

# An Optimization-Based Approach to the Interpretation of Single Line Drawings as 3D Wire Frames

YVAN G. LECLERC AND MARTIN A. FISCHLER

*Artificial Intelligence Center, SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025*

Received

## Abstract

Line drawings provide an effective means of communication about the geometry of 3D objects. An understanding of how to duplicate the way humans interpret line drawings is extremely important in enabling man-machine communication with respect to images, diagrams, and spatial constructs. In particular, such an understanding could be used to provide the human with the capability to create a line-drawing sketch of a polyhedral object that the machine can automatically convert into the intended 3D model.

A recently published paper (Marill 1991) presented a simple optimization procedure supposedly able to duplicate human judgment in recovering the 3D “wire frame” geometry of objects depicted in line drawings. Marill provided some impressive examples, but no theoretical justification for his approach. Here, we introduce our own work by first critically examining Marill’s algorithm. We provide an explanation for why Marill’s algorithm was able to perform as well as it did on the examples he presented, discuss its weaknesses, and show very simple examples where it fails. We then provide an algorithm that improves on Marill’s results. In particular, we show that an effective objective function must favor both symmetry and planarity—Marill deals only with the symmetry issue. By modifying Marill’s objective function to explicitly favor planar-faced solutions, and by using a more competent optimization technique, we were able to demonstrate significantly improved performance in all of the examples Marill provided and those additional ones we constructed ourselves. Finally, we examine some questions relevant to the implications of this work for understanding the human ability to interpret line drawings.

## 1 Introduction

The interpretation of line drawings has been an important focus for research in machine vision since the field’s inception. There seems to be little question that human subjects can easily recover 3D models from 2D line drawings depicting many classes of objects. One such class of special interest has been called the “blocks world.” This class consists primarily of polyhedral solids in 3D Euclidean space and the projections of the visible edges of these objects onto a 2D plane (which we call the line drawing). Given a single line drawing of a blocks world scene, normal human subjects will usually arrive at the same 3D interpretation, even though there may be a very large number of possible 3D objects that could have produced the given drawing.

Beginning with the work of Guzman in 1968, there has been a concerted effort by vision researchers to

develop an algorithmic procedure that could duplicate human performance in interpreting line drawings, at least with respect to blocks world objects. A significant body of work in this area was produced by such prominent scientists as Clowes (1971), Huffman (1971), Waltz (1972), Mackworth (1973), Kanade (1980), Draper (1981), and Sugihara (1982, 1984). However, the problem as originally formulated, devising a procedure for recovering psychologically plausible 3D models from line drawings, remains unsolved. (A psychologically plausible reconstruction of a line drawing is the one that virtually all people will accept.)

The earliest work by Guzman was heuristic in nature, failed in many cases where humans had no trouble in finding appropriate interpretations, and did not actually return a 3D model, but rather partitioned the scene into separate polyhedral objects. Clowes, Huffman, Waltz, Mackworth, and Kanade formalized and

extended the work of Guzman, but did not solve the original problem. They were (usually) able to label the edges of the line drawing to correctly reflect a consistent 3D interpretation if one existed, or could assert that the drawing did not correspond to a realizable blocks world scene. Mackworth and Kanade explicitly exploited the planarity of the faces of blocks world and “Origami” objects (by employing a “gradient space” representation) to accomplish a form of semiquantitative recovery. In addition to consistent edge labeling, they could also constrain the relative orientation of the faces of the target 3D model. The labels could describe the edges as being convex, concave, occluding, and so forth, but still, for the general case, no explicit 3D model was returned (without introducing additional constraints) and the algorithms would make occasional errors.<sup>1</sup>

In a series of papers, Sugihara reformulated the realizability and recovery problems for line drawings of polyhedra (both with and without hidden lines removed) in purely algebraic terms. He required as input a specification of the vertexes defining each of the individual planar faces of the polyhedra, and also required that the implied line drawing be a general-position projection of the polyhedra. With this approach he succeeded in providing an algebraic criterion as a necessary and sufficient condition for a line drawing to represent a physically realizable polyhedral object. He could also constrain the space of feasible solutions, and obtain a unique solution if enough additional constraints were provided. These additional constraints were obtained from information beyond that provided by the line drawing (e.g., shading or texture information). Sugihara’s work was an important advance, but again it fell short of the original goal. It will rarely be the case that a unique reconstruction is implied by the line drawing, and thus the primary objective of duplicating human performance in this regard is not met.<sup>2</sup>

Our motivation for writing this article was supplied, in part, by a recent publication authored by T. Marill (1991). He refocused on the original problem of human interpretation of single line drawings as 3D structures; he did not restrict his universe to blocks world objects nor did he demand that the line drawings be complete. The surprising thing about his work was that he used an optimization approach involving (seemingly) an almost trivial objective function, and the simplest possible descent algorithm to find a solution, and yet provided examples of reconstructed objects that were, intuitively, extremely good. (Figure 1, examples A

through I, shows the line drawings used in Marill’s experiments.) However, his paper provided no justification for why the algorithm should work, and thus no basis for judging its generality or insight into how it could be improved (should this be desirable).

The first reference we have found that presents the case for choosing between various interpretations of a line drawing based on an objective function is Hochberg and McAlister (1953). In their paper, they “showed that: (1) some variants of the Necker cube are more likely to be described as 2D figures, and some are more likely to be described as 3D; and (2) these differences could be predicted by an objective and plausible coding scheme. Within this scheme, the economy of description was assessed by (among other measures) the number of lines and angles contained within the coding. Thus, the costs and benefits of 2- versus 3-D interpretations could be assessed. Figures that could be coded more simply under a depth interpretation were, in fact, seen in depth; those that could not be simplified in this way were seen to lie in the picture plane” (Pomerantz & Kubovy 1981, pp. 439–440).

Barrow and Tenenbaum (1981) suggested ideas similar to Marill’s for interpreting line drawings (both for simple closed curves and polyhedra), but did not pursue the ideas in greater depth. More recently, Barnard and Pentland (1983) and Pentland and Kuo (1990) have pursued Barrow and Tenenbaum’s approach for simple curves and line drawings of surfaces by finding the smoothest curve (or surface) corresponding to the line drawing.

In this article we introduce our own work by first critically examining Marill’s algorithm. We provide an explanation for why Marill’s algorithm was able to perform as well as it did on the examples he presented, discuss its weaknesses, and show very simple examples where it fails (figure 1, examples J through N). We then provide an algorithm that improves on Marill’s results for all nine of his examples, and also successfully deals with the simple cases where Marill fails. Finally, we examine some questions relevant to the implications of this work for understanding the human ability to interpret line drawings.

We see the work described here as being of both theoretical and practical interest. The practical utility of this work is its relevance to man-machine communication about 3D structures via line drawings—in particular, providing the human with the capability to create a line-drawing sketch of a polyhedral object that the machine can automatically convert into the intended

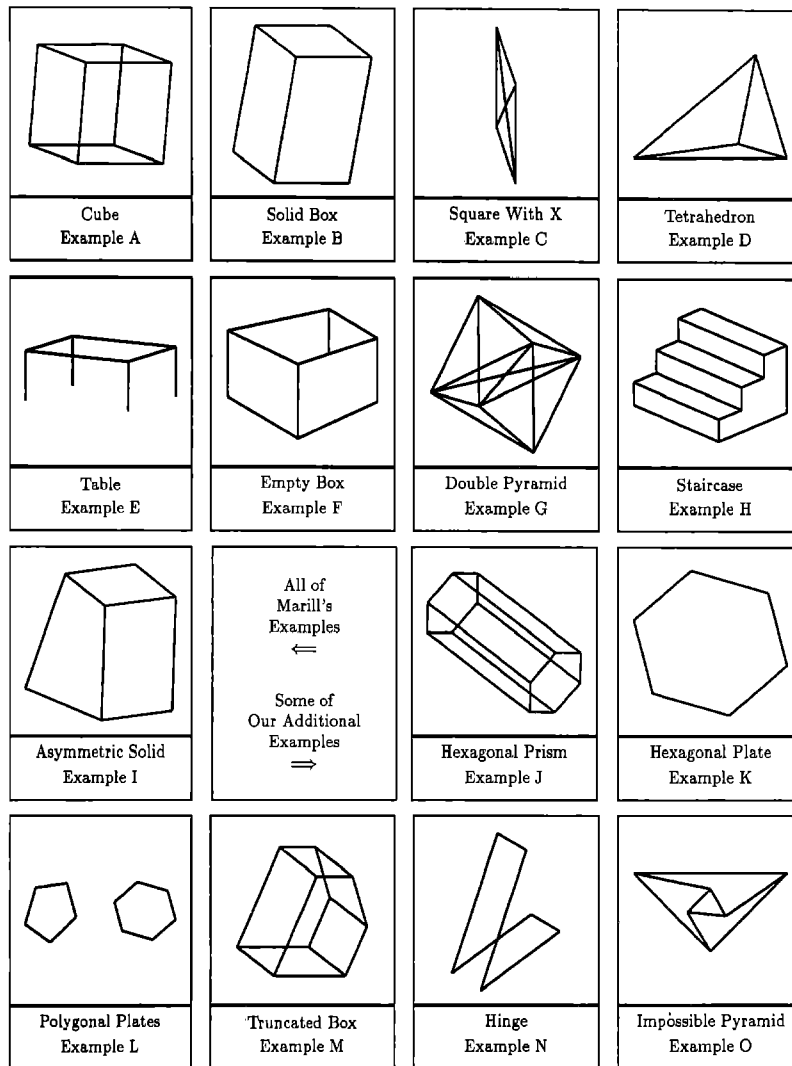


Fig. 1. The line drawings examined in this article. Examples A through I are taken from Marill's paper. Examples J through N are line drawings introduced here for which Marill's algorithm failed to recover a psychologically plausible 3D model. Example O is a line drawing for which a psychologically plausible 3D model is not feasible.

3D model. Deficiencies in providing a complete theory are not fatal, since auxiliary information can always be supplied interactively to resolve ambiguities, but the underlying theory should reduce this "side communication" to a minimum.

## 2 Marill's MSDA Algorithm

Marill's algorithm consists of two components, an objective function and a simple descent optimization procedure for finding a local minimum of this objective function. The objective function is simply the standard deviation of all of the angles (SDA) in the recovered 3D

object with respect to their common mean. Marill calls the minimization of the SDA the MSDA principle.

The input line drawing is specified as a set of points (vertices) and lines; each point is represented by an  $(x, y)$  coordinate pair, and each line is represented by an integer pair corresponding to the sequence numbers of the two points it joins. The representation of the recovered 3D object involves supplying a third ( $z$ ) coordinate for each of the originally specified points. This is what we call the orthographic extension of the line drawing.<sup>3</sup> It is actually a wire frame rather than a solid object.

To evaluate the objective function for a given proposed solution, every pair of lines terminating on a

point (as defined in the input specification) is considered to form a separate angle. Thus, if five lines terminate on the same point, every potential 3D solution contains ten angles at this point that contribute to the objective function. Note that the intersection, between two lines that happen to cross at *intermediate* points of their extent in the line drawing, is not treated as a vertex, and does not contribute to the objective function (even if the lines were to lie in the same plane in the 3D reconstruction). Similarly, two *distinct* vertexes can have the same  $(x, y)$  coordinates in the line drawing, but then the line segments terminating on the distinct vertexes do not interact to form angles (even if the vertexes coincide in the 3D reconstruction).

