Reconstruction of a 3D Object From a Main Axis System

D. Kang, M. Masry and H. Lipson *
Sibely School of Mechanical and Aerospace Engineering
Cornell University
Ithaca, USA, 14853
{dk227,mark.masry,hod.lipson}@cornell.edu

Abstract

This paper presents an algorithm for the reconstruction of a three-dimensional object from a single two-dimensional freehand sketch composed of strokes connected at vertices. The proposed algorithm uses the angular distribution of the strokes in the sketch plane to determine one or more orthogonal three-dimensional axis systems whose projection correlates with observed stroke orientations, and then uses these axis systems to calculate a plausible depth for each vertex to reconstruct a 3D object from the sketch. The proposed approach is effective for reconstructing objects that are mostly comprised of orthogonal features, as commonly found in many engineering-oriented sketches. We demonstrate an implementation of the algorithm using Levenberg-Marquardt optimization that permits reconstruction of a typical object with over 100 strokes in interactive time.

Introduction

Freehand sketches – informal drawings of shapes using lines and curves – are often the simplest, most effective way of communicating shape information. A recent proliferation of pen-based hardware and software operating at high sampling rates (Berger 2002) has made it possible for natural sketching to be done with pen-based digital devices. Of particular interest to engineers and designers are 2D line drawings of 3D objects that represent projections of 3D objects onto a 2D viewing plane. Though human observers are able to reconstruct unambiguous 3D shapes from the 2D sketch of a single view without additional information, CAD systems lack this ability entirely.

The process of reconstructing a 3D object from a 2D sketch specified by a set of strokes connected at various vertices requires that depth values be assigned to each vertex. Any arbitrary set of depths assigned to the vertices of the graph constitutes a 3D configuration whose projection will match the given sketch exactly. In principle, each such assignment yields a valid candidate 3D reconstruction.

The reconstruction process itself can be computationally intensive, and many reconstruction solutions do not take into account the inherent consistency of sketches of 3D objects. Strong trends in the angular direction of the 2D strokes that make up a sketch can easily be identified by a human observer, and suggest a considerable amount about the shape of the associated 3D object. This paper presents a fast algorithm that reconstructs a 3D object from freehand line sketch by exploiting these angular trends.

*The authors would like to thank Microsoft for their generous support of this project

Previous work

Several systems (e.g. (Stahovich, Davis, & Schrobe 1998; Davis 2002)) have recently been constructed with sketch-based interfaces. The reconstruction of 3D shapes from 2D sketches has also been the focus of much research. Line labelling approaches (Huffman 1971; Clowes 1971) classify each line as convex, concave or occluding edge without explicitly reconstructing 3D shapes. The methods in (Mackworth 1973; Wei 1987) construct relationships between the slope of sketch lines and the gradients of the associated 3D faces in an attempt to constrain the number of possible interpretations. Interactive methods, in which 3D objects are incrementally constructed by attaching facets sketched by the user in 2D, are specified in (Lamb & Bandopdhay 1990; Fukui 1998). A gesture-based system for interactively constructing 3D rectilinear models is given in (Zeleznik, Herndon, & Hughes 1996); another such interface for the interactive construction of free-form 3D solids is given in (Igarashi, Matsuoka, & Tanaka 1999). Other approaches to the reconstruction problem require the assumption that 3D elements in a scene are specified entirely by known primitives (Wang & Grinstein 1989). Though restrictive, this allows the reconstructed scene to be specified with a convenient solid geometry.

Optimization-based reconstruction methods determine depth values as the solution that optimizes a target function. These methods are more general than the approaches above and can be used to reconstruct relatively complex 3D objects. The relationship of a 2D sketch to an underlying 3D object can be characterised by systems of linear equations for which the existence of solutions is a sufficient criterion for reconstruction. Linear programming optimization techniques may provide these solutions (Sugihara 1986; Grimstead & Martin 1995). Another approach is taken in (Lipson & Shpitlani 1996), in which 2D sketches are converted to line and vertex graphs, which are analyzed for regularities such as parallelism, perpendicularity and symmetry. Regularities in the 2D sketch plane are then weighted according to the probability that they correspond to 3D geometrical relationships, and summed to produce an overall compliance function that estimates how well the 3D construction conforms to the regularities in the 2D sketch. Reconstruction proceeds by optimizing this compliance function. There are also statistical approaches to optimization-based reconstruction (e.g. (Lipson & Shpitlani 2000; 2002)). The correlation between the 2D angles formed by lines in the sketch plane and the angle between these lines in 3D space are learned from a large number of computer-generated 3D shapes and the corresponding projections of these shapes onto a viewing plane. These 2D-3D geometric correlations are then used to determine the most likely 3D shape corresponding to a set of 2D angles, by optimizing over possible assignments of depth values.

