ELSEVIER

Image and Vision Computing 15 (1997) 263–276

Interpreting line drawings of curved objects with tangential edges and surfaces

Martin C. Cooper

IRIT, University of Toulouse III, 31062 Toulouse, France

Received 6 February 1995; accepted 1 October 1996

Abstract

This paper extends previous work on the interpretation of line drawings of curved objects to include objects with tangential edges and surfaces. We present a catalogue of labelled junctions and show that the existence of a legal global labelling can be determined in linear time. However, minimising the number of invisible vertices necessary to produce a legal global labelling is shown to be NP-hard. The potential combinatorial ambiguity inherent in the drawings can be mitigated by using information from the intensity image to identify points at which the object surface is tangential to the viewing direction.

Keywords: Line drawing labelling; Smooth edges; Constraint satisfaction

1. Introduction

The semantic labelling of line drawings was one of the first problems studied in Artificial Intelligence [1]. The results obtained showed great potential, since apparently intelligent behaviour was demonstrated by systems using a relatively small amount of information, a catalogue of legal junction labellings [2,3]. This catalogue was derived by considering all physically-possible three-dimensional vertices viewed from all possible viewpoints.

Unfortunately, the earliest work in this domain was limited to perfect line drawings of polyhedral objects, thus preventing its application to line drawings derived from real images of complex objects. Indeed, many workers made the obvious deduction that these early systems worked well *because* of the unrealistic assumptions that they made about objects and line-drawing formation. However, other workers showed that certain of the restrictive assumptions about line drawings could be relaxed without greatly reducing the performance of the labelling algorithms. In each case, other sources of information were found to reduce the ambiguity introduced by relaxing the assumptions:

1. Waltz [4] increased the number of line labels to allow line drawings with shadows. In fact, as artists know, shadows can be a source of information about the three-dimensional shape of objects. He also considered accidental alignment and cracks.
2. Falk [5] compared the drawing with known object models to compensate the loss of information when interpreting an imperfect line drawing (in which, for example, certain lines are missing).
3. In line drawings of curved objects, the line of sight may be tangential to the object surface. The resulting depth discontinuity is known as an extremal edge and its projection in the drawing is known as an extremal line. For example, the straight lines in a drawing of a cylinder are extremal lines. (See the appendix for a glossary of technical terms used in this paper.) Shapira and Freeman [6,7] introduced a new label for an extremal line and used multiple drawings of the same object to provide extra information. Malik [8] observed that, by detecting discontinuities of curvature of lines in the drawing, transitions from extremal edges to surface-normal discontinuity edges can be located.
4. In an earlier paper [9], the author showed that smooth edges (discontinuities of surface curvature) can be detected and labelled, thus generalising Malik's work which disallowed smooth edges.

Clearly, junctions are not the only source of information in a drawing. Pairs of similar curves, pairs of skew-symmetric curves, colinear points, sets of lines with a common vanishing point [10] and (probabilistic) information about the shape of objects are all potential sources of information. For example, under orthographic

0262-8856/97/$17.00
PII S0262-8856(96)01135-3

projection parallel lines in the drawing are likely to be projections of parallel lines in 3D. Recognition by Components [11] is based on five shape-inferring structures: colinearity, curvilinearity, skew symmetry, parallel curves and vertices.

This paper extends the class of curved objects with smooth edges, to include tangential edges and surfaces. New junction-types and new labels are introduced, but the main new source of information is the intensity image. Extremal edges and certain phantom vertices can be detected and identified in the intensity image.

We retain the following assumptions about the line drawings analysed:

1. Objects are opaque, have Lambertian surfaces and no surface markings.
2. Object surfaces are C^3 patches separated by surface-normal discontinuity edges and curvature discontinuity edges (smooth edges). In any neighbourhood centred on a vertex or edge, the object is topologically equivalent to its interior. Thus no two parts of the object are joined at a point or along a line; filaments and sheets are disallowed.
3. General viewpoint, general positions of the light sources and general positions of objects (i.e. an infinitesimal perturbation in the position of one of the viewpoint, a light source or an object does not change the configuration of the line drawing).
4. The drawing is a perfect projection of the 3D scene. There are thus no missing lines due to contrast failure and no shadows.
5. Objects do not contain improbable vertices.

We will justify the classification of each vertex as improbable or not as we construct our catalogue of labelled junctions. As a simple rule, a vertex is improbable if it requires an unnecessary coincidence. For example, polyhedral vertices are points of intersection of three or more planes, but since more than three planes meeting at a point represents a coincidence, we consider only trihedral vertices.

2. Basic catalogue for curved objects

In this section we give the basic catalogue of labelled junctions for curved objects composed of piecewise C^3 surface patches (Fig. 1). This is the catalogue given by Cooper [9] (a revised version of Malik's [8] catalogue) to which we have added the four T-junction labellings derived from the vertex N of Fig. 2. Many authors assume that this type of T-junction (which we call a non-occlusion T-junction) does not occur in the line drawing. Under this assumption all T-junctions in the drawing are due to occlusion and the drawing can immediately be segmented into separate objects. Unfortunately, non-occlusion T-junctions can occur in practice, and so we have to include these four T-junction labellings.

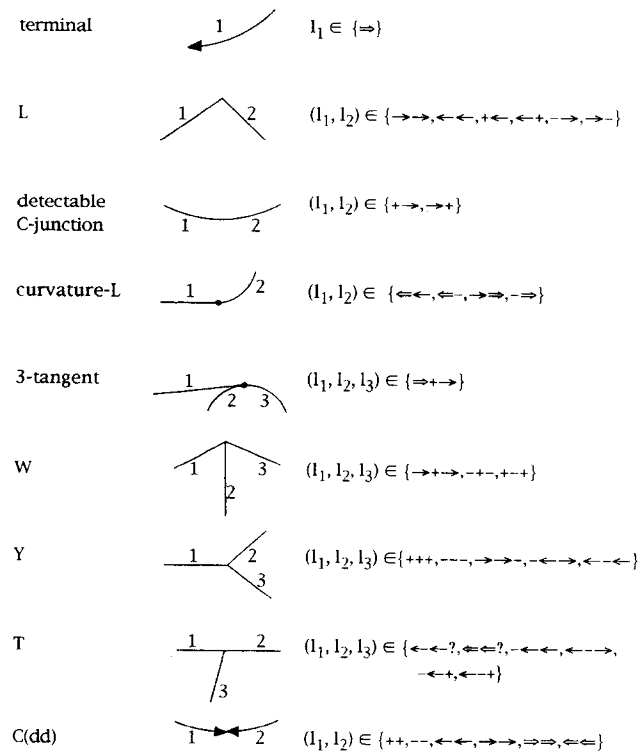

Fig. 1. Catalogue of labelled junctions for curved objects composed of C^3 surface patches.

We have also added a new junction, denoted $C(dd)$, which represents a C-junction at which both line-ends disappear (i.e. fade out) as they approach the junction. We allow this possibility because the brightnesses of the two curved surfaces which project on either side of a line in the drawing can vary continuously and will often be identical at isolated points.

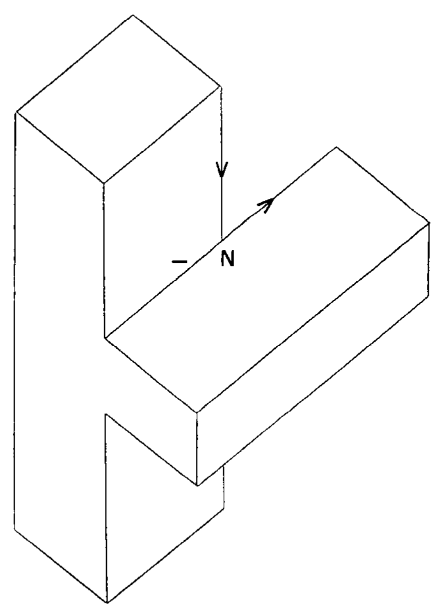

Fig. 2. Example of a vertex which projects into a non-occlusion T-junction in the drawing.

For typographical reasons an extremal line is represented here by $\Leftarrow$. In drawings extremal lines are labelled by double-headed arrows. The labelling of terminal junctions is discussed by Koenderink and van Doorn [12]. A dot represents a discontinuity of curvature of a line. However, in the case of a 3-tangent junction, lines 2 and 3 have continuous curvature. We explain in a later section why the C-junctions given in this catalogue are known as detectable C-junctions. The possibility of convex detectable C-junctions was demonstrated in a previous paper [9]. In the list of labellings for a T-junction, a question mark means any label.