Thus, given a line drawing with  $n$  vertexes, each possible orthographic extension is represented as a  $z$  vector having  $n$  components; the corresponding angles and SDA are computed to evaluate the proposed solution. Marill uses a descent technique to search for a best answer, recognizing that this is simply a heuristic and that this approach will find only a single local minimum of his objective function. The input object has all of its  $z$  values initially set to zero; that is, it is a flat object lying in the  $(x, y)$  plane. At each stage in the search, the SDA of the current  $z$  vector is computed and the program then looks at the children of the current vector. These  $2n$  children are all of the vectors one step size away from the current vector, and are formed by both adding and subtracting a specified value ( $\Delta z$ ) to each of the  $n$  components in the current  $z$  vector. The value of the SDA is computed for each of these  $2n$  children, and the child with the minimum SDA is selected as the new current vector. This process is repeated until no improvement in the SDA is obtained, and the resulting  $z$  vector is returned as the solution for the first of three rounds of descent. Each additional round uses a smaller  $\Delta z$  and begins with the result of the preceding round. Marill experimentally found effective values of  $\Delta z$  for his three rounds to be 1, 0.5, and 0.1.

Figure 2 shows a line drawing, its internal representation as described above, and the reconstructions using Marill's algorithm and the algorithm we describe in section 3.

In the top left window of the figure is the input line drawing (with the vertexes numbered for reference by the written representation below). The four windows on the top right show two views of Marill's reconstruction and two views of our reconstruction. In the middle of the figure is a table showing the internal represen-


tation of the input line drawing. In the first row are the  $(x, y)$  coordinates of the vertexes, in the order shown on the drawing.<sup>4</sup> In the second row are the integer pairs representing the lines in the drawing. In the third row are the sequences of vertexes corresponding to the planar faces derived according to the rules of appendix A (see section 3). The reconstructions are discussed in section 3.3.

### 2.1 Marill's Examples

Marill described the application of his algorithm to examples A through I of figure 1. We categorize these examples along the following dimensions (based on the appearance of the input drawing and on the characteristics of the recovered 3D object):

- a. —Three-dimensional [A B D E F G H I]  
—Flat [C]
- b. —Blocks world (planar-faced solids with occluded edges not rendered) [B H I]  
—Origami (planar-faced, possibly hollow) [C F]  
—Wire frame of blocks world object (all edges of a blocks world object are given, and additional lines between vertexes of a planar face may be added) [A D G]  
—Restricted wire frame (every closed circuit of lines, without interior lines in the given input representation, corresponds to a planar face) [E]  
—Nonplanar wire frame (none of the above)
- c. —Symmetric [A B C E G H]  
—Asymmetric [D F I]
- d. —All angles (approximately) equal [A B E F H]  
—A few distinct but mostly repeated angles (C G I)  
—Mostly unequal angles [D]

For the purposes of our discussion, we use Marill's categorization and augment it with our own subjective evaluation where we disagree or need to add additional attributes to those Marill provides. It is important to remember that Marill always returns a wire frame as his solution, regardless of the categorization of the object. Thus, we would call the wire frame of a blocks world object a correct solution if it was a geometrically correct representation of the 3D geometry of the edges of the psychologically plausible blocks world object whose orthographic projection corresponded to the input line drawing, even though the wire frame does not provide an explicit representation of the grouping of lines into faces, and so forth.



Points	(1.97 -1.00)	(1.32 -1.75)	(0.67 -1.75)	(0.68 -1.00)	(1.34 -0.25)	(1.98 -0.25)									
	(-0.68 1.00)	(-1.34 0.25)	(-1.98 0.25)	(-1.97 1.00)	(-1.32 1.75)	(-0.67 1.75)									
Lines	(0 1)	(1 2)	(2 3)	(3 4)	(4 5)	(5 0)	(6 7)	(7 8)	(8 9)	(9 10)	(10 11)	(11 6)	(0 6)	(1 7)	
	(2 8)	(3 9)	(4 10)	(5 11)											
Faces	(0 1 7 6)(1 2 8 7)(2 3 9 8)(3 4 10 9)(4 5 11 10)(5 0 6 11)(0 1 2 3 4 5)(6 7 8 9 10 11)														

	Zs	Lengths	Angles (Mean / Range)	SDA <sup>2</sup>	DP
Original Object	0.00 0.10 0.87 1.53 1.43 0.66 -2.23 -2.12 -1.36 -0.69 -0.80 -1.57	1.0 to 4.0	100.0 90.0 to 120.0	0.060923	0.000000
Marill's Reconstruction	0.00 0.46 -2.15 -1.48 -2.19 0.72 -0.37 0.33 -2.61 -1.92 -2.36 0.31	1.0 to 3.4	84.0 47.5 to 111.2	0.110660	0.044710
Our Reconstruction	0.00 0.12 0.96 1.66 1.55 0.71 -1.99 -1.87 -1.04 -0.35 -0.47 -1.31	1.0 to 3.9	100.0 88.6 to 122.6	0.061289	0.000000


Fig. 2. Example J. This line drawing was created by orthographically projecting a specific 3D wire frame object. In this case, the object was a regular hexagonal prism. Although arbitrary line drawings can be used as input to the reconstruction algorithms described in this article (with greater or lesser success in reconstruction), all of the examples introduced here were created by starting with specific 3D objects. The panels in the upper right show two views of the object reconstructed by Marill's algorithm. The first view is of the object rotated about the vertical axis by 30 degrees, and the second is of the object rotated about the horizontal axis by 90 degrees. The two panels in the lower right show two views of the object reconstructed by our algorithm. The table below this is the internal representation of the line drawing used by the reconstruction algorithms. Note that intersections such as those between lines (1 7) and (2 3) are not represented. Marill's algorithm uses only the first two components of this representation. The third component (faces) is derived from the line drawing using the algorithm described in section 3.1. The table at the bottom shows the results of the reconstructions in written form.

Examples A, B, E, F, and H can all be visualized as approximately equiangular three-dimensional objects. That is each of the objects has an equiangular 3D wire frame as a psychologically plausible solution. Since these equiangular solutions exactly satisfy Marill's minimum standard deviation of angles (MSDA) criterion, it is obvious why Marill's objective function should prefer what we accept as the correct solutions in these cases. In the other four cases, supposedly representative examples of the ability of Marill's algorithm to deal with complicated structures having unequal angles, reasonably correct solutions are also recovered, and it is this performance we wish to understand.

### 2.2 The Performance of the MSDA Principle

Given its overall simplicity, it would be quite remarkable if the MSDA principle generally converged to a psychologically plausible reconstruction. Unfortunately, it is rather easy to find examples where this is not the case, contrary to Marill's implied competence for the principle.

Examples J through N of figure 1 are line drawings for which Marill's algorithm converged to solutions that are clearly psychologically implausible, even though these drawings are not significantly more complicated or more asymmetric than the examples that Marill used



Points	(0.96 -0.27) (0.24 -0.89) (-0.72 -0.61) (-0.96 0.27) (-0.24 0.89) (0.72 0.61)
Lines	(0 1) (1 2) (2 3) (3 4) (4 5) (5 0)
Faces	(0 1 2 3 4 5)

	$Z_s$	Lengths	Angles (Mean / Range)	$SDA^2$	$DP$
Original Object	0.00 0.32 0.24 -0.15 -0.48 -0.40	1.0 to 1.0	120.0 120.0 to 120.0	0.000000	0.000000
Marill's Reconstruction	0.00 0.22 -0.12 0.00 0.22 -0.12	0.9 to 1.1	116.2 116.2 to 116.2	0.000000	0.030363
Our Reconstruction	0.00 0.34 0.28 -0.11 -0.43 -0.38	1.0 to 1.0	120.0 119.6 to 120.4	0.000029	0.000000

Fig. 3. Example K.


(figures 2, 3, 4, 5, and 6 illustrate both Marill's reconstructions and our reconstructions, as described in section 3). In Examples J and K it would appear that the fault could lie with Marill's use of a descent algorithm because the SDA of the psychologically plausible answer is less than or equal to the SDA for the solution Marill actually obtains. Thus, one can argue that a more competent global search strategy could have found the psychologically plausible answer using the same objective function. However, Examples L, M, and N are line drawings for which the SDA of Marill's solution is significantly lower than that of the psychologically plausible solution. Thus, the MSDA principle is clearly not adequate to *reliably* handle even simple line drawings.

Before discussing ways of augmenting the MSDA principle to obtain a more competent principle and algorithm, we attempt to explain the performance of

MSDA for line drawings depicting objects that are not equiangular.

### 2.3 Evaluating the Performance of the MSDA Principle

It is not immediately obvious why the MSDA principle should prefer a psychologically plausible answer if the object depicted in the line drawing contains two or more significantly different angles (e.g., C, D, G, I, and J). Marill offers no explanation for this phenomenon, and thus no way to judge the conditions under which his algorithm should be expected to succeed or fail. In this section we provide a partial explanation for cases (such as C, G, J, K, and L) that have critically important attributes—the psychologically plausible reconstruction is a 3D *planar-faced* object whose faces are either equiangular or form "complete-star" configurations (see appendix B).



Points	(0.81 -0.19)	(-0.02 -0.96)	(-0.82 -0.41)	(-0.49 0.71)	(0.52 0.85)	(3.96 -0.27)					
	(3.24 -0.89)	(2.28 -0.61)	(2.04 0.27)	(2.76 0.89)	(3.72 0.61)						
Lines	(0 1)	(1 2)	(2 3)	(3 4)	(4 0)	(5 6)	(6 7)	(7 8)	(8 9)	(9 10)	(10 5)
Faces	(0 1 2 3 4)(5 6 7 8 9 10)										

	Zs	Lengths	Angles (Mean / Range)	SDA <sup>2</sup>	DP
Original Object	0.00 0.29 0.94 1.06 0.47 0.47 0.15 0.23 0.63 0.95 0.87	1.0 to 1.2	114.5 108.0 to 120.0	0.010876	0.000000
Marill's Reconstruction	0.00 0.30 0.93 1.04 0.46 0.00 0.31 -0.24 0.00 -0.31 0.24	0.9 to 1.2	108.0 107.7 to 108.2	0.000005	0.165552
Our Reconstruction	0.00 0.00 0.00 0.00 0.00 0.08 0.41 0.35 -0.04 -0.36 -0.30	1.0 to 1.2	114.5 97.9 to 120.4	0.018157	0.000000

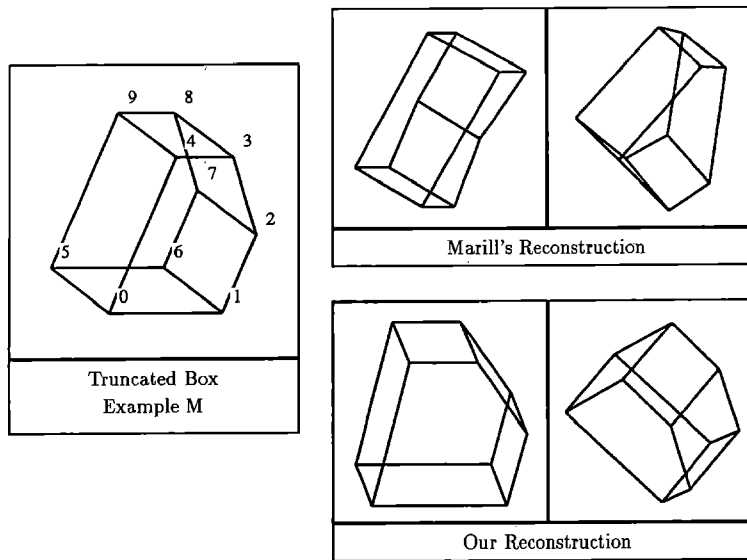
Fig. 4. Example L. Note that Marill's unacceptable reconstruction has an SDA that is significantly lower than that of the psychologically plausible original object. Thus, the MSDA principle itself has failed in this instance.