While flexible, optimization-based methods suffer from two drawbacks:

1. The optimization surface itself may contain many local minima that make it difficult to determine a global minimum.
2. The computational complexity of the optimization process is generally polynomial in the number of vertices and lines in a sketch.

This paper presents a novel approach to optimization-based reconstruction that avoids these issues in cases where the sketch is relatively structured. The proposed method uses the angular distribution of all strokes in the sketch to determine a three-dimensional main axis that is then used to calculate the missing depth values. The derivative of the optimization function used to determine the main axis can be expressed analytically, allowing the use of fast optimization methods. The sketch's depth values are the determined using the connectivity of the graph specifying the sketch.

Sketch Reconstruction

The reconstruction process takes as input a 2D sketch composed of strokes connected at vertices that represents the 2D orthographic projection of a 3D object onto the plane $z = 0$. The 2D sketch can be interpreted as a connectivity graph with edges specified by the strokes and vertices specified by the sketch vertices. It is assumed that the graph is *connected* e.g. that a path can be constructed from each vertex to every other vertex. It is further assumed that none of the vertices or strokes in the sketch completely obscure other elements of the same kind, and that at least one vertex is connected to three strokes. These assumptions are generally true for all 2D projections of single 3D objects. Since the (x, y) coordinates of each vertex are given in the sketch, reconstructing a 3D object requires assigning a z coordinate (also termed the *depth* value) to each vertex, subject to constraints on the characteristics of the resulting 3D object.

Given a 2D sketch with the criterion specified above, the proposed reconstruction algorithm can be summarised as follows:

1. Select all strokes in the 2D sketch
2. Build a histogram of the angular distribution of the selected strokes in the sketch plane
3. Select the vertex where the angular distribution of the connected strokes has the highest correlation with the histogram of Step 2
4. Assign weights to all strokes as a function of the strokes attached to the selected vertex
5. Generate a maximum weight spanning tree (MST) for the graph and select the strokes making up the MST
6. If the strokes making up the MST differ from those previously selected, go to Step 2
7. Set the depth of the selected vertex to 0
8. Determine the depth of the other endpoints of the strokes attached to the selected vertex to reconstruct an axis system
9. Determine the depth values of all other vertices in the sketch using the reconstruct axis system and the MST

The algorithm has three major components: the construction of an angular distribution histogram of a set of strokes and the selection of a representative vertex (Steps 2 and 3), the assignment of a weights to each stroke and the construction of a maximum spanning tree for the graph (Steps 4 and 5), and the reconstruction of a 3D axis system and the subsequent determination of the depth values of the sketch vertices. Each of these components is discussed in detail below. The problem of reconstructing sketches with more than one primary axis is also addressed.

Many alternative methods of constructing the 3D axis system are possible within the ADG/MST formulation presented above. For example, rather than select an axis vertex as in step 2, an alternative approach (Lipson & Shpitlani 1996) is to construct an independent, unattached 3D axis system at this step using a more general optimzation-based procedure. This approach produces an orthogonal axis system except in degenerate cases, but is more computationally intensive than the algorithm presented in this paper.

Angular Distribution Graphs

Since orthogonality is the prevailing trend in most engineering drawings, and the easiest to identify, a statistical analysis of the direction of strokes in the sketch is performed to determine whether these are consistent with the projections from an underlying orthogonal axis system. The Angular Distribution Graph (ADG) for a set of strokes is a discrete histogram of the 2D angles of the strokes relative to the sketch plane.

Each angle is taken to be the mean of a Gaussian distribution with a fixed variance to reduce sensitivity to noise; the resulting distribution is then sampled and added to the histogram. Peaks in the ADG show prevailing sketch angles. The ADGs of most polyhedral 3D objects have clear peaks. The reconstructed objects axis system should thus have a spatial orientation such that it projects onto the sketch plane at angles corresponding to maxima in the ADG. Figure 1 shows a 2D sketch and the ADG of all strokes.