3. Tangential edges and surfaces

The basic catalogue of Fig. 1 was derived under the assumption that there were no tangential surfaces, no tangential edges and no edge tangential to a surface. In this section we consider all these cases, while still imposing the general rule of no unnecessary coincidences. For example, we keep the restriction that vertices are formed by the intersection of a maximum of three surfaces: S_1, S_2, S_3. The consequences of relaxing this trihedral condition are discussed in a later section. Let E_{ij} represent the intersection of the surfaces S_i and S_j; if S_i and S_j are not tangential, E_{ij} is a simple 3D curve.

There are three generic cases to consider:

- S_1 is tangential to S_2
- E_{12} is tangential to S_3
- E_{12} is tangential to E_{13}

However, if no two surfaces are tangential, then the two latter cases are equivalent. In fact, in this case,

E_{ij} is tangential to S_k for all $i,j,k = 1,2,3$ and
E_{ij} is tangential to E_{ik} for all $i,j,k = 1,2,3$

as is illustrated in Fig. 3(a).

There are three cases to consider when two surfaces are tangential, according to the dimensionality of their intersection:

- the surfaces kiss at a single point,
- the surfaces touch along a simple curve, or
- the surfaces merge to form a single surface.

The case in which they kiss at a single point P is illustrated in Fig. 3(b). When they touch or merge along a simple curve, we call this a TIC (Tangential Intersection Curve). Three events can occur along a TIC to form a vertex:

1. A TIC can intersect another surface S, as illustrated in Fig. 3(c). The possibility that the TIC is tangential to S is disallowed by our general rule of no unnecessary coincidences.
2. Two TICs can intersect, as illustrated in Fig. 3(d); this is equivalent to a vertex formed by three tangential surfaces. The possibility that the two TICs are also

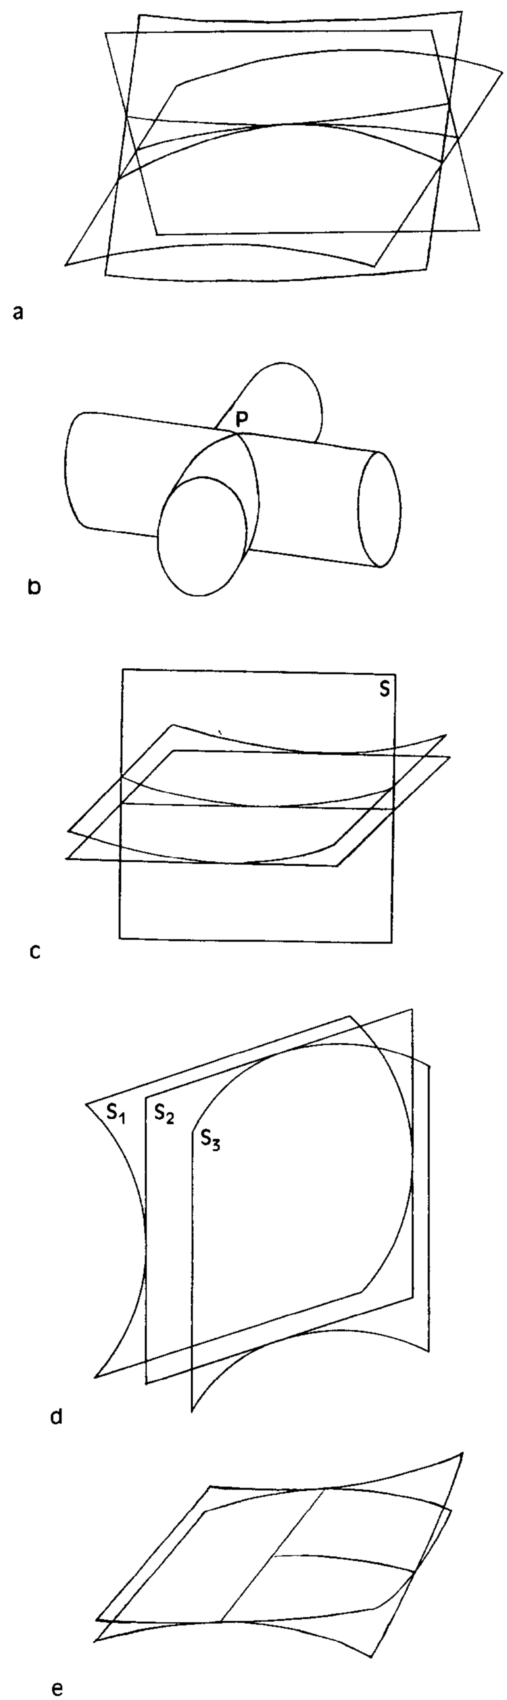

Fig. 3. Five ways in which tangential edges or surfaces can meet to form a vertex.

tangential is disallowed, by our general rule of no unnecessary coincidences. In Fig. 3(d), S_1 and S_3 lie on opposite sides of S_2; another possibility is that they lie on the same side, and hence intersect.

3. The two surfaces which touch or merge can also intersect, as illustrated in Fig. 3(e). Again, we disallow the possibility that this intersection is tangential to the TIC.

In summary, we can differentiate five cases:

(a) E_{ij} tangential to S_k for all $i,j,k = 1,2,3$, and E_{ij} tangential to E_{ik} for all $i,j,k = 1,2,3$.
(b) S_1 kisses S_2.
(c) S_1 tangential to S_2 and their TIC intersects S_3.
(d) S_i tangential to S_j for all $i,j = 1,2,3$.
(e) S_1 tangential to S_2 and they intersect.

In each of the cases given in Fig. 3, 3D space is divided by the surfaces into a certain number of subspaces. Each subspace can contain material or can be empty. By considering all combinations, we can find all possible vertices involving tangential edges or surfaces. Certain highly unlikely vertices can be disallowed at this point. By considering the remaining vertices from every possible viewpoint, we can write down the corresponding list of labelled junctions. The sets of vertices shown in Figs. 4–7 correspond, respectively, to cases (a), (b), (c) and (d) above. The vertices generated by case (e) project into terminal junctions labelled either + or −.

We isolate two labelled junctions in Fig. 8. They are of capital importance because they are *not* visible in the drawing and are not even detectable by analysis of the intensity image. We denote such label transitions on a line as undetectable C-junctions.

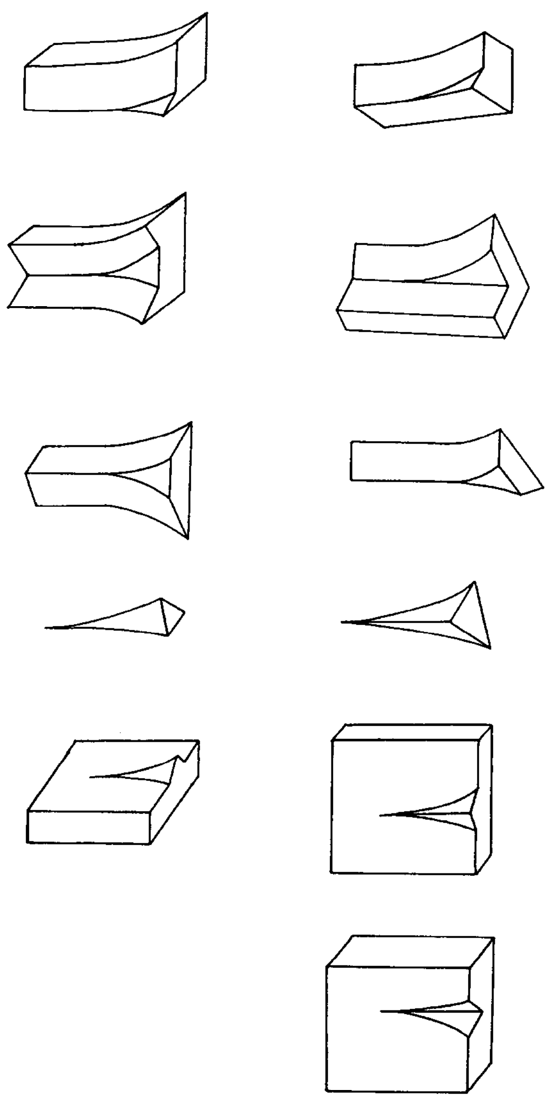

Fig. 4. Vertices formed when the intersection curves of three surfaces are all tangential.