To establish the role played by the above geometric attributes, we define the *planar orthographic extension of a simple closed 2D circuit* in a line drawing to be any orthographic extension for which the corresponding 3D contour is planar. If a line drawing contains more than one simple closed 2D circuit, then a *planar orthographic extension of the entire line drawing* exists if we can cover the line drawing with a set of simple closed 2D circuits such that (a) every angle in the drawing is included in at least one circuit, and (b) each circuit projects to a 3D planar contour.<sup>5</sup>

In appendixes B, C, and D, we provide a number of theorems that are pertinent to understanding the effectiveness of the MSDA principle applied to planar orthographic extensions. The main theorem, appendix D, asserts that solutions with certain symmetries correspond to the global minimum of the SDA over all planar orthographic extensions (the specific symmetry condition we examine is that all faces must either be

equiangular or form complete-star configurations).

Consequently, if there were some way to consider as possible solutions only the *planar* orthographic extensions of a line drawing (such as the psychologically plausible solutions for examples A, B, C, G, J, K, and L), these solutions would be global minima of the SDA because of the angular symmetry they exhibit. We show in example L that Marill's algorithm is not constrained to search only for planar solutions; while it will also find solutions with nonplanar faces that have lower SDAs than the planar solutions, there is still the possibility that MSDA shows at least a weak inherent preference for planarity. While we cannot completely rule out this possibility, it appears that the geometric constraints inherent in the specific examples Marill selected, rather than MSDA itself, are largely responsible for finding planar-faced solutions. Specifically, triangles in the line drawing will always produce planar faces in the orthographic extension, and as we prove



Points	(0.15 -0.06) (0.80 -0.06) (0.99 0.38) (0.86 0.81) (0.54 0.81) (-0.18 0.19) (0.46 0.19)
	(0.66 0.63) (0.53 1.06) (0.20 1.06)
Lines	(0 1) (1 2) (2 3) (3 4) (4 0) (5 6) (6 7) (7 8) (8 9) (9 5) (0 5) (1 6) (2 7) (3 8) (4 9)
Faces	(0 5 9 4)(1 0 5 6)(2 1 6 7)(3 2 7 8)(4 0 5 9)(4 3 8 9)(0 1 2 3 4)(5 6 7 8 9)

	Zs	Lengths	Angles (Mean / Range)	SDA <sup>2</sup>	DP
Original Object	0.00 0.77 0.61 0.06 -0.32 0.28 1.04 0.88 0.34 -0.04	0.5 to 1.0	96.0 90.0 to 135.0	0.071281	0.000000
Marill's Reconstruction	0.00 0.68 0.67 -0.17 -0.37 0.31 0.97 0.93 0.18 0.01	0.4 to 1.0	95.4 73.5 to 125.0	0.047822	0.004897
Our Reconstruction	0.00 0.57 0.49 0.10 -0.20 0.37 0.94 0.86 0.46 0.16	0.4 to 1.0	96.0 80.7 to 132.2	0.059677	0.000000

Fig. 5. Example M. Note that our reconstruction has a slightly lower SDA than that of the original object, indicating the preference of our algorithm for equiangular faces.

in appendix B, a closed four-sided polygonal space curve with 90-degree angles at each vertex will always be a planar configuration. Since in Marill's examples listed above, all the faces satisfy these two geometric conditions, we see why both the desired planarity and symmetry are present in the computed solutions.<sup>6</sup>

Marill offers only two examples (D and I) that are not clear instances of the above analysis (all angles equal, or symmetric planar faces). His solution for example I is at least questionable since it does not recover the wire frame of a polyhedral solid (our algorithm finds such a solution; there is a further discussion of this subject in sections 3.3 and 4). However, this solution has almost all of its angles equal to 90 degrees, and so it needs no further explanation if we accept it as correct.

Marill's solution to the asymmetric drawing of example D looks very reasonable; it has all its angles fairly well distributed between 40 and 70 degrees, and we have not found a more symmetric (equiangular) ortho-


graphic extension for this line drawing. However, because the input line drawing is a completely connected set of triangular faces, all solutions are constrained to have planar faces. Thus, a large range of psychologically plausible objects is accessible to any reasonable algorithm.

In summary, there is an understandable reason why Marill's MSDA principle will sometimes tend to select planar symmetric 3D wire frames when a purely equiangular solution is not possible. But we also see that MSDA will make unacceptable errors, even in simple cases, because it is not constrained to prefer solutions with planar faces unless the geometry of the line drawing itself forces planarity.

### 3 Our Planarity Enforcing MSDA Algorithm

What's missing in the MSDA principle is a means for enforcing the planarity of specified faces. There are two





Points	(-0.58 0.24) (0.95 1.36) (1.50 1.04) (-0.02 -0.08) (0.30 2.89) (0.86 2.56)
Lines	(0 1) (1 2) (2 3) (3 5) (5 4) (4 0)
Faces	(0 1 2 3)(0 4 5 3)

	$Z_s$	Lengths	Angles (Mean / Range)	$SDA^2$	$DP$
Original Object	0.00 0.64 -0.12 -0.77 -0.47 -1.24	1.0 to 2.8	75.0 45.0 to 90.0	<b>0.137078</b>	<b>0.000000</b>
Marill's Reconstruction	0.00 1.21 -1.17 -0.41 -1.18 1.26	2.0 to 3.3	63.8 62.7 to 65.0	<b>0.000293</b>	<b>0.196270</b>
Our Reconstruction	0.00 1.93 1.69 -0.21 -2.21 -2.40	0.7 to 3.6	89.8 88.8 to 91.3	<b>0.000291</b>	<b>0.000000</b>

Fig. 6. Example N. The SDA of Marill's unacceptable reconstruction is again significantly lower than that of the psychologically plausible original object.

parts to this problem: (1) finding those faces in the line drawing that should be planar in the 3D reconstruction, (2) and enforcing the planarity of these faces during, or at least by the end of, the optimization process.

### 3.1 Finding Planar Faces

The following algorithm for finding the planar faces is based on a set of psychological assumptions presented in appendix A. The requirements of items 3, 4, and 5 from appendix A have been composed into the following algorithm. (In the following discussion, we define a face in the line drawing to be a sequence of vertexes.)

First, all simple (nonself-intersecting) closed circuits containing more than three lines are found. (Triangles are necessarily planar, so they need not be considered.)

Those circuits that are either: (1) completely empty of both lines and vertexes (such as the faces of example B); or (2) both convex (in the line drawing) and free of internal circuits (such as all the faces of example J) are considered to be planar faces of the wire frame; call this initial set  $\mathcal{P}_0$ . A circuit is defined to be an internal circuit to a convex circuit if: (1) all of its vertexes lie within the convex circuit; and (2) it terminates in two nonadjacent vertexes of the convex circuit.

Added to  $\mathcal{P}_0$  are those circuits, defined by the following algorithm, that are not subsets of any circuit in  $\mathcal{P}_0$ . First, all triples of consecutive lines such that the first and third lines are parallel are found (the two planar faces of example N fall into this category, as do the "table legs" of example E). Then, if possible, each triple of lines is extended with additional consecutive lines such that all even-numbered lines are parallel to

each other and all odd-numbered lines are parallel to each other. An example of a closed circuit found this way is the side of the staircase facing the viewer in example H; the side of the staircase opposite the viewer is an example of an open circuit found using this same rule.

Finally, pairs of parallel lines lie on planar faces in general position, so the four vertexes of the pair of lines are defined to form a planar face (whether or not the vertexes are connected by lines in the line drawing). If the pair of lines are not already a subset of a previously found planar face, these are added to  $\mathcal{P}_0$ .<sup>7</sup>

The above procedure is remarkably robust in dealing with unconstrained line drawings. For example, we have yet to find a case where this procedure proposes a psychologically implausible planar face (it even found faces in our test cases that we had not originally recognized as being planar—such as the back side of the staircase in example H). However, it will sometimes miss finding a concave planar face leaving the 3D model underconstrained, and this can result in the reconstruction of a psychologically implausible 3D wire frame. If we know that the line-drawings to be processed are restricted to the projections of blocks world objects with all planar intersections included in the drawing (i.e., no hidden lines removed), then we can be assured that no faces are missing by (omitting some details here) first employing the above procedure, next removing all lines' edges from the drawing that are assigned to two faces, and then repeating this whole process on the reduced line drawing until all the edges have been assigned to exactly two faces (there are some special-position configurations in which three or more faces have a single edge in common that we presently do not deal with). For this more constrained universe of line-drawings where we correctly and completely identify all the planar faces, we have yet to encounter a case where our algorithm produces a psychologically implausible 3D model.

### 3.2 Enforcing Planarity

The second requirement, enforcing planarity, is accomplished by adding a term to the objective function that is zero when all the designated planar faces are actually planar, and increases in value as the faces deviate from planarity (call this term DP). The new objective function,  $E(\lambda)$ , is a linear combination of the previously defined SDA term and the new DP term.<sup>8</sup>

$$E(\lambda) = \lambda SDA^2 + (1 - \lambda)DP$$

Note that minimizing  $E(\lambda)$  favors planar faces, but strict planarity is not necessarily assured. This is not quite what we would like in the ideal case. Ideally, we would like to find the orthographic extension of the line drawing with the lowest SDA that has exactly planar faces (i.e., for which  $DP = 0$ ).<sup>9</sup> To achieve this, we use a continuation method (Leclerc 1989; Witkin et al 1987), which is a sequence of descent steps applied to  $E(\lambda)$ , for decreasing values of  $\lambda$ . The sequence begins with the initial condition that Marill suggests ( $z = 0$  for all points) and with some initial  $\lambda_0 \leq 1$ . Then,  $\lambda$  is decreased by a given amount and the descent algorithm is applied anew, starting at the solution found for the previous value of  $\lambda$ . This is repeated until  $\lambda$  is sufficiently close to zero so that no additional changes occur with further reductions in  $\lambda$ .

Why not simply start with  $\lambda$  close to zero in the first place? The reason is that when  $\lambda$  is sufficiently close to zero, the local minima of  $E(\lambda)$  are determined only by the planarity component. Thus, simply starting with  $\lambda$  close to zero would not allow us to find solutions with low SDAs (in fact, when  $\lambda = 0$ , the original line drawing, which is planar, is a local minimum of  $E(\lambda)$ ). Although we cannot affect the shape of  $E(\lambda)$  when  $\lambda$  is small, we can choose the starting point for the descent algorithm. Thus, the purpose of the continuation method is to choose a sequence of starting points that are first strongly influenced by the SDA term, but which eventually become dominated by the DP term. The method is not guaranteed to find a global minimum of the objective function, but has yielded excellent answers for all the examples discussed in this paper.

