

3-D Interpretation of Imperfect Line Drawings

Ronald Chung

Kin-lap Leung

Department of Mechanical & Automation Engineering
The Chinese University of Hong Kong, Shatin, Hong Kong
E-mail: rchung@mae.cuhk.hk

Abstract

Recovering three-dimensional shape of an object from a single line drawing is a classical problem in computer vision. Methods proposed range from Huffman-Clowes junction labeling, to Kanade's gradient space and skew symmetry analysis, to Sugihara's necessary and sufficient condition for a realizable polyhedral object, to Marill's MSDA shape recovery procedure, and to the recent Leclerc-Fischler's shape recovery procedure which assures planar faces. Yet all these assume perfect line drawings as the input. We propose a method that through the use of iterative clustering interprets an imperfect line drawing of a polyhedral scene. It distinguishes the true surface boundaries from the false ones like the surface markings, fill-in the missing surface boundaries, and recovers 3-D shapes satisfying constraints of planarity of faces and parallel symmetry of lines, all at the same time. Experiments also show that the 3-D interpretation agrees with human perception.

1 The Line-Drawing Interpretation Problem

Humans interpret many line drawings or cartoon-like pictures with ease and without the feeling of much loss of information about the shapes of the objects. This shows that even if other cues like shading and texture are absent, contours alone already convey much information about shapes, and it is especially true for objects which display significant degree of regularity. Such an ability, if emulated in a machine, can have applications ranging from robotics to cartography.

There have been a great deal of work on such a so-called shape-from-contour problem. Huffman [2] and Clowes [1] separately proposed a junction-labeling method to recover a three-dimensional (3-D) description of polyhedral objects, though qualitatively. Mackworth [6] and Kanade [3] proposed gradient space analysis and skew symmetry to recover quantitative 3-D description. Sugihara [9] provided an algebraic criterion as a necessary and sufficient condition for a line drawing to represent a physically realizable polyhedral object. Marill [7] also proposed a method that reconstructs polyhedral objects, which is based on minimizing a simple objective function about the standard deviation of the internal angles of each surface. Recently, Leclerc and Fischler [4] added a planarity term

to the objective function, which enforces the vertices of the same surface to be coplanar. This new objective function gives extremely good results.

However, all these methods assume perfect line drawings. In real images the contours extracted are not necessarily surface boundaries; they can be surface markings, cracks, shadows, imaging noise, and others. Human perception also seems not to require a perfect line drawing to function; it would tell surface boundaries and markings apart, and it would fill-in the missing boundaries for many line drawings.

Use of region segmentation techniques [8] would allow a number of closed regions of defined homogeneity to be extracted from an image. Such segmentation is likely to return corrupted line drawings: lines which are not surface boundaries are present and some surface boundaries are missing. Questions are, how surface boundaries can be distinguished from other contours, how missing boundaries can be filled-in, and how 3-D information can be inferred about the objects. The traditional belief is that these can be done in a sequential manner: segmentation comes first which identifies the true surface boundaries, then the shape recovery techniques mentioned above follow. Here we take a different point of view. We argue that segmentation has to return boundaries which are discontinuities of depth or orientation, and at the same time 3-D interpretation has to know where the true surface boundaries are, a typical chicken-and-egg problem. From this perspective, the two processes should work in parallel cooperatively to come up with a consistent output.

We address the problem of interpreting an imperfect line drawing of a polyhedral scene in this paper. The input is a set of closed regions which may come from region segmentation or from perceptual grouping techniques applied towards detected edges of the image. The goal is to distinguish surface boundaries from surface markings and other false lines, to infer missing lines, as well as to recover 3-D shapes of the objects. We use an interactive clustering procedure that allows the segmentation and shape recovery processes to come into mutual agreement, during which one process acts on the intermediate result of the other process in turns.