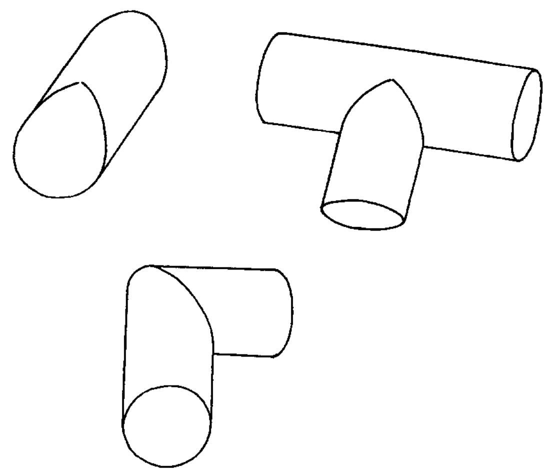

Fig. 5. Vertices formed when two surfaces kiss at a point.

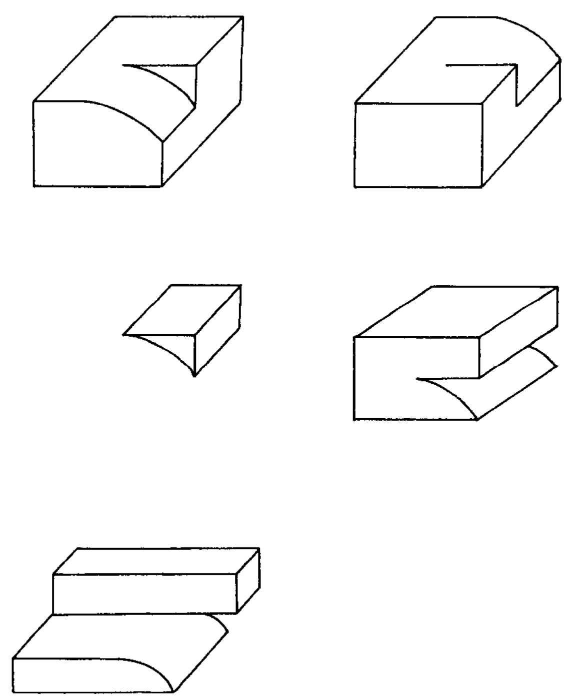

Fig. 6. Vertices formed when a surface intersects the tangential intersection curve of two other surfaces.

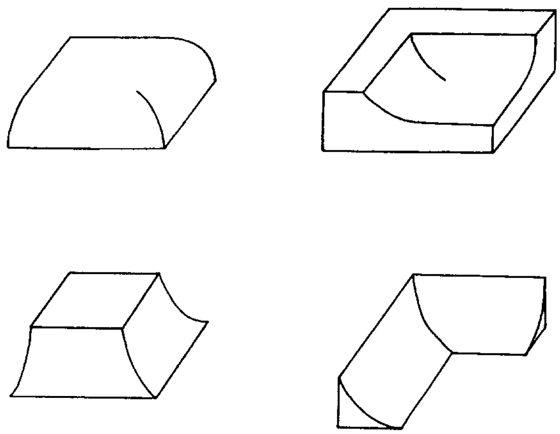

Fig. 7. Vertices formed when three surfaces are all tangential.

Junctions can also be formed by the projection of the intersection of a TIC with an extremal edge; in other words, the surfaces meeting at the TIC are tangential to the viewing direction. An example is shown in Fig. 9.

When two surfaces are tangential they may or may not have the same curvature. Figs. 4–7 show the vertices that result under the condition that any pair of tangential surfaces have the same curvature at their points of intersection. Figs. 10–12 show the extra vertices that result when tangential surfaces do not necessarily have identical curvature. The vertices in Figs. 10, 11 and 12 correspond, respectively, to cases (c), (d) and (e) above. The catalogue of labelled junctions corresponding to the set of vertices shown in Figs. 4–12 is given in Fig. 13. We have, of course, included in the catalogue of Fig. 13 the projections of non-tangential intersections of extremal edges, surface-normal discontinuity edges and smooth edges, as well as junctions caused by occlusion or the termination of a smooth edge [9].

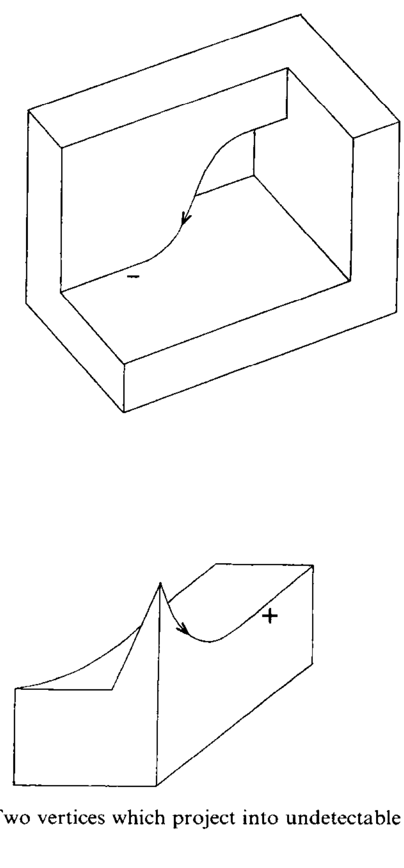

Fig. 8. Two vertices which project into undetectable C-junctions.

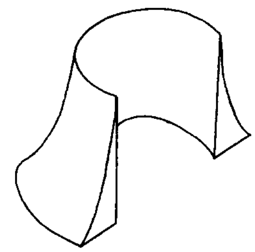

Fig. 9. Vertex formed when the tangential intersection curve of two surfaces is tangential to the viewing direction.

Broken lines represent discontinuities of surface curvature (smooth edges). The projection of a smooth edge in a line drawing is called a roof line or ramp line, because its cross-section in the intensity image has the form of a roof or ramp [13]. Unbroken lines represent surface-normal discontinuity edges (such as the edges of a cube) or extremal edges [14] (such as the boundary of a sphere). The projections of these edges in a line drawing are known as step lines because of the form of their cross-section in the intensity image. A solid arrowhead at the end of a line indicates that this line disappears in the intensity image: in the case of a disappearing roof line, the difference in gradients tends to zero; in the case of a disappearing step line, the step height tends to zero.

In the junctions $L(dd)$, $C(dd)$, $Y(dnn)$ and $W(ndn)$ those lines which disappear as we approach the junction are denoted by d and those which do not as n. Each 0 suffix (in L_0, W_0, Y_0 and W_{00}) represents an angle of 0 degrees and hence a pair of tangential lines. The notation x/y represents a type-y junction composed of step and

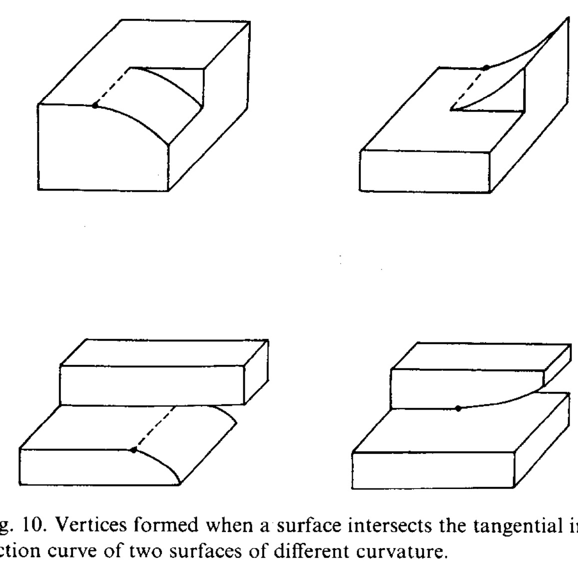

Fig. 10. Vertices formed when a surface intersects the tangential intersection curve of two surfaces of different curvature.