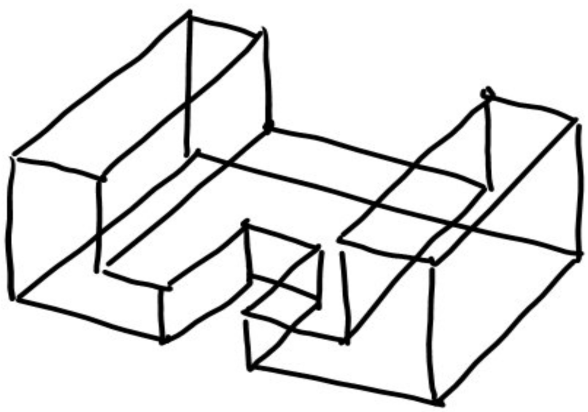

(a)

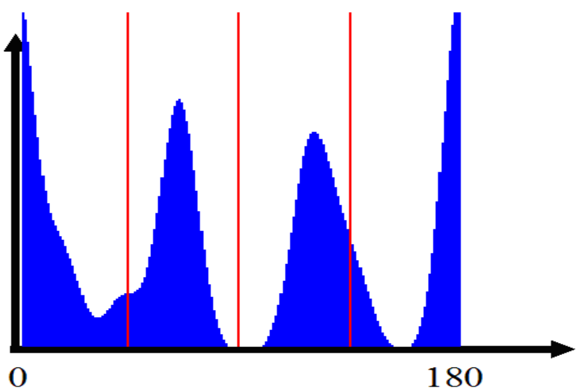

(b)

Figure 1: (a) A 2D sketch with a single distinct axis system (b) its Angular Distribution Graph (ADG)

The local ADG of a vertex is the ADG of all strokes attached to the vertex.The first step in the reconstruction process is to select a vertex whose local ADG is most similar to the ADG of a representative set of strokes. The similarity between any two discrete ADGs can be measured using linear correlation. The vertex whose local ADG has the highest correlation with the ADG of the representative set of strokes is chosen to be the origin of an axis system used to reconstruct the 3D object. The three strokes attached to this vertex represent the projection of the axes onto the sketch plane.

Determining a Maximum Spanning Tree

The depth of each vertex is determined by the projection of the connected strokes onto the main axis system. Given that the graph is connected, it is possible to construct a spanning tree that connects each vertex to the axis origin. Depth values are then propagated along this tree, beginning at the axis endpoints. Since strokes whose direction are not parallel to one of the reconstructed axes cannot be reliably reconstructed, these must be avoided when constructing the spanning tree. The weight assigned to the vector $\mathbf{v}_n = (x_n, y_n)$ underlying stroke n is given by t

$$\max_{\mathbf{v}_a \in \{\mathbf{v}_{a1}, \mathbf{v}_{a2}, \mathbf{v}_{a3}\}} \frac{x_n x_a + y_n y_a}{\sqrt{x_n^2 + y_n^2}\sqrt{x_a^2 + y_a^2}} \quad (1)$$

where $\{\mathbf{v}_{a1}, \mathbf{v}_{a2}, \mathbf{v}_{a3}\}$ are the 2D vectors underlying the strokes in the axis system. A Maximum weight Spanning Tree (MST) constructed with these weights connects all of the vertices in the sketch while avoiding strokes that are least similar to the sketch ADG; using it to determine the depth values of each vertex therefore maximizes the regularity of the reconstructed 3D object.

The maximum weight spanning tree is determined using Prim's algorithm (Cormen, Leiserson, & Rivest 1999). The algorithm begins using only the main axis vertex and the projected axes. The tree is then iteratively expanded by selecting the connected stroke with the highest weight. The representative ADG and MST for the sketch are constructed iteratively (Steps 2 to 5) in order to minimize the effects of atypical strokes. If this process does not converge on a single MST after a several iterations, the 2D sketch does is considered not to have a single underlying axis system. A method for determining multiple underlying axis systems is outlined below. Sketches without one or more underlying axis systems cannot be reconstructed with the proposed method.