We define the deviation from planarity term, DP, as the sum of terms  $DP_i$ , where  $DP_i$  is zero when face  $f_i$  is planar, and increase as the face deviates from planarity. We have found two useful definitions of the  $DP_i$ . The first is a strong planarity term that will not allow a face to fold from one planar configuration to another planar configuration, but applies only to convex faces. To see how a face can fold from one planar configuration to another one within the context of the optimization we are performing, consider a line drawing of a square. When all of the  $z$  values of the vertexes are zero, the face is planar. By letting the  $z$  values of the first and third vertexes become arbitrarily large, the face "folds" into a configuration that, in the limit, is also planar. In order to detect and avoid this folding whenever possible, we define  $DP_i$  to be the following function ( $DP_i$  is based on item 6 in appendix B):

Let  $n$  be the number of sides in the face, and  $\alpha_j$  be the angle at the  $j^{\text{th}}$  vertex. Then,

$$DP1 = \left[ (n - 2)\pi - \sum_j \alpha_j \right]^2$$

A weaker measure of planarity,  $DP2$ , applicable to all faces, is based on the observation that the normals defined by pairs of consecutive pairs of lines should lie in the same direction (this is analogous to the notion of torsion for a curve):

$$DP2 = \sum_j \left[ 1 - \left( \frac{(l_{j-1} \times l_j) \cdot (l_j \times l_{j+1})}{\|l_{j-1} \times l_j\| \|l_j \times l_{j+1}\|} \right) \right]^2$$

where  $l_j$  is the  $j^{\text{th}}$  line of planar face  $f_i$  and  $j - 1$  and  $j + 1$  refer to the previous and next lines in the face, respectively (i.e., the subscripts are taken modulo the number of lines in the face).

The combined  $DP$  term is the sum of: (1) the sum of  $DP1$  over all convex faces, and (2) the sum of  $DP2$  over all nonconvex faces divided by the number of angles in all of the nonconvex faces.

### 3.3 Results

Figures 2 through 6 illustrate the results of our *planarity enforcing MSDA* algorithm, and allows one to compare them with both Marill's reconstructions and the original 3D objects that were used to generate the line drawings. The "original 3D objects" presented in our figures are the psychologically plausible solutions that we expect the program to recover. We started with actual 3D wire frames, rather than arbitrary line drawings as an experimental expedient, since most random line drawings will not induce the perception of a 3D configuration in human subjects.

The reconstructions are illustrated both graphically (as two views in the upper third of each figure) and in tabular form in the lower third. The first column of the table lists the  $z$  coordinates of each object, the second column is the range of lengths of the lines of each object, the third column is the mean and range of the angles formed by all line pairs meeting at a common vertex, the fourth column is the standard deviation of angles (SDA) of each object, and the fifth column is the deviation from planarity (DP) of each object. To simplify the comparison of the results, the recovered  $z$  coordinates have been normalized so that the first

point always has  $z = 0$ , and the second coordinate is always positive (this normalization procedure has no effect on the objective function).

We also applied our algorithm to examples A through I from Marill's paper. Since his algorithm produced approximately planar-faced solutions by itself in all cases but example I, it isn't surprising that our algorithm produced solutions almost identical to his. The greatest deviation from his result was for example I, because Marill's algorithm recovered a significantly nonplanar face for the leftmost face of the line drawing.

In all of the examples, the  $\Delta z$ s we used for Marill's algorithm (both as a stand-alone algorithm and within the continuation method) were 0.125, 0.0625, 0.03125, 0.015, and 0.007. We used a smaller initial  $\Delta z$  than Marill suggests because the larger one often forced the algorithm out of the valley of attraction of the current local minimum. Decreasing  $\Delta z$  by a factor of two generally allowed the algorithm to run in the fewest number of iterations. Using a smaller final  $\Delta z$  allowed the algorithm to produce significantly more accurate solutions. In the continuation method,  $\lambda$  was started at 0.25, and was decreased by a factor of two a total of ten times.

Example J (figure 2) illustrates Marill's reconstruction for a line drawing of a rectangular hexagonal prism. This reconstruction not only appears psychologically implausible from these two views, but, as we discuss in the following section, the reconstructed object does not appear rigid when rotated in real time. It would appear that at least part of the reason for this result is that the recovered faces are clearly nonplanar, as shown by the value of  $DP$  in the table. The reconstruction obtained by using the planarity enforcing MSDA algorithm is almost identical to the original hexagonal prism.

In example K, we see that the MSDA principle is ambiguous for simple line drawings. Marill's reconstruction takes the line drawing of a planar hexagonal plate ( $SDA = 0.0$ ) and reconstructs a nonplanar object, also with  $SDA = 0.0$ . By enforcing planarity, however, our reconstruction is quite close to the original hexagonal plate.

In examples L and N, we see further evidence that the MSDA principle by itself is inadequate for even simple line drawings. In both examples, Marill's reconstruction has a significantly lower SDA than the original object, and we consider both of these reconstructions to be psychologically implausible. Our reconstruction of example L is quite close to the original object, modulo an additive constant and flip of the  $z$  coordinates of the second object (which is invisible to the objective

function). Example N is a fairly ambiguous figure, and our reconstruction favored a “hinge” with all angles close to 90 degrees (the original object had a “hinge-angle” of 45 degrees). Because of the ambiguity of the figure, there exists a family of reconstructions that we consider psychologically plausible, including ours.

Example M shows the reconstruction of a figure for which some of the planar faces are not equiangular. Again, because some of the faces had more than four sides, Marill’s algorithm failed to recover a psychologically plausible object. Our reconstruction is reasonably good, but it did adjust the right angles in the large face by as much as 13 degrees in order to make the angles in that face closer to being equal. Nonetheless, we consider the reconstruction to be psychologically plausible.

### 3.4 Stability and Robustness of the Planarity Enforcing MSDA Algorithm

We have examined the stability and robustness of our algorithm in two ways. The first was to examine the behavior of the algorithm applied to different projections of the same 3D objects, but always using the same initial conditions for the optimization, namely  $z = 0$  for all vertexes. The second was to examine the behavior of the algorithm for different initial conditions.

We ran the planarity enforcing MSDA algorithm on at least 32 randomly chosen projections of the 3D objects used to create the line drawings of examples A through N.<sup>10</sup> For virtually every projection of each of these objects, the algorithm reconstructed the object as well as it did for the original projection. For example, figure 7

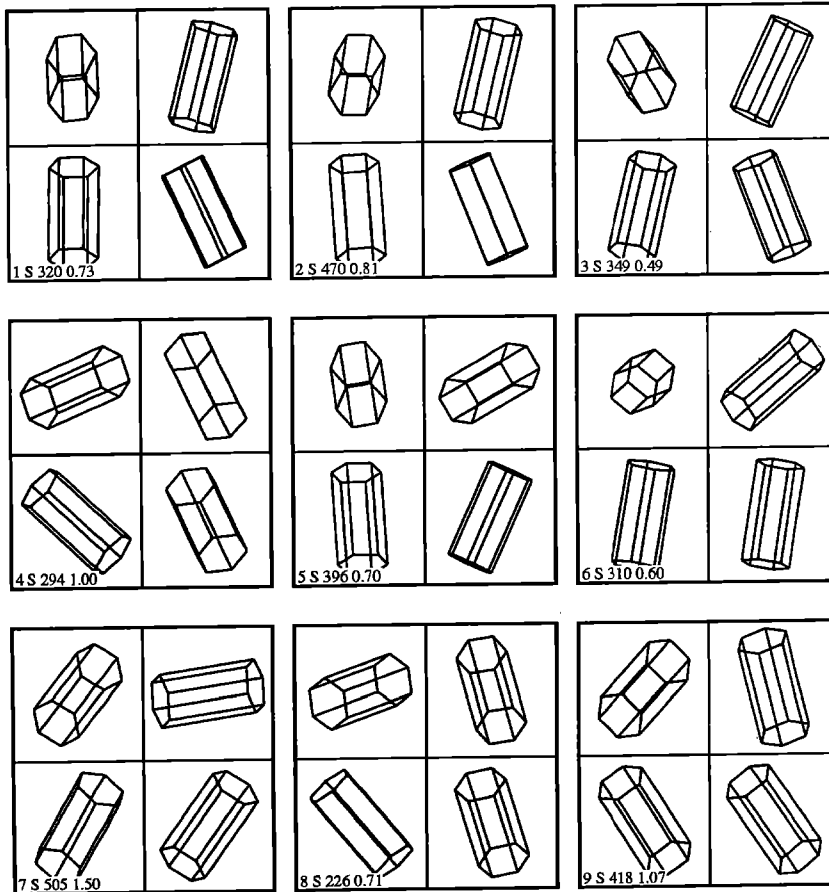


Fig. 7. Nine projections of the hexagonal prism, and our corresponding reconstructions. The projections used as original line drawings are shown in the lower left-hand corner of each group of four. The original line drawing is annotated by: (1) the projection number; (2) the letter S when the planar faces found for that line drawing were the same as for the original projection, and D otherwise; (3) the number of iterations required for convergence; and (4) the largest difference between corresponding angles in the reconstruction and the original object, in degrees. The other three line drawings are three views of the reconstruction.

shows nine projections and the corresponding reconstructions for the hexagonal prism (Marill's algorithm failed for all of these projections). An example of a near failure is shown in figure 8, where the eighth projection of the staircase is almost in special position, producing the largest error, and using the greatest number of iterations. In fact, when the rule adding all pairs of parallel lines as planar faces is removed, the algorithm leaves the  $z$  values virtually unchanged from their initial values (not illustrated here). In summary, in approximately 500 trials, either the planarity enforcing MSDA algorithm correctly reconstructed the original object, or it left the line drawing as an "uninterpreted" flat object.

By comparison, the MSDA algorithm is relatively unstable, even for the line drawings one might expect it to get right. For example, figure 9 shows nine projections and the corresponding reconstructions using the MSDA algorithm, for a cube in which all of the angles

should be exactly equal. Note that projections 1 and 9 produce psychologically implausible reconstructions.

In a second set of experiments, we used a random-number generator to provide twenty sets of initial  $z$ s in the range  $-1$  to  $1$  for examples A through N.<sup>11</sup> With the exception of example D, which was always correctly reconstructed, the MSDA algorithm failed to converge to a psychologically plausible solution in at least four of the twenty trials on each of the other line drawings, and produced an average of ten failures per line drawing. In other words, the SDA term by itself has many local minima that descent algorithms will fall into.