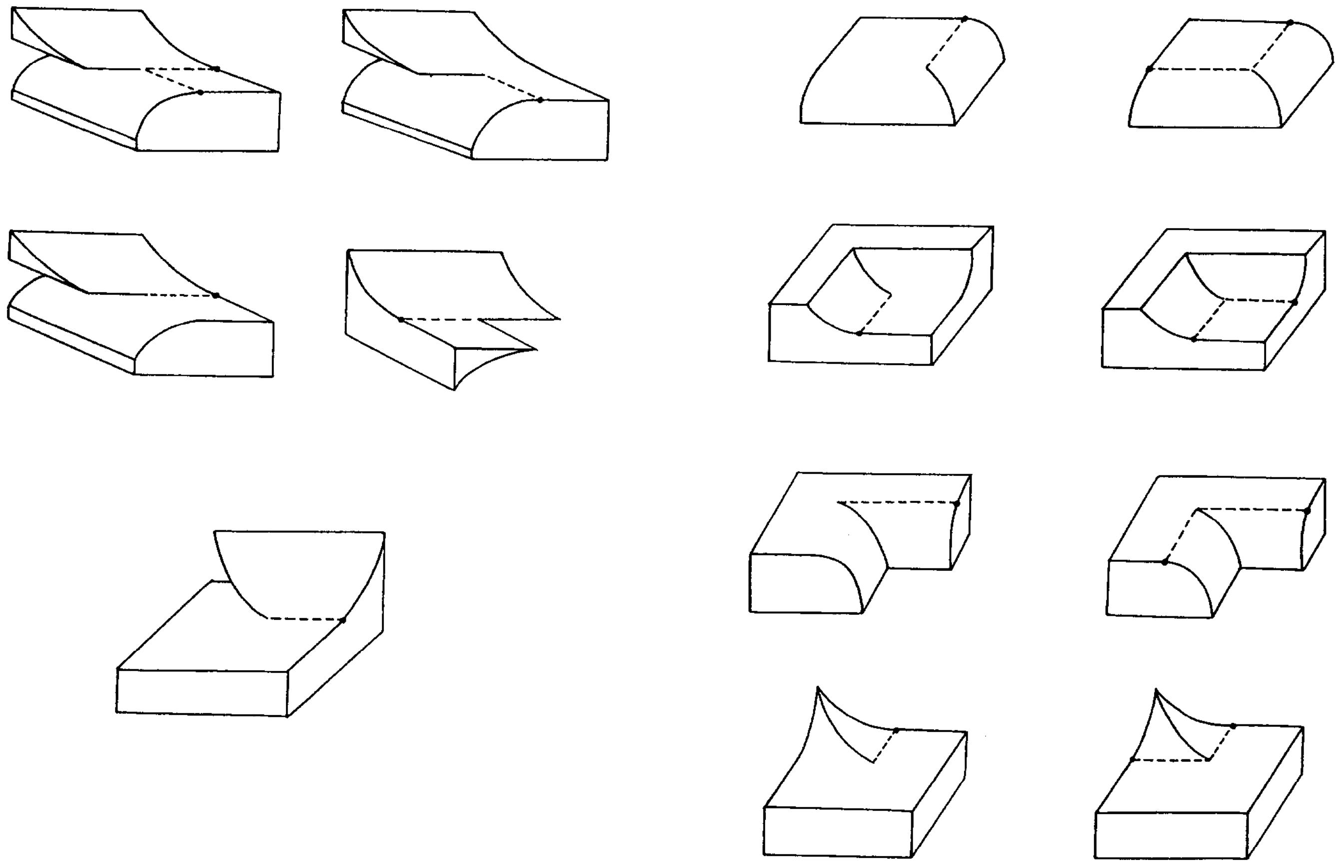

Fig. 11. Vertices formed when three surfaces are all tangential and at least two have different curvature.

ramp lines, in which the step lines alone form a type-x junction. When there is possible ambiguity in determining which lines of the type-y junction are the step lines, then this ambiguity is resolved by denoting in brackets each line as r or s (for ramp and step). At a 3-curvature junction all three lines have distinct curvatures. In a curvature-T junction there is a discontinuity of curvature along the bar of the T. At a Y_0 junction either all three lines have identical curvature or there is a discontinuity of curvature between each pair of lines; there is therefore no possible confusion with a 3-tangent junction.

At each point on a 3D ramp edge E there is a discontinuity in the normal curvature, in the direction orthogonal to E. The label $<$ ($>$) assigned to the projection L of E means that the surface to the right of L is more concave (convex) than the surface to the left of L [9].

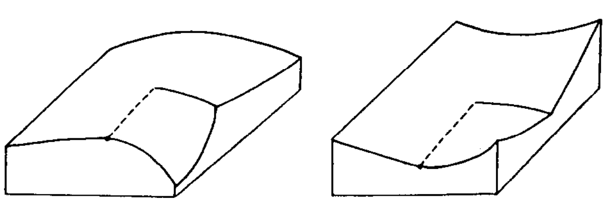

Fig. 12. Vertices formed by the intersection of two tangential surfaces of different curvature.

For reasons of brevity, we have not given the list of labelled junctions for the reflected versions of the following junctions: terminal, curvature-L, 3-tangent, W_0, curvature-L/3-curvature, terminal/L, terminal/$T(srr)$, L/T, $L{:}_0/W_0$. These can easily be generated by the reader.

The junction caused by the intersection of two ramp lines allowed in a previous study [9] has been disallowed in the present study since it requires the intersection of four distinct surfaces, and we here limit ourselves to a maximum of three intersecting surfaces at each vertex.

Some of the vertices illustrated in Figs. 4, 6, 7 and 9 occur in common manufactured objects such as car bodywork, pipework and crockery. On the other hand, some other vertices are highly improbable. Nevertheless, giving a junction catalogue defined by general rules provides greater mathematical accountability and rigour than a junction catalogue derived from an arbitrary sample of objects encountered in a specific application. This research program has as a central aim to introduce more mathematical accountability into computer vision. The danger of considering some highly improbable vertices is that the junction catalogue could be 'diluted' by labellings which hardly ever occur in practice, thus rendering drawings more ambiguous than necessary. This dilution was found to be surprisingly small. Most of the vertices shown in Figs. 4–12 project into new junctions not

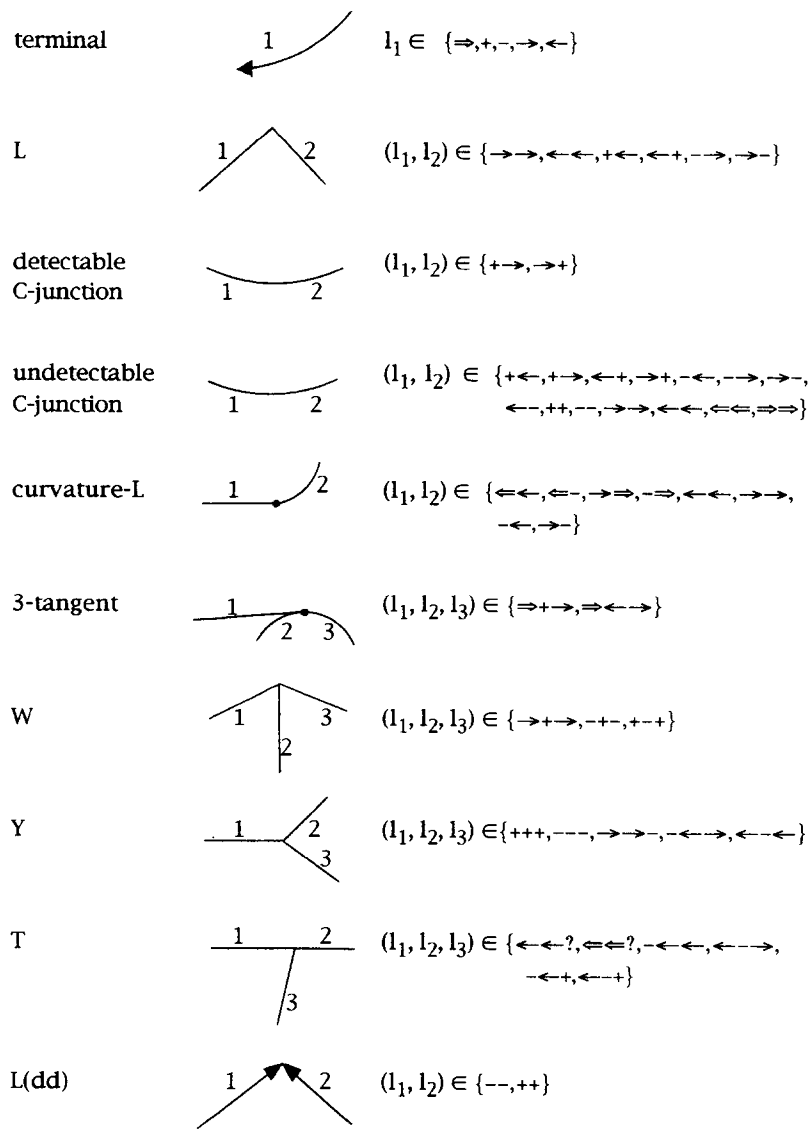

Fig. 13. Catalogue of labelled junctions for curved objects composed of C^3 surfaces in which both edges and surfaces may be tangential.

occurring in the catalogue of Fig. 1. In particular, the set of labellings for L, W, Y and T junctions are unchanged in the catalogue of Fig. 13 compared with the catalogue of Fig. 1.