Reconstruction of the Main Axis

The origin of the main axis system is assumed to have a depth of zero. The depth of the three attached endpoints (x_1, y_1, y_1), (x_2, y_2, y_2) and (x_3, y_3, y_3) at the end of each axis stroke must be determined in order to reconstruct the main axis system. The x and y values are specified in the sketch. The values of z_1, z_2, and z_3 are determined using an optimization-based approach under the assumption that the axis strokes are perpendicular in 3D space. The optimization cost is a function of z_1, z_2, and z_3:

$$\begin{aligned} & f(z_1, z_2, z_3) \\ & = \cos^2\theta_{21} + \cos^2\theta_{32} + \cos^2\theta_{31} \\ & \quad + \omega\left(\left(r_{21} - \tfrac{L_2}{L_1}\right)^2 + \left(r_{32} - \tfrac{L_3}{L_2}\right)^2 + \left(r_{31} - \tfrac{L_1}{L_3}\right)^2\right) \\ & = \tfrac{\mathbf{p}_1 \cdot \mathbf{p}_2}{|\mathbf{p}_1||\mathbf{p}_2|} + \tfrac{\mathbf{p}_3 \cdot \mathbf{p}_2}{|\mathbf{p}_3||\mathbf{p}_2|} + \tfrac{\mathbf{p}_1 \cdot \mathbf{p}_3}{|\mathbf{p}_1||\mathbf{p}_3|} \\ & \quad + \omega\left(\left(r_{21} - \tfrac{|\mathbf{p}_2|}{|\mathbf{p}_1|}\right)^2 + \left(r_{23} - \tfrac{|\mathbf{p}_2|}{|\mathbf{p}_3|}\right)^2 + \left(r_{31} - \tfrac{|\mathbf{p}_3|}{|\mathbf{p}_1|}\right)^2\right) \end{aligned} \quad (2)$$

where $\mathbf{p}_n$ is a three-dimensional vector representing axis stroke n, L_n is the vector magnitude of $\mathbf{p}_n$, and r_{mn} is the ratio of the length of strokes m and n as measured in the sketch plane. The weighting factor ω allows a tradeoff between the angular and length constraints.

Since the vectors in the main axis are assumed to be perpendicular to one another in 3D space, the angle between them should be $90°$; the difference between the ratio of their lengths in the sketch plane and the ratio of their lengthen 3D space should be 0. The optimization goal is therefore to minimize the cost function $f(z_1, z_2, z_3)$. The minimizing parameters are given by the solution to the system equation

$$(\mathbf{J^T J} + \lambda \mathbf{I})\delta\mathbf{x} = \mathbf{J^T}e \quad (3)$$

where $\mathbf{I}$ is the identity matrix, λ is an optimization parameter and $\mathbf{J}$ is the Jacobian matrix given by the partial derivatives of the error equation about z_1, z_2, and z_3, respectively. The Jacobian matrix is given by

$$\mathbf{J} = \left[\frac{\delta f}{\delta z_1}, \frac{\delta f}{\delta z_2}, \frac{\delta f}{\delta z_3}\right] \tag{4}$$

The partial derivative $\frac{\delta f}{\delta z_1}$, for example, is given by

$$\begin{aligned} \frac{\delta f}{\delta z_1} &= -2\left(\frac{\mathbf{p}_1\cdot\mathbf{p}_2}{|\mathbf{p}_1||\mathbf{p}_2|}\right)\left(\frac{z_2}{|\mathbf{p}_1||\mathbf{p}_2|} + \frac{\mathbf{p}_1\cdot\mathbf{p}_2}{|\mathbf{p}_2|}\right)\left(\frac{z_1}{|\mathbf{p}_1|^3}\right) \\ &\quad - 2\left(\frac{\mathbf{p}_3\cdot\mathbf{p}_1}{|\mathbf{p}_3||\mathbf{p}_1|}\right)\left(\frac{z_3}{|\mathbf{p}_3||\mathbf{p}_1|} + \frac{\mathbf{p}_3\cdot\mathbf{p}_1}{|\mathbf{p}_3|}\right)\left(\frac{z_1}{|\mathbf{p}_1|^3}\right) \\ &\quad + 2\omega\left(r_1 - \frac{|\mathbf{p}_2|}{|\mathbf{p}_1|}\right)\left(\frac{z_1|\mathbf{p}_2|}{|\mathbf{p}_1|^3}\right) \\ &\quad + 2\omega\left(r_3 - \frac{|\mathbf{p}_1|}{|\mathbf{p}_3|}\right)\left(\frac{z_1}{|\mathbf{p}_1||\mathbf{p}_2|}\right) \end{aligned} \tag{5}$$