On the other hand, the planarity enforcing MSDA algorithm succeeded in converging to a psychologically plausible solution in all trials but one (it failed in one trial of example N, the hinge.)<sup>12</sup> This extremely robust performance was somewhat unexpected. We believed that the initial condition,  $z = 0$  for all vertexes was an important component of the continuation

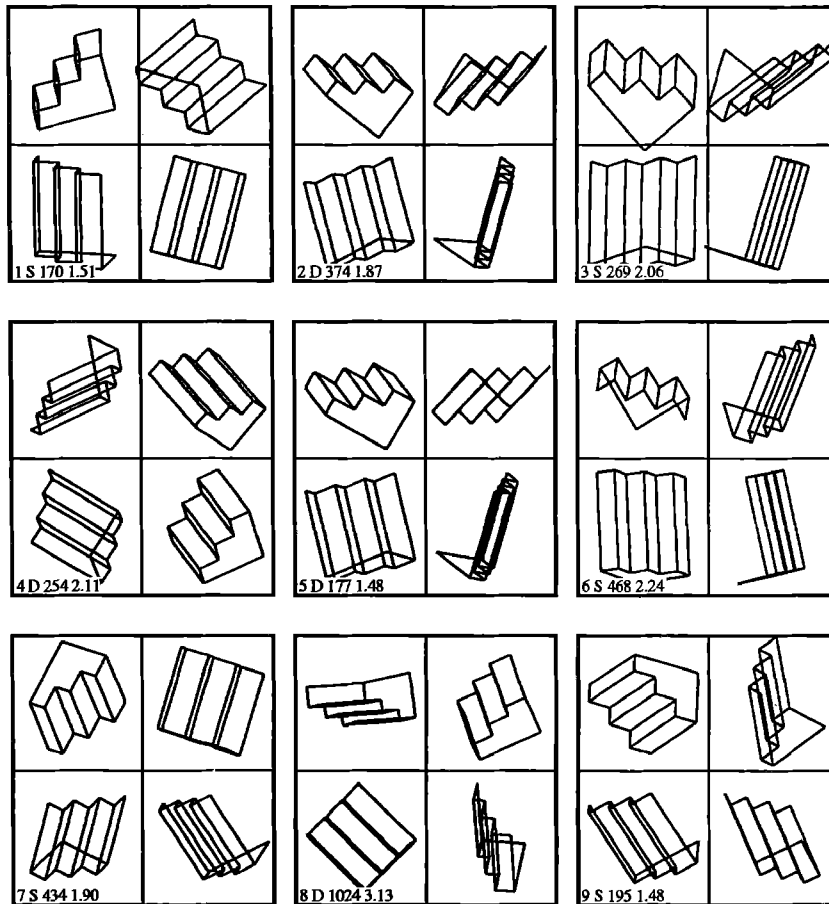


Fig. 8. Nine projections of the staircase, and our corresponding reconstructions. Note that the eighth projection is very nearly in special position, with many vertexes and lines overlapping in the line drawing. The continuation method had the largest error and used the greatest number of iterations for this case. When the rule adding all pairs of parallel lines as planar faces is removed, the continuation method prefers the original line drawing (all  $z$ s constant) as the interpretation, which is certainly psychologically plausible.

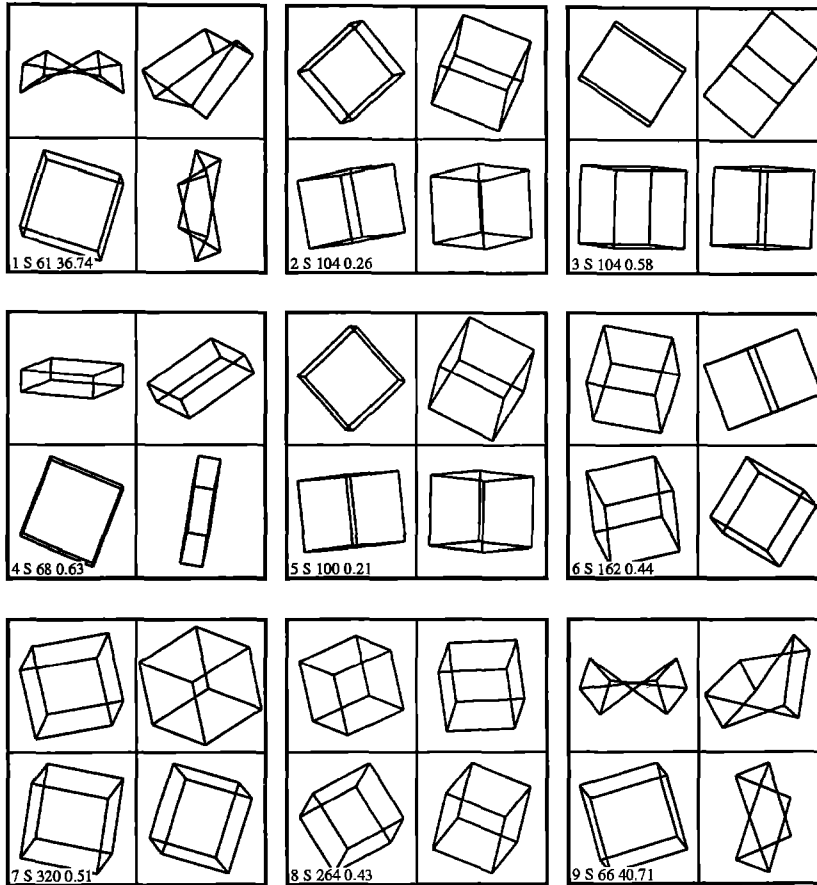


Fig. 9. Nine projections of the cube, and the corresponding reconstructions using Marill's algorithm. Note that projections 1 and 9 produced psychologically implausible reconstructions.

method. However, it would appear from the results of these experiments that the imposition of the planarity term in the continuation method severely curtails, or eliminates, psychologically implausible minima. One might conjecture that, for most line drawings, there is one<sup>13</sup> (or perhaps a very few) psychologically plausible local minima in the SDA when the  $z$ s are constrained to a *planar* orthographic extension.

### 3.5 Reconstruction Time

The specific descent algorithm defined by Marill, and described here, has the nice property that it's easy to describe and easy to implement, no matter what the objective function may be; however, it is typically quite inefficient. One of the better descent algorithms is the conjugate gradient algorithm. To estimate achievable

run times, we implemented the conjugate gradient algorithm described in *Numerical Recipes* (Press et al. 1986). The algorithm requires an objective function (in this case,  $E(\lambda)$ ) and the gradient of the objective function (in this case, a function that returns a vector whose  $i^{\text{th}}$  element is the partial derivative of  $E(\lambda)$  with respect to  $z_i$ ). Analytically deriving the gradient of  $E(\lambda)$  is rather painful, so instead we used a simple numerical approximation; this involves evaluating the objective function for each vertex, which is expensive. A more efficient implementation that only recomputes those components of the objective function that change when a given vertex changes could reduce the following run times by a factor of four or better.

Table 1 gives the number of iterations/run time (in seconds) for three example line drawings. These experiments were run on a Symbolics 3645, so we would expect about a factor of ten improvement if algorithms

were implemented in C on a modern workstation such as a SUN SPARC-2 (according to some simple benchmarks that we ran). The last column gives the expected run time for an optimized conjugate gradient algorithm running on a SPARC-2.

Table 1. Number of iterations/run time for three examples.

Example	Original Descent on Symbolics	Conjugate Gradient on Symbolics	Conjugate Gradient on SPARC-2
Cube	187/199	15/15	15/0.375
Tetrahedron	46/9	14/16	14/0.4
Hexagon prism	406/1306	33/186	33/4.65

Note that the conjugate gradient algorithm improves the run-time considerably for all but the simple tetrahedron line drawing. On a SPARC-2, the run times are such that the time required to reconstruct a line drawing is small relative to the time it would take to manually enter the drawing. That is, the run times are well within “interactive time.”

### 3.6 A Reduced Search Space Technique for Obtaining Exact Planar MSDA Reconstructions

In the planarity enforcing MSDA algorithm described in section 3.2, planarity is not strictly enforced, but rather, nonplanarity is penalized during the optimization process. This approach almost always produces faces that are very nearly planar at the end of the optimization process. There is a very efficient way to *strictly* enforce planarity during the MSDA optimization for line drawings of strictly planar-faced wire frames, described below. The problem with this approach is that if the line drawing does not actually correspond to a planar-faced wire frame, or if the line drawing is not accurate, the resulting reconstruction will typically be psychologically unacceptable—we lose the graceful degradation provided by the planarity enforcing MSDA.

The following method for strictly enforcing planarity is based on the observation that there are far fewer degrees of freedom in a planar-faced object than there are vertexes (to reemphasize, this method is only applicable to line drawings of strictly planar-faced wire frames). One way of expressing this observation is in terms of a subset of vertexes, that we call the *free vertexes*, whose  $z$  values uniquely determine the  $z$  values of all of the other *dependent* vertexes by virtue of the

planarity of certain faces. For instance, given the planar faces of the hexagonal prism of figure 2, specifying the depth of the four vertexes 0, 1, 2, and 6 uniquely determines the depth of the other vertexes: the depth of vertexes 3, 4, and 5 are determined by constraining them to lie on the same planar face as vertexes 0, 1, and 2; similarly, vertex 11 is determined by vertexes 0, 5, and 6; vertex 10 by vertexes 4, 5, and 11; and vertexes 7, 8, and 9 by vertexes 6, 10, and 11.<sup>14</sup>

Having determined the free vertexes, one can then apply the MSDA principle to the reduced search space. For the case of the simple descent algorithm, the only change to the algorithm is that only the free vertexes are directly modified during the optimization, and that the depth of all of the dependent vertexes are recomputed whenever a free vertex is modified. Applying this *method of free vertexes* to the hexagonal prism reduces the number of iterations from 406 to 39, and the run time from 1306 seconds to 47 (the run time is reduced by a greater proportion than the number of iterations because the DP term has effectively been removed from the objective function).

Thus, the advantage of using the method of free vertexes is that it reduces the search space and run times considerably—oftentimes an order of magnitude or more. The disadvantage of using this approach is that, unlike the planarity enforcing MSDA algorithm, it requires a virtually perfect line drawing of a planar-faced object to ensure that the resulting reconstruction is planar. For example, adjusting the  $(x, y)$  coordinates of even one vertex by a small amount in a line drawing such as the cube (example A), can cause the 3D wire frame to be highly nonplanar for some choices of  $z$  coordinates of the free vertexes. Consequently, the method of free vertexes can produce reconstructions that are not psychologically plausible. Nonetheless, there are certain situations in which this approach can be effective, both for special kinds of line drawings, and for line drawings that are first processed to make them precise projections of the intended 3D object.

## 4 Implications for Human Vision

Line drawings provide an effective means of communication about the geometry of 3D objects. It is a matter of some debate as to whether the interpretation of line drawings is a learned skill, or whether line drawings are isomorphic to some intermediate construction of the human visual system (HVS) in its normal processing

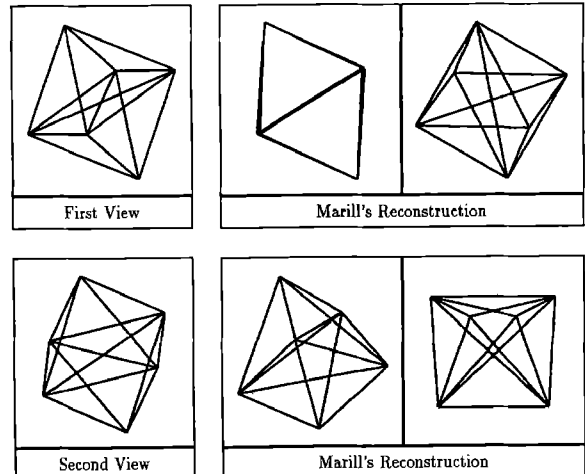
of imagery, but in either case an understanding of how humans interpret line drawings is extremely important in enabling man-machine communication with respect to images, diagrams, and spatial constructs. In this section we address two related questions arising out of the investigation described in earlier sections: (a) under what conditions is a line drawing actually given some intended 3D interpretation, and (b) under what conditions does a moving rigid (wire frame) object actually appear rigid.