4. Labelling algorithms

It is known that the problem of determining whether a line drawing has a legal global labelling according to the Huffman–Clowes scheme is NP-complete [15]. Surprisingly, the same problem for the labelling scheme of Fig. 13 can be solved in linear time. Greater freedom in the possible shapes of object surfaces (curved as opposed to planar) necessarily implies slacker constraints compared with drawings of polyhedra. It is this slackening of the constraints which converts an NP-complete problem into a tractable problem.

Since a sequence of undetectable C-junctions can produce a transition from any of the surface-normal discontinuity labels (namely $+$, $-$, $\rightarrow$, $\leftarrow$) to any other surface-normal discontinuity label, the distinction between these four labels is irrelevant when determining whether or not a line drawing has a legal global labelling according to the catalogue of Fig. 13. We can thus replace each of the labels $+$, $-$, $\rightarrow$, $\leftarrow$ by a single label 'd', representing a surface-normal discontinuity edge, in each of the junction labellings of the catalogue. The resulting line drawing labelling problem can be expressed as a Constraint Satisfaction Problem (CSP) [16] in which each line-end is a variable and each junction is a constraint on the combinations of labels which can simultaneously be assigned to the line-ends which meet at the junction. The degree-3 junctions, which correspond to ternary constraints in the CSP, can all be decomposed into unary and/or binary constraints. For example, the

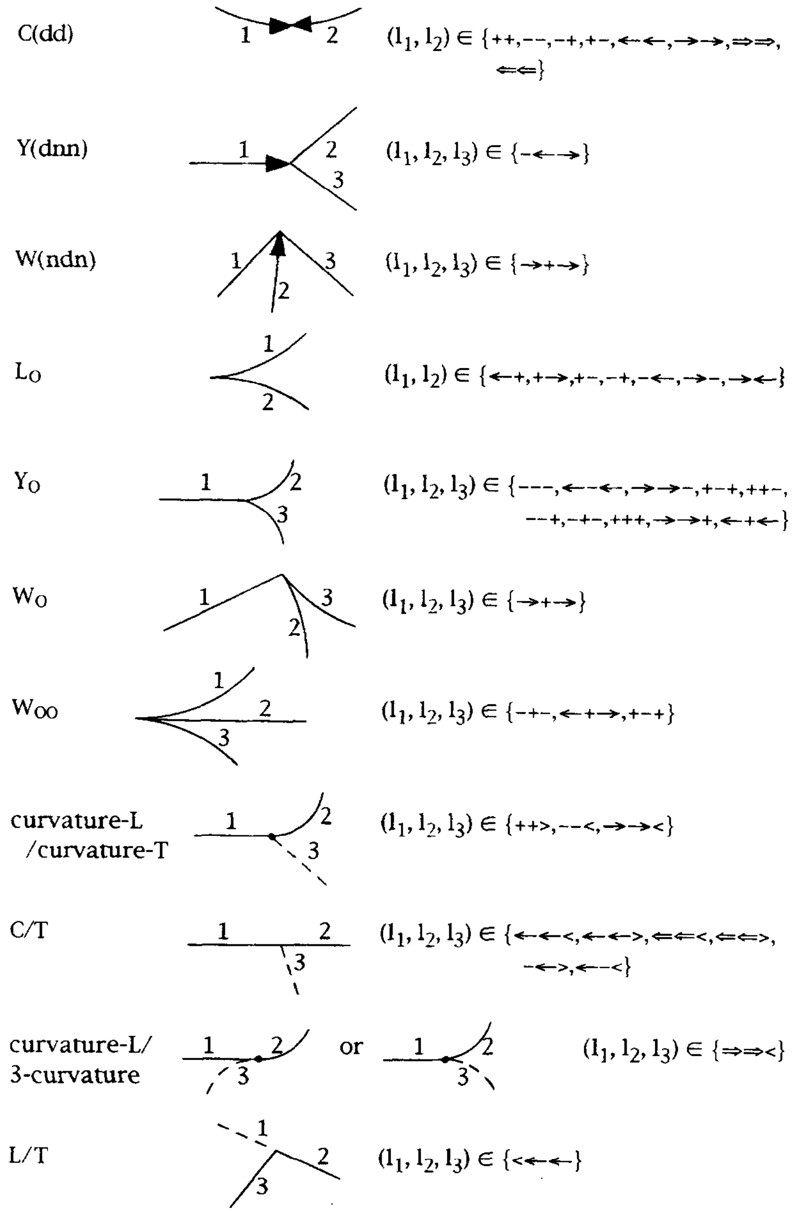

Fig. 13. Continued.

T-junction constraint, given by the list of labellings

$\{ddd,\ dd\Leftarrow,\ dd\Rightarrow,\ \Leftarrow\Leftarrow d,\ \Leftarrow\Leftarrow\Leftarrow,\ \Leftarrow\Leftarrow\Rightarrow\}$

on line-ends $1,2,3$

can be replaced by the binary constraint

$\{dd, \Leftarrow\Leftarrow\}$ on line-ends 1,2

and the unary constraint

$\{d, \Leftarrow, \Rightarrow\}$ on line-end 3.

It is easily verified that each of the resulting binary constraints is a 0/1/all constraint [17,18]. These are constraints C_{ij} on variables i,j in which each possible label for i is consistent with zero, one or all of the possible labels for j and each possible label for j is consistent with zero, one or all of the possible labels for i. Constraint satisfaction problems composed of only unary and 0/1/all binary constraints can be solved in polynomial time: the existence of a legal global labelling can be determined in $O(m^2)$ time and the minimal consistent labelling can be determined in $O(m^3)$ time, where m is the number of variables, i.e. the number of line-ends in the drawing [17].

In fact, curvature-L and terminal/$W(rsr)$ are the only junctions which are not decomposable into unary constraints and/or the binary identity constraint $\{dd, \Leftarrow\Leftarrow, \Rightarrow\Rightarrow\}$. All pairs of line-ends joined by an identity constraint can simply be identified. We can then establish arc consistency, which is a linear operation since the number of constraints is $O(m)$ [19,20]. If any

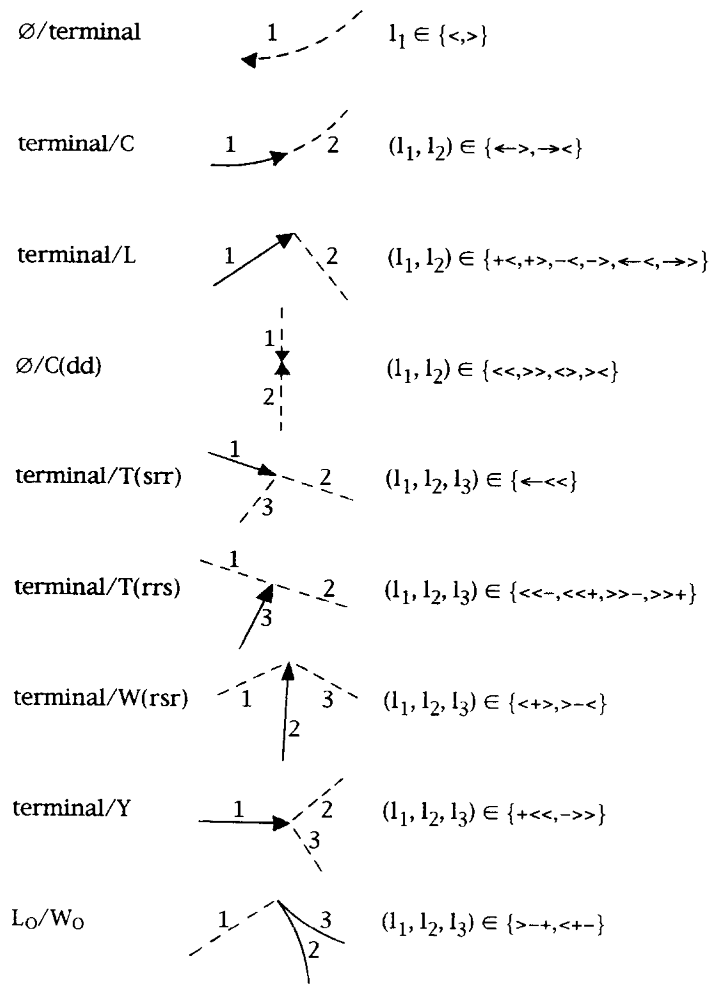

Fig. 13. Continued.