2 An Iterative Clustering Approach

Recovering 3-D description from a 2-D line drawing is an ill-posed problem; there are an infinite number of different 3-D objects that can give the same 2-D projection. Assumptions about the object shape are necessary to constrain the problem. Since we are dealing with polyhedral objects, planarity of all surfaces is naturally a fundamental constraint. Under the general viewpoint assumption where observed features and their interrelationships are stable upon slight perturbations of the viewpoint, 2-D parallelism under orthographic projection can only be projected by 3-D parallelism. This is also in coherence with the Gestalt school of psychology which proposed that parallel symmetry plays an important role in human perception. Planarity of surfaces and 2-D parallelism be projected by 3-D parallelism are therefore the two constraints that a sensible 3-D output should satisfy, and they largely narrow the solution space to a few number of solution points.

However, the original image itself, as a planar object, is certainly one that

satisfies all the above constraints. Such a solution is unlikely in reality, and it often does not agree with human perception. We need another term that discourages such an interpretation. We use a measure of three-dimensionality of an object, which is zero if the object is planar, and is of large value if the object is a solid one with small eccentricity in the 3-D space.

The problem can therefore be formulated as a constraint satisfaction problem: given a graph $G = (V, E)$ where $V = \{v_p : v_p = (x_p, y_p)\}$ are the vertices in the image and $E = \{(v_p, v_q) : v_p, v_q \in V\}$ are the edges among the vertices, the goal is to find a depth measure z_p for each vertex v_p , as well as a set ξ of surfaces in the scene, which is a family of subsets of E (in fact a covering of E), such that they satisfy the above constraints. To put it more formally, we want to come up with a hypergraph $H = (V', E, \xi)$, where $V' = \{v'_p : v'_p = (x_p, y_p, z_p)\}$ and ξ is a set of hyperedges on E , such that it satisfies:

1. Planarity Constraint

- (a) The vertices in each surface $\varepsilon \in \xi$ are coplanar.

We formulate it as: the measure

$$\text{PLN}(\{z_p\}, \xi) = \frac{\sum_{\varepsilon \in \xi} \sum_i \{1 - [\frac{(\vec{e}_{\varepsilon, i-1} \times \vec{e}_{\varepsilon, i}) \cdot (\vec{e}_{\varepsilon, i} \times \vec{e}_{\varepsilon, i+1})}{\|\vec{e}_{\varepsilon, i-1}\| \|\vec{e}_{\varepsilon, i}\|^2 \|\vec{e}_{\varepsilon, i+1}\|}]^2\}}{\sum_{\varepsilon \in \xi} \sum_i 1}$$

where $\vec{e}_{\varepsilon, i}$ is the i th edge in 3-D of surface ε , should ideally be zero and practically be a small value.

- (b) No two hyperedges sharing edges, *i.e.*, no two neighboring surfaces, are coplanar.

2. Parallel Symmetry Constraint: Parallel lines in the line drawing should be projected by parallel lines in 3-D.

We formulate it as: the measure

$$\text{PRL}(\{z_p\}) = \frac{\sum_{i,j} \{1 - [\frac{\vec{e}_i \cdot \vec{e}_j}{\|\vec{e}_i\| \|\vec{e}_j\|}]^2\}}{\sum_{i,j} 1}$$

where \vec{e}_i and \vec{e}_j are a pair of edges in 3-D which appear to be parallel in the line drawing, should ideally be zero and practically be a small value.

3. Maximum Three-dimensionality: A solid object is preferred over a sheet object.

We formulate it as: the measure

$$\text{3-Dness}(\{z_p\}) = \frac{\sum_p \sum_i [\frac{(\vec{e}_{p, i-1} \times \vec{e}_{p, i}) \cdot \vec{e}_{p, i+1}}{\|\vec{e}_{p, i-1}\| \|\vec{e}_{p, i}\| \|\vec{e}_{p, i+1}\|}]^2}{\sum_p \sum_i 1}$$

where $\vec{e}_{p, i}$ is the i th edge in 3-D of the p th vertex, should be maximized,

If ξ is known, the problem becomes simple. We find z_p 's that minimizes the value of the overall objective function

$$E(\{z_p\}, \xi) = -3\text{-Dness}(\{z_p\}) + \alpha \text{PRL}(\{z_p\}) + \beta \text{PLN}(\{z_p\}, \xi) \quad (1)$$

which is a weighted sum of the above constraints.