Some, but not all, line drawings are perceived by human subjects as being three dimensional. What attributes of the drawing promote such an interpretation, and what are the constraints on the nature of the resulting 3D construction? Partially because human introspection is involved, this is a very difficult question to answer. For example, if the drawing is recognized as a known or previously encountered 3D object, it might be visualized this way even though it violates conditions necessary for an unfamiliar object to be perceived as being three dimensional. Gestalt psychologists have suggested that if the drawing offers a simpler construct when seen as three dimensional than when seen as being flat, it will be perceived as being three dimensional; however, an effective computational procedure to evaluate "simpler" has yet to be provided (and there is also the problem of producing the corresponding 3D construct). One might consider that minimizing angular variance is an example of a simplicity principle, but we have not yet been able to define a formal complexity metric, as was done, for example, in the work of Leclerc (1989).

It appears to be much more productive to show a human subject a candidate 3D reconstruction and ask if it corresponds to some given line drawing than it is to tabulate introspective judgments about whether objects appear to be 2D or 3D. The former approach, in fact, is how Marill presents his results to the reader. Obviously, he can not show an actual 3D reconstruction, but only a projection. If he showed the reconstructed object projected without some spatial relocation, then all we have is the original line drawing back again—and no determination can be made; Marill shows two projections of his reconstructed objects, rotated by a few degrees, for evaluation by the reader. Now we know that every orthographic extension is a *geometrically* feasible reconstruction, so on what basis does the human judge acceptability (i.e., what we have called a *psychologically* plausible reconstruction). It is easy to hypothesize a whole list of conditions that

should be met—mostly different instantiations of the idea that regularities (such as parallel lines or equal angles and lengths) observed in the line drawing are not accidental, and should be preserved in the reconstructed object; orthographic projective invariants, such as parallelism, should then also be preserved in the reprojections of the spatially relocated object. One could write computational procedures to search for such invariants, but this approach seems incompatible with the universality of the human evaluation process (e.g., none of the invariants we happened to think of may be present in the instances we are considering). A more powerful idea is to require that the computational procedure that produced the original reconstruction give the same result when applied to any of its general position reprojections—that is, a consistency criterion. This is exactly the condition that obtains when we observe a moving or rotating object to be rigid; when we see a (continuous) sequence of projections that we perceive as being isomorphic to the same geometric reconstruction, we perceive the object as being rigid.<sup>15</sup>

Applying the above ideas to an evaluation of the MSDA algorithm, we find two serious deficiencies in



*Fig. 10.* Illustration of the failure of Marill's algorithm to recover geometrically similar 3D models from two different projections of the same 3D project. The top row shows the input line drawing of the 3D object as seen from one viewpoint (similar to example G), and two views of Marill's reconstructed object. The bottom row shows the input line drawing of the same 3D object as seen from a different viewpoint, and two views of Marill's reconstructed object. The two reconstructed objects not only appear different, but are in fact significantly different geometrically, as we verified by examining their internal representation. In contrast, applying our algorithm to both of these input line drawings, as well as ten other randomly chosen views produced reconstructions with an angular error of less than thirteen degrees from the original object.






Fig. 11. The illusion of nonrigidity for a rotating wire frame with nonplanar faces. The wire frame, Marill's reconstruction of example J, is rotated about a vertical axis in the center of the object. The rotation angle is written in the lower left-hand corner of each box.

the algorithm. First, when presented with two different orthographic projections of an object, the MSDA algorithm sometimes fails to recover 3D wire frames that are even remotely similar to each other (see figure 10). Second, when we use the computer to create a rotating display of some of the reconstructions obtained with the use of the MSDA algorithm, we see what appears to be the movement of a nonrigid object (see Figure 11).

The latter observation led to a number of casual experiments to determine the factors affecting the perception of nonrigidity in displays of rotating 3D wire frames. We found that wire frames with pronounced nonplanar faces (where one would have expected a planar face from the line drawing) appear to be nonrigid. Marill's solution for example I (asymmetric solid) does appear rigid under rotation, even though the faces are slightly warped. However, his solution is very nearly planar; if we force a bit more distortion into the solution, the object then appears to deform under rotation. Thus, it would appear that strict (or at least near) planarity for the appropriate faces is a *necessary* condition for the perception of rigidity.

However, planarity by itself was not sufficient to create a perception of rigidity. For example, if one chooses random values for the free vertices of a certain line drawing (see section 3.6), one produces an object whose faces are strictly planar. However, unless the resulting figure is also a local minimum of the SDA, the resulting 3D wire frame does not appear rigid when

rotated. Similarly, the wire frames of some line drawings with all of the  $z$  coordinates set to zero appeared nonrigid when rotated (e.g., example A). Furthermore, all of the hundreds of solutions produced by the planarity enforcing MSDA algorithm that we looked at appeared rigid under rotation. Thus, we tentatively conclude that a wire frame must not only be planar to be perceived as rigid, but must satisfy additional constraints, such as being a local minimum of the SDA.

## 5 Future Work

There are a number of directions that we have begun to explore or that we plan on exploring in the near future.

The first of these, for which we have some preliminary results, is a redefinition of the objective function in which the angles are partitioned into groups that should be equiangular in 3D. This becomes necessary either when there are angles in the line drawing that are not a part of any planar face or when the angles in a planar face are not all equal in 3D (in either of these cases, the symmetric preference theorem of appendix D does not hold). An example of the first case is the hinge (figure 6), in which angles (1 0 4) and (2 3 5) are not a part of any planar face. An example of the second case is the truncated box (figure 5), in which angles (1 2 3) and (2 3 4) should be equal to each other but not equal to the other angles in planar face (0 1 2 3 4), and similarly for face (5 6 7 8 9).




Fig. 12. An illustration of the need to group angles together than should be equiangular, rather than applying the planarity enforcing MSDA principle to all angles. (a) The reconstruction using an appropriate equiangular grouping. (b) The reconstruction using the planarity enforcing MSDA principle applied to all angles.

By changing the definition of the SDA term to be the sum of the standard deviation of the angles in each equiangular group (weighted by the number of angles in that group), we have improved the reconstruction of these two objects considerably. Defining a simple, yet robust, set of rules that can automatically determine the equiangular groups for a line drawing, as we did for the planar faces of the line drawings in this paper, is still an open question. A simple rule is to group together all angles that are a part of a convex face. This is illustrated in figure 12. The reconstruction is accurate to 3 degrees, whereas using the SDA over all angles gives a relatively poor reconstruction.

A second direction that we plan on exploring is to implement a preprocessing step that would take a rough sketch and enforce various constraints in 2D, such as (1) parallelism between designated pairs of lines, or between designated lines and axes; (2) equality in length between designated lines, or between lines and fixed lengths; and so forth. The paradigm would be similar to the one for the interpretation of the line drawing, namely some set of rules would be used to determine which lines should be parallel or of equal length (with outside intervention always possible), and an optimization step would then enforce the constraints while moving as little as possible from the original line drawing. The ideal is to be able to do as much of this as possible without intervention for an interactive user.

A third direction is to explore the relationship between what we have done and previous work in understanding the 3D shape of curves, such as (Barrow & Tenenbaum 1981; Stevens 1981; Witkin 1981; Barnard

& Pentland 1983; Malik & Maydan 1989; Pentland & Kuo 1990).

An intriguing relationship between Barrow and Tenenbaum's work on single curves and our work on planar faces is as follows. They defined the problem of interpreting curves in a manner similar to the way that we and Marill did: by defining an objective function over the  $z$  coordinates of the object and minimizing that objective function using a descent algorithm.<sup>16</sup> Their objective function was the integral of the change in curvature squared plus the torsion squared. Thus, an ideal curve for their objective function is a planar circle, since both terms in the integral are then zero everywhere (when the end-points are removed from the integral, the arc of a planar circle is also an ideal curve for their objective function). Analogously, one of the ideal curves for our definition is a regular planar polygon (or an arc of a regular planar polygon), since then both the SDA and DP are zero. Thus, the similarities are that the SDA plays a role similar to the integral of squared change in curvature, and the DP plays a role similar to the integral of squared torsion. Some of the differences are that both the SDA and DP1 terms are global measures of symmetry and planarity, while the curvilinear measures are integrals of local measures. A second difference is that the SDA is also zero for some nonregular and even nonconvex polygons.

Pentland and Kuo (199) applied Barrow and Tenenbaum's idea to distinctly nonplanar curves and surfaces by leaving out the torsion component. It is somewhat surprising that this worked since both Barrow and Tenenbaum's and our own experience indicate that

planarity is a key ingredient in making the optimization approach work. We will explore this question in the near future.

Finally, we would like to find some computational-effective procedure for using the rigidity under rotation criterion in the 3D recovery process, rather than as a final check on proposed solutions.

## 6 Discussion

Traditional blocks-world problems are mathematical in nature, they deal with issues of existence and consistency based strictly on geometric considerations; they make no reference to what people actually see. The problem defined by Marill is psychological; since every line drawing has an infinite number of mathematically valid orthographic extensions and no invalid ones, on what basis does the HVS select a particular extension as being psychologically acceptable? Marill proposed an intriguingly simple criterion for duplicating human preference, but we have shown that, while it often produces an acceptable answer, it is unreliable even in very simple situations.

Marill's work has similarities to the Huffman-Clowes-Waltz approach that focused on how polyhedral vertexes can appear in a line drawing and, hence, the constraints such vertexes impose on the implied 3D model; Marill considers only the constraints implied by line intersections at specified vertexes in the line drawing. Mackworth, Kanade, and Sugihara found it necessary to introduce constraints based on the explicit assignment of vertexes to planar faces. We show here the need for introducing a similar explicit requirement for planarity (actually, in the context of optimizing an objective function, our constraint is soft in that it can be violated). However, in our case, the requirement for planarity is justified on psychological grounds rather than as a means for achieving a geometrically more competent algorithm.

The preference of the HVS to interpret a line drawing as the most symmetric polyhedral (planar-faced) object consistent with the drawing is well established in the psychological literature. Marill appeared to have discovered a simple computational procedure for finding such solutions for any given line drawing, but on closer examination, it became apparent that his MSDA principle does not enforce (or even prefer) planar solutions.<sup>17</sup> Because of this deficiency, MSDA is unreliable even in very simple situations. We were able to prove

(appendix D) that if a planarity preference is explicitly added to the MSDA objective function, then indeed, the nonobvious preference for symmetric solutions is also present. However, we are now forced to address the problem of how to provide the auxiliary information necessary to partition the drawing into the coherent components corresponding to the 3D planar faces. It appears that the HVS selects some subset of the contours in the line drawing as corresponding to the planar faces in the 3D model, and if we do not supply this information to a recovery algorithm (either explicitly or by providing a set of conditions implying the same information), we will fail to recover psychologically acceptable models.

Most of the work in the blocks-world tradition employed perfect *labeled* line drawings with the assignment of vertexes to faces given as part of the input specifications. If we follow the same approach (although we are not concerned with having perfect line drawings since our recovery method employs optimization, which can tolerate deviations from any of the constraints embodied in the objective function), then we at least have provided a tool for simplifying man-machine communication using the language of line drawings. However, there is obvious theoretical value in understanding the criterion for human selection of the circuits in the line drawing that correspond to planar faces in the 3D model.<sup>18</sup> In part, this importance is related to the issue of how the HVS recovers the shape of a moving object. Even though there are a few well-known exceptions, it is widely believed that the HVS will assume an object to be rigid and correctly recover its shape if this is indeed the case.<sup>19</sup> However, the rigid wire frames with nonplanar faces provide a whole class of counter-examples to this belief—they appear to be nonrigid when observed in motion (even at very low speeds where maintaining correspondence of vertexes from one projection to the next is no problem). The nonrigidity appears to result from the HVS making incorrect decisions about how the drawing can be partitioned into planar faces (see appendix E).