unary constraint is now an empty set, then no legal global labelling exists. Otherwise, the resulting CSP consists of:

(i) isolated variables (line-ends) to which we can arbitrarily assign a label from the corresponding unary constraint,
(ii) strings of curvature-L constraints and unary constraints, for which there is always a legal labelling thanks to arc consistency,
(iii) circuits of curvature-L junctions (without any unary constraints), for which there is always a legal labelling, namely $(d, d, \ldots, d)$.
(iv) strings of the binary constraint $\{<>, ><\}$, derived from the decomposition of the terminal/$W(rsr)$ constraint, and unary constraints; there is always a legal labelling thanks to arc consistency,
(v) circuits of the binary constraint $\{<>, ><\}$, for which there is a legal labelling iff the circuit is of even length.

We can thus clearly determine the existence of a legal global labelling in linear time. In fact, it is easy to show that, if the drawing has at least one legal global labelling, then minimality is automatically achieved by the above algorithm of decomposition, identification and arc consistency.

In a minimal consistent labelling, every element of every constraint can be extended to a legal global labelling [16]. However, even after establishing minimality, the drawing will often still be very ambiguous. It is clear that not all legal labellings of a line drawing are equally likely. One approach is to try to minimise the number of undetectable C-junctions in the global labelling. The resulting optimisation problem can be solved by branch and bound, although the worst case time complexity is clearly an exponential function of the number of lines in the drawing.

In fact, this optimisation problem is NP-hard, since any polynomial time algorithm to solve it could also be used to determine whether a line drawing has a legal labelling according to the Huffman–Clowes labelling scheme [2,3]. This is because the list of legal labellings for L, Y and W junctions are identical in the Huffman–Clowes catalogue and in the catalogue of Fig. 13. NP-hardness follows from the fact that the essential constructions in the proof of NP-completeness of the Huffman–Clowes scheme [15] only make use of L, Y and W junctions. T-junctions are also employed in the proof, to deal with 'hanging' lines, but we can simply let these 'hanging' lines fade out in the drawing; this has the same desired effect of adding no new constraints.

Heuristic optimisation methods, such as probabilistic relaxation [21], avoid the combinatorial explosion but at the cost of not necessarily finding the optimal solution. A probabilistic relaxation algorithm based on the catalogue of Fig. 13 was run on a sample of line drawings each containing only approximtely 10 junctions. Results were disappointing since the resulting unique labelling was often an apparently random patchwork of legal labellings for parts of the drawing. The lack of a global view often prevented the algorithm from finding a legal global labelling. This was no doubt partly due to the insufficient information supplied by the catalogue (see following section) and partly due to the non-convexity of the search space. We also note in passing that a probabilistic relaxation algorithm has an unfortunate bias towards labels which occur in many legal junction labellings. Thus, for example, the label '-' for the lines which meet at a Y-junction propagates more strongly to adjacent junctions, since it occurs in twice as many legal labellings for the Y-junction as other labels. The adjacent junctions then propagate back this *same* information at the next iteration. The algorithm tends to converge to a state in which the Y-junction is labelled '- - -', even though all six labellings for a Y-junction are a priori equally likely.

5. Predictive power of catalogues of junction labellings

To compare the utility of different catalogues of legal junction labellings, we choose to compare the probability, given a random line drawing and a random global labelling of this drawing, that the labelling is a legal labelling of the drawing. This probability is given by

$$\left(\prod p_t^{n_t}\right) \cdot q^{m_s} \cdot r^{m_r} \qquad (1)$$

where the product is over all junction types t (W, T, Y, etc.) except undetected C-junctions, and

- p_t = probability that a random labelling is a legal labelling for a type t junction;
- n_t = number of type t junctions in the drawing;
- q = probability that a random ordered pair of labels is a legal transition due to the possible occurrence of a sequence of undetected C-junctions along a step line;
- m_s = number of step lines in the drawing;
- r = probability that a random ordered pair of labels is a legal transition on a ramp line;
- m_r = number of ramp lines in the drawing.

In a random global labelling of the drawing, it is assumed that each of the step-line labels is equiprobable for each step line, and that each of the ramp-line labels is equiprobable for each ramp line. In fact, in a global labelling, a label is assigned to each line-end. Each line thus has two labels and the total number of labels to be determined is $m = 2(m_s + m_r)$. We are solving the sparse labelling problem [22]. When there are no ramp lines in the drawing, $m_r = 0$. In fact $r = 0.5$ since the only 'transitions' which are possible on a ramp line are $<<$ and $>>$. If it is known that the drawing contains only polyhedra, then there are no C-junctions and q = probability that the random labels at the two ends of a line are the same (1/4, since there are 4 step-line labels). The values of p_t, q and r are functions of the catalogue.

To produce a measure which is independant of the size of the drawing, we take the mth root of the probability given by Eq. (1). We define the predictive power (pp) of the catalogue to be minus the logarithm of the resulting value:

$$\mathrm{pp} = -\left(\sum (n_t/m)\log_2 p_t\right) - (m_s/m)\log_2 q - (m_r/m)\log_2 r$$

We can interpret pp as the average number of bits of information provided by the catalogue per line-end in a random drawing.

For simplicity, we assume that the different junction types in the catalogue are all equiprobable. Thus n_t is the same for all values of t, and the values of m_s/m and m_r/m can be deduced by counting the number of step and ramp lines in each junction type. This gives the values shown in Table 1 for certain catalogues of legal junction labellings.

Polyhedra are assumed to be labelled according to the Huffman–Clowes scheme [2,3], to which we have added the junction labellings corresponding to non-occlusion T-junctions (such as N in Fig. 2). The C^3 surfaces catalogue is the refined version of the catalogue due to Malik [8] for objects with piecewise C^3 surfaces, given in Fig. 1. The tangential C^3 surfaces catalogue refers to the catalogue given in Fig. 13, for objects with possibly tangential edges and surfaces. For the moment it is assumed that we can detect no C-junctions, and hence the set of legal labellings for a C-junction is the union of the sets of labellings for detectable and undetectable C-junctions.

Table 1
Comparison of the quantity of information supplied by three different catalogues of labelled junctions

	pp	$\mathrm{pp}_{\max}/2$
Polyhedra	2.14	2.00
C^3 surfaces	2.69	2.58
Tangential C^3 surfaces	2.07	2.19

Assuming that each junction (including the undetected C-junction) has at least one labelling, the value of pp satisfies $0 \leq \text{pp} \leq \text{pp}_{\max}$, where

$$\text{pp}_{\max} = 2(m_s \log_2 a_s + m_r \log_2 a_r)/(m_s + m_r)$$

where a_s is the number of different step-line labels (4 for polyhedra, and 6 for the other catalogues) and $a_r = 2$ is the number of ramp-line labels. The value of $\text{pp}_{\max}$ follows from the inequalities

$$\sum n_t \log_2 p_t \geq -2(m_s \log_2 a_s + m_r \log_2 a_r)$$

$$-\log_2 q \geq 2 \log_2 a_s$$

$$-\log_2 r \geq 2 \log_2 a_r$$

It can easily be shown that the expected number of legal global labellings for a random line drawing is given by

$$2^{m((\text{pp}_{\max}/2)-\text{pp})}$$

This tends to zero as the number of lines tends to infinity if and only if $\text{pp} > \text{pp}_{\max}/2$. Although, for example, the C^3 surfaces catalogue provides more bits of information than the polyhedra catalogue, it is slightly less likely to produce a unique and unambiguous labelling, since pp is closer to $\text{pp}_{\max}/2$ than for the polyhedra catalogue.