For any α, β , a locally optimal solution of z_p 's can be found using the *hill-climbing method*. The method is simple, and it does not require the derivative expression of the objective function which can be complicated. It can be described as the following. Given any initial value of the solution vector $\vec{Z} = (z_1, \dots, z_p, \dots, z_{\|V'\|}, 2\|V'\|)$ new solution vectors are formed by adding or subtracting a small value Δz from one of its entries. The vector giving the minimum value of the objective function is selected as the new current vector. The procedure repeats until the current solution vector is stable. Since what comes out is only the locally optimal solution closest to the initial solution vector, the initial values of z_p 's are crucial. It will be discussed in a later section.

Such a locally optimal solution of z_p 's does not assure the absolute constraints of planarity and parallelism be satisfied exactly. To assure that, we can use the *penalty method*: the weights α and β are increased from small values gradually in small steps until they are at large values, while the locally optimal solution of z_p 's is computed at each α, β -setting and carried forward. As the weights increase, the absolute constraints dominate the objective function in a smooth fashion.

What is missing in the above formulation is that we do not know ξ , the segmentation solution. Each innermost cycle of the edges, appeared to be a single surface, can be just part of a large surface if some of the apparent edges are merely surface markings. On the other hand, it can be the boundary of more than one surface if some surface boundaries inside it are missing.

Segmentation requires knowledge of the 3-D shape, while estimating 3-D shape requires knowledge of segmentation. The ultimate solution will be the state where the two come to a mutual agreement. One possible method is to have iterations over a number of intermediate shapes and segmentations until they get to a stable, consistent state. Fortunately, even without knowledge of segmentation, as seen from the objective function (1) above, the parallelism and three-dimensionality constraints can still be applied to the input line drawing and a coarse shape description can be constructed. Such a shape description does not have the notion of planar surfaces, but can be close enough to the true shape for subsequent segmentation purpose because of the highly restrictive constraint of parallelism. Obviously, such a formulation would fail if there are no parallel lines in the input line drawing. However, we conjecture that even humans have difficulty in perceiving shape from a line drawing which is without much regularity or symmetry.

With such an approximate shape and the line drawing, we can group the edges into different surfaces under the criterion that edges on the same surface should be more or less coplanar. To facilitate finding out which planes in 3-D the edges vote for, we can subdivide each vertex into a number of *L-subvertices*, where each *L*-subvertex is merely the corresponding vertex point itself plus a pair of the vertex's component edges. Defining such *L*-subvertices has the advantage that each of them uniquely defines a plane in the 3-D space. We can then treat the segmentation problem as a clustering problem, in which the extracted *L*-subvertices are to be grouped into different surfaces according to how close their corresponding planes are.

The resulting segmentation may not be entirely correct, but it introduces the concept of approximate surfaces to the original line drawing. As a result, it allows

markings contained in the extracted surfaces to be deleted and missing boundaries of some open surfaces to be added. In other words, a new line drawing is constructed, which is likely to be closer to the true line drawing than the original one. Such a line drawing can be input to the approximate shape recovery process described above and all procedures are repeated until the shape recovered and the revised line drawing are stable, *i.e.*, the segmentation and shape recovery processes agree with each other. We then add the planarity term to arrive to the final shape.

Our approach therefore consists of iterations over three steps, as outlined in Figure 1:

- **Step 1: Approximate Shape Recovery**

Given the current line drawing, initially the input line drawing, find an approximate shape that optimizes the parallelism and three-dimensionality constraints. The penalty method plus the hill-climbing algorithm described above can be used.

- **Step 2: L -subvertex Clustering**

Given the current approximate shape and line drawing, extract the L -subvertices and group them into different surfaces using the nearest-neighbor clustering algorithm [5].