## 7 Summary

Marill's recently published paper claimed that the simple procedure he described could duplicate human judgment in recovering the 3D wire frame geometry of objects depicted in line drawings. He provided some impressive examples, but no theoretical justification to

back his claims. In this article, we critically examined the merits of Marill's algorithm, provided at least a partial explanation for its competence, identified weaknesses, showed how it could be improved, and discussed the implications of this work for clarifying some important problems in human perception.

In particular, we provided a number of theorems that show that minimizing the standard deviation of angles is (potentially) a simple and effective method for selecting symmetric solutions when the constraining line drawing (which is the projection of a wire frame that may be incomplete) permits such interpretation. On the other hand, we showed that Marill's algorithm could fail in simple cases, that he employed an optimization procedure that was often too weak to find the correct answer even when it was within the competence of the objective function, and that the algorithm would often produce wire frames with nonplanar faces (something no human would intuitively accept in perceiving a straight-line drawing as a 3D configuration).

We argued that an important condition in testing or evaluating the psychological plausibility of a reconstruction is that its reprojections (after spatial relocation) result in the same object being produced by the recovery algorithm. For the human visual system, this is equivalent to the condition that the recovered object appear rigid when observed during movement or rotation. The perception of rigidity for wire frames appears to be highly correlated with the presence or absence of strongly nonplanar faces. By modifying Marill's objective function to explicitly favor planar-faced solutions, and by using a more competent optimization technique, we were able to demonstrate significantly improved performance in all of the examples Marill provided as well as those additional ones we constructed ourselves. The robustness of our algorithm was demonstrated by obtaining consistent psychologically plausible reconstructions in hundreds of experiments involving variations in viewpoint and initial conditions for the approximately 20 objects in our database.

### Acknowledgements

The work reported here was partially supported by the Defense Advanced Research Projects Agency. We gratefully acknowledge the valuable discussions with Aaron Bobick and Thomas Strat regarding both the content and organization of this article. A number of improvements and clarifications in its final version were suggested by Thomas Marill in a private communication.

### Notes

1. Gradient space, originally conceived of by James Clerk Maxwell in 1864 (see (Whitely 1986)) and rediscovered by D.A. Huffman, provides only necessary conditions for planar realizability of general polyhedral objects with hidden lines removed, and thus consistent edge labeling is possible for impossible blocks world and Origami objects. Further, the labeling/recovery algorithms were not always competent to find an existing solution.
2. There were some other problems of lesser significance for our purposes. For example, the algebraic formulation was sensitive to computation round-off errors, and digitization errors in specifying the line drawing; a realizable object could be rejected because of such minor numeric inaccuracies. Sugihara dealt with this problem by adding an optimization step to his algorithm, which could find a feasible reconstruction if the input drawing was an almost correct specification.
3. Marill, on the other hand, calls the set of all possible  $z$ s the orthographic extension.
4. For simplicity, the vertexes are represented by only two digits of precision in the table. However, we used the full 32-bit precision of the projection in the internal representation used by the algorithms.
5. We note that while there generally can be many different ways of covering a line drawing, those of blocks-world objects with hidden lines removed will be covered uniquely if we demand that the interior of the 2D circuits be free of any lines. We also note that it is not always possible to cover a line drawing with *simple closed* circuits corresponding to the specified planar faces of a given orthographic extension (see example N). It may also be the case that a given covering has no nontrivial orthographic extension with planar faces as specified, as in example O.
6. One face, in example H, is an exception to this statement. However, there are enough other geometric constraints in this particular case to enforce planarity.
7. Because this rule typically produces many additional planar faces, it was not used in figures 2 through 6. For these line drawings, the results are virtually identical with or without these additional planar faces. However, the rule was used in the stability and robustness experiments of section 3.4.
8. The SDA term is first squared to make it commensurate with the DP term. Note that squaring the SDA term has no effect on the minimization when  $\lambda = 1$  (i.e., the simple MSDA algorithm), because the SDA term is positive, and squaring is a monotonic function of the positive reals.
9. This assumes the line drawing is perfect. We later discuss how such perfect drawings can be obtained in an interactive environment.
10. Since we had only the original line drawing for each of Marill's examples, we used the reconstruction from each line drawing as the 3D object for the random projections.
11. The line lengths for these drawings were approximately in the range of 2 to 5.
12. For all line drawings except the truncated box and the hinge, the largest absolute difference in angles between any trial and the reconstruction with  $z = 0$  was less than one degree. For the truncated box, the largest error was less than fifteen degrees. For the hinge, one of the trials caused the hinge to "fold" with arcs pairs (1 0 4) and (2 3 5) going to zero degrees. Otherwise, the largest error was less than seven degrees.
13. Modulo a change in sign in the  $z$  coordinates.

14. The set of free vertexes is by no means unique. For example, any three vertexes from one hexagon plus any vertex from the other hexagon will do for this line drawing. We have a simple algorithm for finding a set of free vertexes, but have not yet proven that it is correct, so we do not present it here.
15. The successive reconstructions are not independent; to the extent that they allow a range of interpretations, the parameters selected for one interpretation will influence the parameter selections for successive interpretations.
16. Some differences are that Barrow and Tenenbaum considered arbitrary, but *known* perspective transforms in their paper, while Marill used only orthographic projections. In either case, the set of state variables is equivalent. In addition, Barrow and Tenenbaum did not consider the use of a continuation method.
17. Marill, of course, only returns the wire frame. But in the case of a blocks-world object, competent algorithms exist for finding all the valid completions of the wire frame as a solid polyhedral object (Strat 1984, Markowsky and Wesley 1981).
18. As noted in section 3.1 and appendix A, we have made some initial progress toward the solution of this problem and have developed an algorithmic procedure that can successfully handle all of the examples discussed in this paper, but we recognize that this is still far short of a complete solution.
19. For example, by using Ullman's result that three distinct orthographic projections of four noncoplanar points in a rigid configuration are sufficient to uniquely determine the structure and motion up to a reflection about the image plane.

## References

- Barnard, S.T., and Pentland, A.P., 1983. Three-dimensional shape from line drawings, *Proc. 8th Intern. Joint Conf. Artif. Intell.*, Karlsruhe, W. Germany.
- Barrow, H.G., and Tenenbaum, J.M., 1981. Interpreting line drawings as three-dimensional surfaces, *Artificial Intelligence* 17(1-3):75-116.
- Clowes, M.B., 1971. On seeing things, *Artificial Intelligence* 2(1):79-116.
- Draper, S.W., 1981. The use of gradient and dual space in line-drawing interpretation, *Artificial Intelligence* 17:461-508.
- Hochberg, J., and McAlister, E., 1953. A quantitative approach to figure "goodness." *J. Exp. Psychol.* 46:361-364.
- Huffman, D.A. 1971. Impossible objects as nonsense sentences. In Meltzer and Michie, ed., *Machine Intelligence* Vol. 6, p. 295-323, Edinburgh University Press.
- Kanade, T., 1980. A theory of Origami world, *Artificial Intelligence* 13(1):279-311.
- Leclerc, Y.G. 1989. Constructing simple stable descriptions for image partitioning. *Intern. J. Comput. Vis.* 3(1):73-102.
- Mackworth, A.K., 1973. Interpreting pictures of polyhedral scenes, *Artificial Intelligence* 4(2):121-137.
- Malik, J., and Maydan, D., 1989. Recovering three-dimensional shape from a single image of curved objects, *IEEE Trans. Patt. Anal. Mach. Intell.* 11(6):555-566.
- Marill, T., 1991. Emulating the human interpretation of line-drawings as three-dimensional objects, *Intern. J. Comput. Vis.* 6(2):147-161.
- Markowsky, M.A., and Wesley, G., 1981. Fleshing out projections, *IBM J. Res. Develop.* 25(6):934-954.

- Pentland, A., and Kuo, J., 1990. Three-dimensional line interpretation via local processing, *Electronic Imaging* 1249-30. Also Media Lab Technical Report 131.
- Pomerantz, J.R., and Kubovy, M., 1981. Perceptual organization: an overview. In *Perceptual Organization* pp. 423-456, Lawrence Erlbaum Associates; Hillsdale, NJ.
- Press, W.H., Flannery, B.P., Teukolsky, S.A., and Vetterling, W.T., 1986. *Numerical recipes, the art of scientific computing*, Cambridge University Press, Cambridge.
- Stevens, K.A., 1981. The visual interpretation of surface contours, *Artificial Intelligence* 17(1-3):47-74.
- Strat, T.M., 1984. Spatial reasoning from line drawings of polyhedra, *Proc. DARPA Image Understanding Workshop*, New Orleans, p. 230-235.
- Sugihara, K., 1982. Mathematical structures of line drawings of polyhedrons—toward man-machine communication by means of line drawings, *IEEE Trans. Patt. Anal. Mach. Intell.*, 4(5):458-469.
- Sugihara, K., 1984. A necessary and sufficient condition for a picture to represent a polyhedral scene. *IEEE Trans. Patt. Anal. Mach. Intell.*, 6(5):578-586.
- Waltz, D.A., 1972. Generating semantic descriptions from line drawings of scenes with shadows. Technical Report AI-TR-271, MIT.
- Whitely, W., 1986. Two algorithms for polyhedral pictures, *Proc. 2nd Annual Symp. Computa. Geometry*, p. 142-149.
- Witkin, A.W., 1981. Recovering surface shape and orientation from texture, *Artificial Intelligence* 17:17-47.
- Witkin, A.W., Terzopoulos, D., and Kass, M., 1987. Signal matching through scale space, *Intern. J. Comput. Vis.* 1:133-144.

## Appendix A. Psychological Assumptions

The following are some of the basic assumptions that we believe are typically made by people in the reconstructions of wire frames from line drawings, and some constraints relevant to partitioning a line drawing into planar face. They are known to have rare exceptions.

1. Three dimensional wire frames, derived from line drawings, have implied planar faces inside subsets of their closed circuits; they *can* also have struts, such as legs or bracing wires, in or on a planar face. (Strongly nonplanar faces produce psychologically implausible solutions.)
2. Symmetric reconstructions are preferred over non-symmetric ones.
3. Parallel lines in a line drawing are parallel in space. Lines connecting vertexes falling on two parallel lines are in a common plane with the two parallel lines.
4. Many-sided convex closed contours without *internal circuits* (in a 2D line drawing) are likely to correspond to the contours of planar faces in the corresponding 3D orthographic extension (see B4). An

internal circuit to a convex polygon is defined to be a circuit for which all the vertexes are internal to the polygon, and for which the ends of the circuit lie on nonadjacent vertexes of the polygon.

5. A closed simple contour in a line drawing, without internal lines, corresponds to a planar face in the corresponding 3D reconstruction.

An algorithmic procedure for identifying 3D planar faces in the corresponding 2D line drawing of a wire frame has been constructed by composing the requirements of items 3, 4, and 5 into a single algorithm, as defined in section 3. That procedure is sufficient to deal with all of the examples we discuss here, but is not general enough to handle other cases we can think of.

## Appendix B. Projective Invariants

The following are some important projective invariants for planar geometric structures.

1. The sum of the interior angles (measured between 0 and 360 degrees) of a closed planar contour with  $n$  sides equals  $(n - 2)180$  degrees. Thus, since a polygon of  $n$  sides projects to a polygon of  $n$  sides under both orthographic and central projection, the mean value of the interior angles of a given closed planar contour  $[(n - 2)180/n]$  is invariant under both orthographic and central projection.