It should be noted that the values presented here are over-optimistic for two reasons. Firstly, we are unlikely to be able to detect and correctly identify all junctions in the drawing, as we have implicitly assumed. This clearly becomes harder when junctions involve ramp lines. Secondly, we consider a random line drawing to be chosen among all drawings, even those which represent impossible objects, and hence have no global legal labelling. If we consider random drawings chosen among drawings with at least one global legal labelling, then the probability of a random labelling being legal is much greater. Often, for every legal labelling, there are several other legal labellings which differ on only a small number of line-ends. This phenomenon has been observed in other constraint satisfaction problems [23], but its quantitative effect in this case is difficult to estimate.

6. Extra information from the intensity image

It is clear from Table 1 that the catalogue of Fig. 13 provides insufficient information on its own to give an unambiguous labelling of an average drawing, since $\text{pp} < \text{pp}_{\max}/2$. This is especially true when analysing imperfect line drawings derived from real images. To increase the likelihood of obtaining an unambiguous labelling, we can use other sources of information, such as shading, colour or texture information from the intensity image, multiple drawings of the same objects or models of the objects.

Shape-from-shading [24] is an important tool in interpreting images of curved objects. For example, if extremal lines are detected by analysis of the intensity image [22], and if they account for 1/6 of the lines, then this provides an average of up to $(1/6)\log_2 6 = 0.43$ bits of information per line-end. This is only an upper bound since, for example, some extremal lines will have already been identified due to the fact that they terminate at a 3-tangent junction. Detecting those C-junctions which are the result of the viewing direction being tangential to an object surface, by analysis of the intensity image, increases pp by an amount which is dependant on the catalogue. We have also already implicitly assumed that edges which fade out at junctions have been detected by analysis of the intensity image.

To explain how extremal lines can be detected in the intensity image, consider the simple image of a cylinder in Fig. 14. Let f be the brightness function. Then the partial derivative of f tends to plus or minus infinity as we approach the extremal line L from the left. In general, the partial derivative of f in a direction n normal to the (possibly curved) extremal line L tends to plus or minus infinity as we approach L in the direction of n and on the side of L which is the projection of the surface which is tangential to the viewing direction. (In fact, this is true for any direction n which is not actually tangential to the line L.)

Fig. 14 also shows a C-junction C which is the projection of a point at which the viewing direction is tangential to the object surface. This, too, can be detected in the intensity image, since it is an isolated point having the characteristics of an extremal line, i.e. the partial derivative of f in a direction n normal to the line tends to plus or minus infinity as we approach the line along n. Such C-junctions are denoted as detectable C-junctions in the catalogues of Figs. 1 and 13. The

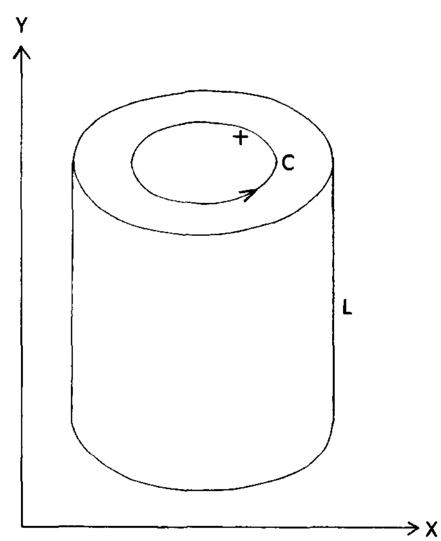

Fig. 14. An external line L and a C-junction C both of which can be detected in the intensity image.

Table 2
Comparison of the quantity of information supplied by different catalogues of labelled junctions when extremal edges and detectable C-junctions are detected in the intensity image

	pp	$\mathrm{pp}_{\max}/2$
Polyhedra	2.14	2.00
C^3 surfaces	3.28	2.58
Tangential C^3 surfaces	2.70	2.20

surface which is tangential to the viewing direction is assumed to be above the junction in Figs. 1 and 13.

Our assumption of C^3 surfaces disallows many improbable objects, but unfortunately also disallows apices of cones. If we decide to allow apices of cones, then these can immediately be detected in the intensity image, since the apex of a cone is simply the junction of two extremal edges.

Table 2 shows the values of pp when extremal lines and detectable C-junctions are detected by analysis of the intensity image, for each of the three catalogues. The value of pp remains unchanged for the case of polyhedral objects, since the line drawings contain no extremal lines or C-junctions.

The relatively high value of pp for the case of objects with tangential edges and surfaces is due to the fact that approximately one-third of all lines are uniquely labelled as extremal lines, and does not by any means imply that most drawings will have a unique legal labelling. Although some C-junctions can be detected by analysis of the intensity image, this is not the case for all C-junctions when objects may have tangential edges or surfaces (see Fig. 8). This means that propagation of the labels $+$, $-$, $\rightarrow$ is impossible, since we have no guarantee that the two ends of a line should be labelled by the same semantic label.

Nonetheless, propagation of all labels *is* possible for line drawings of objects with C^3 surfaces without tangential surfaces and edges (Fig. 1) when detectable C-junctions are indeed detected in the intensity image. For this class of curved objects, we can even conclude from Table 2 that the identification of extremal lines and detectable C-junctions renders the labelling problem easier than the same problem for polyhedra.

One simple solution to the problem of undetectable C-junctions is to use a reduced label set. Wong [25] employs labels which mark the absence or presence of an object on either side of the line. Three line labels are thus possible: absence/presence, presence/absence, presence/presence. These line labels do not change at C-junctions since an object remains present on both sides of the line despite transitions such as $+$ to $\rightarrow$. However, this scheme is only capable of identifying the occluding contour of the object and provides no information about its internal shape. In fact, many workers assume that the occluding contour (the boundary between the object(s) and the background) can be identified in the drawing. If so, then this too provides a valuable extra source of information.

7. Discussion

7.1. Lighting conditions

It should be noted that our model of line drawing formation is unrealistic, since it ignores shadows and missing lines due to contrast failure. Extremal edges and C-junctions cannot be identified if the surface which is tangential to the viewing direction is in complete shadow. Furthermore, ramp lines can be caused by shadows cast by non-point light sources, and the reflection of a point light source off a planar Lambertian surface provides just such a non-point light source. Under diffuse lighting it is well known that concave edges appear in the intensity image as ramp edges. There are therefore several potential sources of 'spurious' ramp lines not caused by smooth edges. In particular we should be wary of interpreting a C/T junction as evidence of occlusion.

An occluding edge disappears due to contrast failure when the occluding and occluded surfaces are parallel in three-dimensional space and the two surfaces have identical surface characteristics, and are subject to identical lighting conditions.

In this paper we have extended previous work by allowing objects of more general shape. We felt it necessary to iron out problems concerning object shape before passing on to more realistic models of line drawing formation. Further work is clearly required to incorporate shadows, specular reflection, non-point light sources, contrast failure, surface markings, transparent objects and imperfect edge detection.

7.2. Object shape

During the investigation of curved objects with tangential edges and surfaces, we were surprised by the myriad of possible vertices. We have necessarily made subjective choices in selecting those vertices which we consider to be the most likely. For simplicity of presentation, we have been forced to omit many vertices which, in certain applications, could be considered more plausible than some of the vertices we have included in the derivation of the catalogue of Fig. 13. As examples of such vertices we considered the following possible constructions:

- the intersection of two smooth edges,
- the tangential intersection of a TIC and a surface-normal discontinuity edge
- the intersection of E and E', where E and E' are two

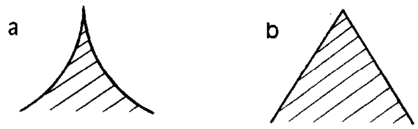

Fig. 15. Cross-sections of two distinct types of surface-normal discontinuity edges.

distinct edges formed by the intersection of the same pair of surfaces S_1 and S_2. Each of E and E' may be a surface-normal discontinuity edge or a TIC.

We found that, by including any combination of these particular examples, the constraints remain decomposable into binary 0/1/all constraints. Our conclusion is that the complexity results in relation to the catalogue of Fig. 13, are robust to small changes in the assumptions about object shape. We conjecture that for any natural class of objects which includes tangential edges and surfaces, the existence of a legal global labelling can be determined in low-degree polynomial time.