Each L -subvertex represents a 4-D normalized vector (a, b, c, d) corresponding to the plane $ax + by + cz + d = 0$ in 3-D, whose values can be obtained from $\vec{n} \cdot ([x \ y \ z]^T - \vec{v}_p) = 0$ where \vec{n} is the cross product of the edges of the subvertex and \vec{v}_p the 3-D position of the vertex. The inter-pattern distance measure for clustering can be the magnitude of the cross-product of such 4-D unit vectors the two L -subvertices represent.

With the surface segmentation, new lines are then added to complete the open surfaces, and lines are removed if they are contained inside surfaces. A new line drawing is formed.

If the new line drawing is different from the previous one, go to step 1, else go to step 3.

- **Step 3: Global Consistency Assurance**

Given a perfect line drawing, find the 3-D shape that is globally consistent and that is optimal: it satisfies constraints of parallelism and planarity exactly and maximizes the three-dimensionality measure. Again, the penalty method plus the hill-climbing algorithm can be used. This can be regarded as a finishing step of the shape description.

Two points are worth noting:

1. Free Vertices:

A vertex with only two component edges is not counted towards three-dimensionality, as it does not involve multiple surfaces. If none of its component edges is parallel to other edges in the line drawing, such a vertex is also not involved in the parallelism constraint. As a result, this vertex does

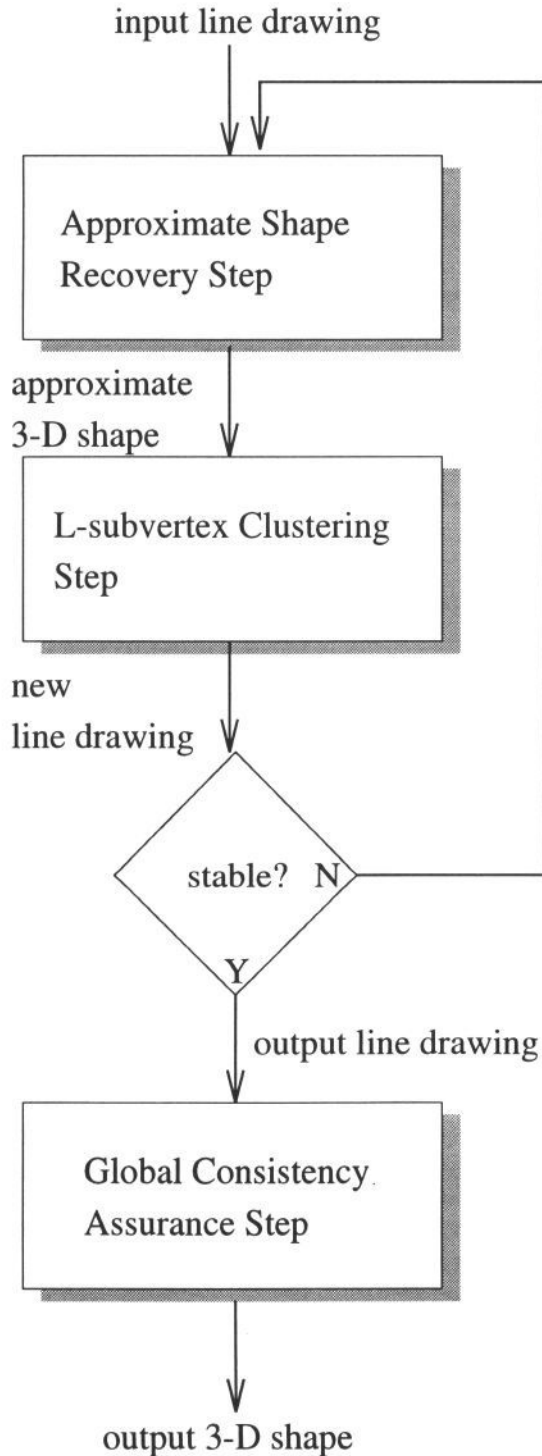


Figure 1: Overview of the clustering approach.

not "move" in step 1. We call such a vertex a *free vertex*. Borrowing from the planarity concept, we can in fact move the free vertex to a reasonable position by assigning it with a normal angle which is the mean of the most coplanar normal angles of the two connected vertices.