We note that Marill measures angles only in the interval between 0 and 180 degrees. To the extent that we are primarily concerned with equiangular closed contours in the application of the above theorem in explaining and using his results, this discrepancy is irrelevant since all the interior angles of such contours are less than 180 degrees.

2. Consider an angle (two line segments sharing a common endpoint) in 3D space and its orthographic projection. We will call the plane containing the angle the source plane, and the plane containing its projection the projection plane. If the angle is translated in the source plane, its projection is also translated, but does not change in magnitude from its original projected value. Now consider a set of  $n$  angles lying on a common source plane, such that the sum of these angles is 360 degrees. If it is also the case that the angles can be translated so that when all their vertexes coincide, they exactly span an angle of 360 degrees, then the mean value of the set of angles  $(360/n)$  is unaltered under orthographic projections. We will call such a collection of angles a "complete-star." (Example C, for instance, contains

a complete-star consisting of the eight 45-degree angles formed at the corner vertexes by the diagonals with the sides of the square. Example G contains this same configuration in its central plane.) We note that if an essentially infinite number of copies of an angle of  $d$  degrees (where  $360/d = k$  and  $k$  is an integer) is uniformly distributed in orientation over a plane, then the mean value of the angles under any orthographic projection of the plane is the constant value  $d$ .

3. We note that if the angle between two line segments is less than 180 degrees, the angle can be closed to form a triangle, and since triangles are preserved under both orthographic and central projection, an angle of less than 180 degrees will never transform under such projections into one of more than 180 degrees. We will call a closed planar contour convex if the region it bounds is convex. Since a convex contour has all internal angles of less than 180 degrees, a convex planar contour remains convex under both orthographic and central projection.
4. We note that the orthographic projection of an arbitrary nonplanar polygonal space curve, with four or more sides, has a probability of projecting to either a nonsimple or concave curve with a probability ( $P$ ) that increases with the number of sides:

$$P > 1 - 0.5^{n-3} \quad \text{for } n \geq 4$$

This expression is based on the following model: Consider a process that generates a chain of 3D random vectors by generating three random numbers for each vector (in spherical coordinates, an angle uniformly distributed between 0 and 360 degrees, a second angle between 0 and 180 degrees, and a length uniformly distributed between 0 and some fixed integer  $L$ ). As each vector is generated we extend the projection of the developing space curve on the  $XY$  image plane. The process stops after some fixed number of steps, which is determined by choosing a random number in some given range; the curve is now closed by connecting the starting point, which could be the origin of the  $XY$  plane, to the last point generated and this determines whether the inside is to the left or right as we follow the chain of edges of the projected polygon. We note that the only relevant factor in whether the projected closed contour is convex or concave is the cylindrical angle giving the rotation of each of the random vectors relative to the  $X$  axis in the image plane. For more than three sides, there is a 50% probability at each vertex that the inside angle is greater than 180 degrees,

which thus produces a concave polygon (the last closing side can be ignored since it does not have the same statistics as the other edges in our random model.) Other probabilistic models would give nonidentical, but similar results. The  $>$  condition is based on additional considerations, such as the projected curve intersecting itself even though the input specification does not record a vertex at the cross-point.

5. Closed four-sided polygonal space curves with 90-degree angles at each vertex are planar contours. To prove this assertion, let the sequence of vertexes be labeled  $a, b, c,$  and  $d$ . Let the plane containing lines  $L_{ab}$  and  $L_{bc}$  (and thus vertexes  $a, b,$  and  $c$ ) be called  $P_1$ . Since all angles are 90 degrees,  $L_{cd}$  must lie in a plane ( $P_2$ ) normal to  $L_{bc}$  at  $c$ . Similarly,  $L_{ad}$  must lie in a plane ( $P_3$ ) normal to  $L_{ab}$  at  $a$ . Vertex  $d$  must then lie on the line ( $L_d$  of intersection of  $P_2$  and  $P_3$ , which is normal to  $P_1$ . We know one solution is to locate  $d$  at the point of intersection ( $d^*$ ) of  $L_d$  and  $P_1$  (where  $a, b, c,$  and  $d^*$  form a rectangle). This is the planar solution and we wish to show that no other solution is possible. We note that a second constraint on the location of  $d$  is that it must lie on a sphere with diameter  $ac$  (i.e., all right angles, with legs passing through points  $a$  and  $c$ , must be inscribed angles of circles through  $a$  and  $c$  with diameter  $ac$ ). We know  $d^*$  lies on the sphere and  $P_1$  is a bisecting plane of the sphere. Thus  $L_d$  is tangent to the sphere at  $d^*$  and  $d^*$  is the only possible solution.

6. A Global Planarity Test for a Space Curve. A planar polygonal curve has a sum of internal angles equal to  $(n - 2)180$  degrees. Thus, if the curve is triangulated using only the existing vertexes along the curve, the sum of the angles of the triangles is also  $(n - 2)180$ .

**Case 1:** Consider a space curve  $S$  that projects to a convex planar curve. If the space curve is itself planar, the sum of its angles (measured between 0 and 180 degrees) is  $(n - 2)180$ . Assume  $S$  is nonplanar, that is, there is a "fold" along one or more edges of some triangulation of its planar projection. Consider the vertex  $V$  at the intersection of one such fold (with respect to the implied triangulation  $T$ ) and  $S$ . The plane through the two edges of  $S$  meeting at  $V$ , and the faces of the triangles of  $T$  that have edges intersecting at  $V$ , form a polyhedral angle. It is known that any face angle of a polyhedral angle is less than the sum of the other face angles. Therefore, the sum of the angles of the space curve is equal (at vertexes with no folding) or less (at vertexes with folding) than the sum of the angles

of the triangles in  $T$  (i.e., less than  $(n - 2)180$ ).

**Case 2:** If the projection of the space curve  $S$  is concave, and we measure angles between 0 and 360 degrees, the sum of the *internal* contour angles in the planar projection will equal  $(n - 2)180$  as in Case 1. However, while the space angles with projections of less than 180 degrees will decrease at folds, the internal angles greater than 180 degrees will increase (i.e., at vertexes where there are folds, the polyhedral angle in the argument given in Case 1 is now formed for the *external* angle of  $S$  at  $V$ ). Thus, since some angles will increase and others decrease, we cannot be sure that the curve is planar even if the sum of its internal angles equals  $(n - 2)180$ . However, we do have a sufficient condition for nonplanarity. That is, the curve is known to be nonplanar if the sum of its internal angles, measured between 0 and 360 degrees, is not equal to  $(n - 2)180$  degrees.

### Appendix C. A Partition Theorem

The variance of a set of  $S$  of  $n$  objects  $\{a_i\}$  is defined as

$$V = \frac{1}{n} \sum_{i=1}^n (a_i - M)^2 = \left[ \frac{\sum_{i=1}^n a_i^2}{n} \right] - M^2$$

where

$$M = \frac{1}{n} \sum_{i=1}^n a_i$$

Let us now partition the  $\{a_i\}$  into  $k$  subsets, such that the subset  $S_j$  has  $n_j$  elements and mean  $M_j$  where:

$$M_j = \frac{1}{n_j} \sum_{S_j} a_i$$

Let  $V_j$  be the variance of  $S_j$  about  $M_j$  and let  $\Delta_j = (M - M_j)$ .

**Theorem:**

$$V = \frac{1}{n} \sum_{j=1}^k n_j [V_j + \Delta_j^2]$$

*Proof:* The expression for  $V$  can be rewritten as

$$V = \frac{1}{n} \left[ \sum_{S_1} [a_i - (M_1 + \Delta_1)]^2 + \sum_{S_2} [a_i - (M_2 + \Delta_2)]^2 + \dots + \sum_{S_k} [a_i - (M_k + \Delta_k)]^2 \right]$$

If we let:

$$V_j^* = \sum_{s_j} [a_i - (M_j + \Delta_j)]^2$$

Then we have:

$$\begin{aligned} \frac{V_j^*}{n_j} &= \sum \frac{a_i^2}{n_j} + M_j^2 + \Delta_j^2 - 2\Delta_j \sum \frac{a_i}{n_j} \\ &\quad - 2M_j \sum \frac{a_i}{n_j} + 2M_j \Delta_j \end{aligned}$$

Given that  $\Sigma(a_i/n_j) = M_j$ , we note that the 4th and 6th terms cancel and the 2nd and 5th terms combine:

$$\frac{V_j^*}{n_j} = \left[ \sum \frac{a_i^2}{n_j} - M_j^2 \right] + \Delta_j^2 = V_j + \Delta_j^2$$

And

$$V_j^* = n_j[V_j + \Delta_j^2]$$

QED

#### Appendix D. Symmetric Preference Theorem

Recall that

1. In appendix B we showed that the average angle of all planar orthographic extensions of a given simple closed 2D contour are the same, and that the average angle of all planar orthographic extensions of a complete-star are also the same;
2. in appendix C we proved a theorem that allows us to compute the SDA of a set of simple closed planar contours (and/or complete-stars) as the sum of two components. The first component is the variance of the angles in a contour or star about the mean angle of that contour or star, summed over all contours and stars. The second component is a weighted sum of the squared differences between the mean angle of each contour and star, and the average of all the angles under consideration.

By (1), the second component of the variance is constant over all planar orthographic extensions because (a) the mean of each contour and star is constant over all such extensions, and (b) the mean of all angles can be computed as the weighted sum of the mean of each contour and star.

Consequently, if we restrict our attention to the planar orthographic extensions of a line drawing, then by (2) above, only the first component of the variance will change over the extensions. Since the first component is zero for an extension comprising only equi-

angular planar contours and stars (such as the solutions for examples A, B, C, G, J, K, and L), and since it is positive otherwise, then such symmetric solutions correspond to the global minimum of the SDA over all planar orthographic extensions.

#### Appendix E. Factors Affecting the Perception of Nonrigidity

If we rotate a *randomly* derived orthographic extension of almost any of the line drawings used as examples in this article, the object appears nonrigid to most observers (even though, of course, the wire frame is actually a rigid object). While there are many possible explanations for this phenomenon, our conjecture is that it is primarily due to special position projections of the wire frame (that occur at one or more poses in its rotation) that lead the HVS to incorrectly assume that some projective invariant (such as parallel lines, see figure 11) is being observed. This, in turn, causes incorrect expectations about the presence and location of planar faces.

We informally looked at some other possible causative factors, but did not observe consistent nonrigidity phenomena. For example, we looked at objects, such as example N that produce compelling 3D interpretations with Necker reversals, but for which the drawing is incomplete—it does not show all the edges that should be visible, for example, where planar faces intersect. There was the possibility that these missing edges in the 3D model (and thus missing lines in the drawing) could cause the appearance of a nonplanar-faced object to be observed. But the hinge, and the few other objects we looked at in this category, appeared rigid.

We also looked at nonplanar orthographic extensions of drawings that generally appeared flat, including blocks-world type drawings that do not have corresponding polyhedral realizations (such as example O). The results here were ambiguous. The rotating objects generally produced illusions of nonrigidity, but since these objects did not always appear 3D, the illusions were generally very weak.

Some other causal experiments include cases where all the lines connecting the vertexes of the wire frames are deleted; we observed that some of the wire frames that originally appeared nonrigid now appeared to be rigid under rotation. And, as a general observation, we have not encountered any examples in which the wire frame of a (nondegenerate) blocks-world object appears nonrigid when in motion.