Another simplification we have made is to neglect the difference between the two types of edges shown in Fig. 15. When viewed from the left, both these edges project into a line which we label '→'. However, we also investigated an alternative labelling scheme in which the tangential occluding edge of Fig. 15(a) was assigned a new label. We could not reasonably make the distinction between these two types of edge without allowing the possibility of a transition of an edge of type Fig. 15(a) into an edge of type Fig. 15(b). This provides a new labelling for undetectable C-junctions, which means that the distinction between these two edge labels is irrelevant to the problem of determining the existence of a legal global labelling. Thus again this problem can be solved in linear time. The predictice power of this alternative labelling scheme is given by

$$pp = 2.20 \quad pp_{max}/2 = 2.50$$

Compared to the predictive power of the catalogue of Fig. 13 ($pp = 2.07$, $pp_{max}/2 = 2.19$), this scheme thus produces a modest increase in information per line-end, at the cost of reducing the probability that the drawing has an unambiguous labelling.

An important restriction we have made is the assumption that no more than three surfaces meet at a vertex. The projections of vertices caused by the intersection of four surfaces can easily be incorporated into the junction catalogue. We considered, however, that the tangential intersection of some of the four surfaces meeting at such a vertex required too many coincidences to be incorporated into the catalogue. As far as the propagation of labels is concerned, we need not distinguish between the surface-normal discontinuity edge labels (+, −, →, ←), which are all grouped into the single label 'd'. With the reduced label set $\{\Leftarrow, \Rightarrow, d, <, >\}$, the only new junction labellings caused by the projection of vertices at which four non-tangential surfaces intersect are four-line junctions with the unique labelling (d, d, d, d). This has no effect on the existence of a linear-time labelling algorithm.

7.3. Realisability

We claim that, for any natural set of curved objects (with or without tangential edges and surfaces), any legally-labelled drawing can be physically realised. Imagine a flat rubber sheet laid over the drawing. By creating minuscule ridges, valleys, folds and ripples we can create convex edges (+), concave edges (−), non-extremal occluding edges (→) and extremal edges (⇒). All the labelled junctions in the catalogue can by physically realised by local deformations of the rubber sheet in the vicinity of the junction. At this point we may have to introduce hidden 3D edges (for example in the case of an L-junction). These edges need never be visible, since we can assume that they terminate almost immediately at a terminal junction (in the case of step lines) or a ø/terminal junction (in the case of ramp lines). We end up with an implausible but physically possible interpretation of the drawing as a flat rubber sheet with small local deformations in the vicinity of lines and junctions.

This argument no longer holds if we have extra information, such as that straight lines in the drawing are necessarily projections of straight lines in the scene.

8. Conclusion

We have chosen to study a class of objects with possibly tangential edges and surfaces which include many man-made objects. The resulting catalogue of labelled junctions has the pleasing property that we can determine the existence of a legal global labelling in linear time.

Our conclusion, by calculation of the information content of line drawings per line-end, is that a catalogue of labelled junctions provides insufficient information by itself to uniquely determine the correct semantic labelling of a line drawing of curved objects involving possibly tangential edges and surfaces. This is mainly due to the presence of undetectable label transitions on lines. Optimisation techniques can be used to minimise the number of invisible label transitions but this optimisation problem is unfortunately NP-hard.

Various sources of extra information exist which can help to reduce ambiguity, including notably the detection of extremal edges and C-junctons by analysis of the intensity image. Further research on the generalisation of this work to include more realistic drawings must necessarily be accompanied by the search for alternative sources of information.

Appendix: glossary of technical terms

Surface-normal discontinuity edge: boundary in 3D between two surface patches of an object presenting a discontinuity in surface normal.

Extremal edge/virtual edge: locus of points in 3D at which the viewing direction is tangential to an object surface.

Extremal line/limb: projection of an extremal edge in the drawing.

Occluding edge: an extremal edge or a surface-normal discontinuity edge such that the surfaces which are visible on the two sides of this edge present a depth discontinuity along this edge and/or belong to different objects.

TIC: tangential intersection curve C of two surfaces which touch or merge along C.

Step edge: a step discontinuity in the intensity image. Under the assumptions of this paper, this is necessarily caused by the projection of a surface-normal discontinuity edge or an extremal edge.

Step line: line in the drawing corresponding to a step edge in the intensity image.

Smooth edge: boundary in 3D between two surface patches with continuous surface normal but presenting a discontinuity of surface curvature.

Ramp edge/roof edge: a gradient discontinuity in the intensity image. Under the assumptions of this paper, this is the projection of a smooth edge into the intensity image.

Ramp line/roof line: line in the drawing corresponding to a ramp edge in the intensity image.

Occluding contour: the contour of the projection of a single object or collection of objects in the line drawing.

Invisible vertex/phantom vertex: a point in 3D which projects into a C-junction.

C-junction/invisible junction/phantom junction: point on a line at which the semantic label of the line changes without any visible reason.

Detectable C-junction: projection of an isolated point at which the viewing direction is tangential to an object surface.

Non-occlusion T-junction: T-junction caused by the projection of an object vertex and not by the occlusion of one edge by another.

References

[1] A. Guzman, Decomposition of a visual scene into three-dimensional bodies. Proc. Fall Joint Comput. Conf., 1968, 291–304.

[2] D.A. Huffman, Impossible objects as nonsense sentences. In: Meltzer, B. and Mitchie, D. (eds.), Machine Intelligence, Edinburgh University Press, UK, 1971, 295–323.

[3] M.B. Clowes, On seeing things. Artificial Intelligence, 2 (1971) 79–116.

[4] D. Waltz, Understanding line drawings of scenes with shadows. In: Winston, P.H. (ed.), Psychology of Computer Vision, McGraw-Hill, New York, 1975, 19–91.

[5] G. Falk, Interpretation of imperfect line data as a three-dimensional scene. Arficial Intelligence, 3 (1972) 101–144.

[6] R. Shapira and H. Freeman, Computer description of bodies bounded by quadric surfaces from a set of imperfect projections. IEEE Trans. Comput., 27(9) (1978) 841–854.

[7] R. Shapira and H. Freeman, The cyclic order property of vertices as an aid in scene analysis. Comm. ACM, 22(6) (1979) 368–375.

[8] J. Malik, Interpreting line drawings of curved objects. Int. J. Comput. Vision, 1 (1987) 73–103.

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

[10] P. Parodi and V. Torre, On the complexity of labeling perspective projections of polyhedral scenes. Artificial Intelligence, 70 (1994) 239–276.

[11] R. Bergevin and M.D. Levine, Extraction of line drawing features for object recognition. Pattern Recognition, 25(3) (1992) 319–334.

[12] J.J. Koenderink and A.J. van Doorn, The shape of smooth objects and the way contours end. Perception, 11 (1982) 129–137.

[13] M. Petrou and J. Kittler, Optimal edge detectors for ramp edges. IEEE Trans. Pattern Analysis and Machine Intelligence, 13(7) (1991) 483–490.

[14] Q. Zhu, Virtual edges, viewing faces, and boundary traversal in line drawing representation of objects with curved surfaces. Computers and Graphics, 15(2) (1991) 161–173.

[15] L.M. Kirousis and C.H. Papadimitriou, The complexity of recognizing polyhedral scenes. J. Computer and System Sciences, 37(1) (1988) 14–38.

[16] E. Tsang, Foundations of Constraint Satisfaction, Academic Press, 1993.

[17] M.C. Cooper, D.A. Cohen and P.G. Jeavons, Characterising tractable constraints. Artificial Intelligence, 65 (1994) 347–361.

[18] L.M. Kirousis, Fast parallel constraint satisfaction. Artificial Intelligence, 64 (1993) 147–160.

[19] R. Mohr and T.C. Henderson, Arc and path consistency revisited. Artificial Intelligence, 28 (1986) 225–233.

[20] C. Bessière Arc-consistency and arc-consistency again. Artificial Intelligence, 65 (1994) 179–190.

[21] A. Rosenfeld, R.A. Hummel and S.W. Zucker, Scene labeling by relaxation operations. IEEE Trans. Systems, Man and Cybernetics, 6(6) (1976) 420–433.

[22] J. Malik and D. Maydan, Recovering three-dimensional shape from a single image of curved objects. IEEE Trans. PAMI, 11(6) (1989) 555–566.

[23] M.C. Cooper, Estimating optimal parameters for parallel database hardware. Int. J. Systems Sci., 23 1 (1992) 119–125.

[24] B.K.P. Horn, Robot Vision, MIT Press, Cambridge, MA, 1986.

[25] E.K. Wong, Model matching in robot vision by subgraph isomorphism. Pattern Recognition, 25(3) (1992) 287–303.