2. Interpretation Ambiguity and Initial Depth Values for optimization:

A phenomenon which is related to the famous Necker's cube problem is, for some line drawings there are 3-D interpretations which both satisfy the constraints of planarity, parallelism, and three-dimensionality, and yet are entirely different. Imagine the line drawing of a rectangular box: the vertex in the middle can be interpreted as convex or concave. However, there is a slight preference of a convex vertex in human perception. A reason may be that the inside of objects are often not visible.

To remove such an ambiguity in our solution, we give a bias to the initial depth values in the optimization procedure of steps 1 and 2: assign zero depth to all vertices on the outsidemost boundary of the line drawing, and small depth values to the inside vertices so that they are initially slightly convex to the viewer. If there are no inside vertices, we assign zero depth to all bihedral vertices and small depth values to the rest.

3 Experimental Results


We have applied the proposed procedure to a number of different line drawings.

Results on an example used in [4] but with surface markings added and some boundaries removed, is presented in Figure 2. The object went through the *L*-subvertex clustering step twice, and the approximately recovered shape was slightly distorted. However, with the addition of the planarity constraint in the global consistency assurance step, the result is a perfect hexagonal object.


Results on a more complicated line drawing with multiple occlusions are shown in Figure 4. It consists of four objects, some individuals of which are used in [7, 4]. Along the occlusion boundaries there are T-junctions, which are where one end of an edge (the stem) meets the middle of another edge (the cap). The stem is an edge of the occluded object, while the cap is an edge of the occluding object. We allocate two depth values to each T-junction, one for the occluding object and one for the occluded object. Then the iteration steps described above are applied to the line drawing as a whole as if there is a single object in the scene. The line drawing required two passes through the *L*-subvertex clustering step. All objects are interpreted correctly even if some of their edges are partially occluded. It should be noted that the hidden lines are not removed in displaying the final shape.

4 Conclusion and Future Work

We have seen how the line drawing interpretation problem can be formulated as a constraint satisfaction problem, and how iterative clustering can help solve it under a proper representation. Future work includes an extension of the approach




input line drawing




output shape (shaded view)

Figure 2: Results for the line drawing of a hexagonal object.




input line drawing




output shape (shaded view)


Figure 3: Results for the line drawing of an L-shaped object.







input line drawing



output shape (shaded view)



two views of the last approximate shape



three views of the output shape

Figure 4: Results for the line drawing of multiple objects.

for line drawings of curved objects, in which case the planarity constraint will become the smoothness constraint, and the parallelism constraint will become the parallel symmetry constraint.

References

- [1] M. B. Clowes. On seeing things. *Artificial Intelligence*, 2(1):79–116, 1971.
- [2] D. Huffman. Impossible objects as nonsense sentences. In B. Meltzer and D. Michie, editors, *Machine Intelligence 6*, pages 295–323. Edinburgh University Press, Edinburgh, 1971.
- [3] T. Kanade. Recovery of the three-dimensional shape of an object from a single view. *Artificial Intelligence*, 17:409–460, 1981.
- [4] Y. G. Leclerc and M. A. Fischler. An optimization-based Approach to the interpretation of single line drawings as 3D wire frames. *International Journal of Computer Vision*, 9(2):113–136, 1992.
- [5] S. Y. Lu and K. S. Fu. A Sentence-to-Sentence Clustering Procedure for Pattern Analysis. *IEEE Transactions on Systems, Man and Cybernetics*, 8(5):381–389, May 1978.
- [6] A. Mackworth. Interpreting pictures of polyhedral scenes. *Artificial Intelligence*, 4:121–137, 1973.
- [7] T. Marill. Emulating the human interpretation of line-drawings as three-dimensional objects. *International Journal of Computer Vision*, 6(2):147–161, 1991.
- [8] R. Ohlander, K. Price, and R. Reddy. Picture segmentation by a recursive region splitting method. *Computer Graphics and Image Processing*, 8:313–333, 1978.
- [9] K. Sugihara. *Machine Interpretation of Line Drawings*. MIT Press, 1986.