The Levenberg-Marquardt method (Heath 2002) yields a fast solution to this nonlinear optimization problem. The method is an iterative variation of the Newton method in nonlinear estimation. The value λ is initialized to a small value. If the value of e obtained for a particular $\delta\mathbf{x}$ reduces the error, $\mathbf{x}$ is incremented to $\mathbf{x} + \delta\mathbf{x}$ and λ is divided by 2. On the other hand, if the e increases, λ is multiplied by 2 and the augmented normal equations are solved again until an increment that reduces the error is obtained.

Vertex Depth Calculation

The depth of each vertex in the sketch is determined by the propagating the depth of the endpoints of the axis strokes to each connected vertex, using the reconstructed axis system and the MST. The value z_p of a vertex (x_p, y_p, z_p) that is connected to an already reconstructed vertex in the MST is determined by first selecting the axis whose projection $\mathbf{v}_a$ maximizes the 2D projection $\frac{x_n x_a + y_n y_a}{\sqrt{x_n^2+y_n^2}\sqrt{x_a^2+y_a^2}}$ and then maximizing solving the 3D projection

$$1 = \frac{x_a x_p + y_a y_p + z_a z_p}{\sqrt{x_a^2 + y_a^2 + z_a^2}\sqrt{x_p^2 + y_p^2 + z_p^2}} \tag{6}$$

for z_p. This equation can also be expressed as a second order polynomial of the form $Az_p^2 + Bz_p + C = 0$. In cases where the 2D projection of the vector onto the chosen axis is not equal to 1, no analytic solution will exist; the reconstructed depth is the value that minimizes the second order polynomial equation. The above computation is then recursively applied as specified by the connectivity of the MST.

Multiple Axis Sketches

The reconstruction algorithm described above assumes that the sketch has a single underlying axis system. If the sketch was drawn using two or more distinct axis systems, however, these must be handled independently. A separate axis selection step was therefore added to the reconstruction algorithm. In the modified algorithm, the axis system determined in Step 3 is used to construct a local MST of connected strokes rooted at the axis origin. Connected strokes are added to the MST if the cumulative ADG of the axis system and the candidate stroke does not contain more than

