L.M. 1990. Effectively labelling planar projections of polyhedra. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(2):123–130.

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

Larrosa, J. and Schiex, T. 2003. In the quest of the best form of local consistency for weighted CSP. In *Proc. IJCAI*, pp. 239–244.

Lipson, H. and Shpitalni, M. 1996. Optimisation-based reconstruction of a 3D object from a single freehand line drawing. *Computer-Aided Design*, 28(8):651–663.

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

Parodi, P., Lancewicki, R., Vijh, A., and Tsotos, J.K. 1998. Empirically-derived estimates of the complexity of labelling line drawings of polyhedral scenes. *Artificial Intelligence*, 105: 47–75.

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

Piquer Vicent, A., Company Calleja, P., and Martin, R.R. 2003. Skewed mirror symmetry in the 3D reconstruction of polyhedral models. *Journal of WSCG*, 11(1).

Schiex, T., Fargier, H., and Verfaillie, G. 1995. Valued constraint satisfaction problems: hard and easy problems. In *Proc. of the 14th IJCAI*, Montreal, Canada, pp. 631–637.

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

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

Varley, P.A.C. 2004. email message.

Varley, P.A.C. and Martin, R.R. 2000. A system for constructing boundary representation solid models from a two-dimensional sketch. In *Proc. Geometric Modelling and Processing 2000*, R. Martin and W. Wang (Eds.). IEEE Computer Society Press, pp. 13–30.

Varley, P.A.C. and Martin, R.R. 2001. The junction catalogue for labelling line drawings of polyhedra with tetrahedral vertices. *International Journal of Shape Modeling*, 7(1):23–44.

Varley, P.A.C. and Martin, R.R. 2002. Estimating depth from line drawings. In *Proc. 7th ACM Symposium on Solid Modelling and Applications*, pp. 180–191.

Varley, P.A.C. and Martin, R.R. 2003. Deterministic and probabilistic approaches to labelling line drawings of engineering objects. *International Journal of Shape Modeling*, 9(1):79–99.

Varley, P.A.C., Suzuki, H., and Martin, R.R. 2004. Interpreting line drawings of objects with K-vertices. In *Proc. Geometric Modelling and Processing 04*, S.-M. Hu and H. Pottmann (Eds.), pp. 249–358.

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