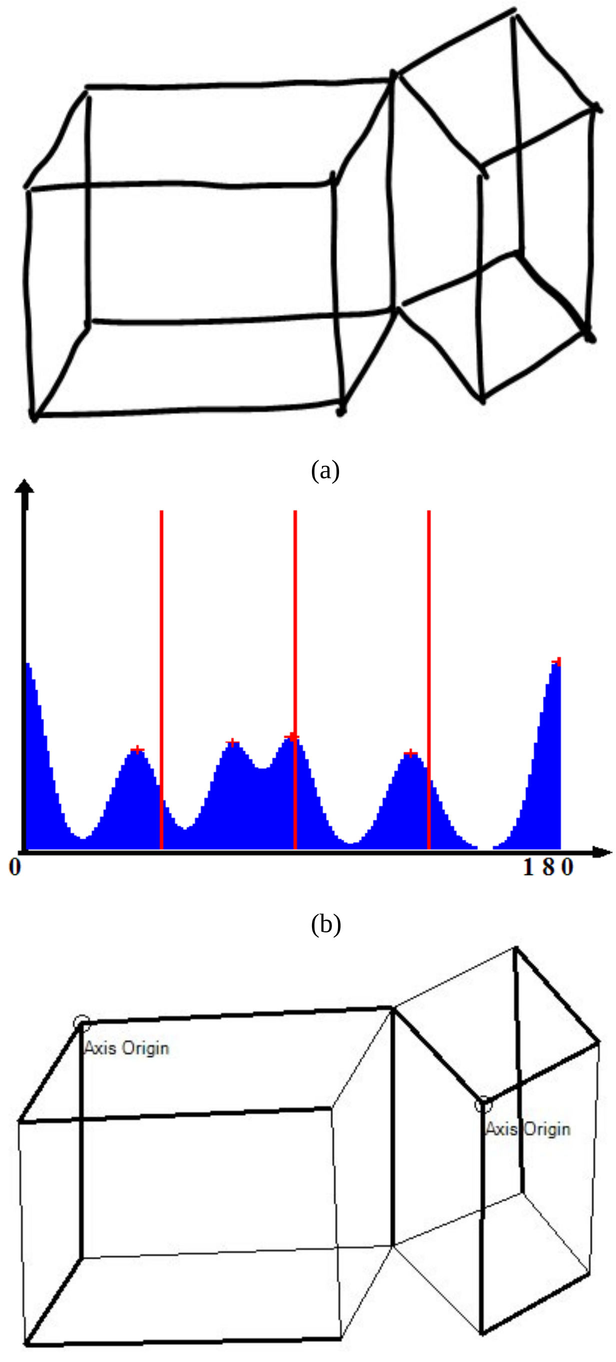

Figure 2: (a) A 2D sketch with two distinct axis systems and (b) its Angular Distribution Graph (ADG) (c) The resulting Maximum Spanning Trees (shown in bold) and axis origins

three peaks. The set of strokes in the MST and disjoint set of all other strokes are are recursively split until each vertex is assigned to at one or more local MSTs.

Figure 2 (a) shows a sketch with two axis systems. The sketch ADG shown in Figure 2 (b) has five distinct peaks, as opposed to the three found in the ADGs of sketches with

ω	0	0.001	0.01	0.1	1
$\mu_{\theta_{mn}}$	$89.90°$	$89.93°$	$88.40°$	$88.3°$	$86.10°$
μ_{length}	1.75	1.53	1.32	1.15	1.03

Table 1: Mean axis angular error ($\mu_{\theta_{mn}}$) and mean length ratios (μ_{length}) as a function of ω measured over all three axes for two axes reconstructions.

a single axis system such as Figure 1. Figure 2 (c) shows the two axis systems isolated by the algorithm, and the the local MST attached to each one. All vertices belonging to a single MST are reconstructed separately, then translated so that they are aligned at vertices belonging to more than one local MST.

Results

The characteristics of the reconstructed shape are determined by those of the main axis systems. An analysis of the reconstruction of a single main axis with coordinates $(100, 0, z_1)$, $(0, 90, z_2)$, $(20, 20, z_3)$ is shown in Figure 3. Small random values are initially assigned to z_1, z_2, z_3, and all constants in Equation 2 are initially set to 1. The reconstructed values are $z_1 = -17.52$, $z_2 = -7.37$ and $z_3 = 112.98$, respectively. The final value of the optimization function was very close to 0, indicating that the reconstructed axes were perpendicular, and the ratio of their projected lengths was equal to the ratio of their lengths. Convergence occurred after approximately 25 iterations.

Table 1 shows the mean angle between the reconstructed axes (in degrees) and the mean ratio between the longest and shortest reconstructed axes as a function of the weighting term ω, measured over the reconstruction of two different main axis systems with $r_n = 1 \forall n$ in both cases. The relative importance of length ratios in the optimization function increases as ω increases, while as the relative importance of the angular terms decreases. The mean length ratio therefore decreases to 1 as ω increases, while the axes shift away from perpendicularity. Because of its effect on the relative length ratios, the value of ω can be used to determine the elongation of the reconstructed 3D shape, and should generally be defined by the user.

Example 2D reconstructions are shown in Figure 4. The reconstruction process for these shapes finished in approximately 0.1 sec on a Pentium-IV tablet PC. Sketches of several hundred strokes were reconstructed in less than 1 second. The proposed algorithm is suited to all sketches composed of strokes with limited deviations from the underlying axis systems. Further work will center on the development of algorithms for reconstructing plausible shapes when strokes exhibit large deviations from the underlying axis systems, as well as the reconstruction of curved strokes, and strokes not attached to the main object.

Conclusions

This paper has presented a fast algorithm for reconstructing a 3D object from a 2D sketch, assuming that the object was constructed with an underlying axis system. The algorithm uses the angular distribution of the strokes in the sketch to reconstruct a 3D axis system given by a three-connected sketch vertex and the attached strokes. The connectivity of the strokes in the sketch is then used to assign depth values to all of the sketch vertices. A method for determining multiple axis systems was also presented. The algorithm can reconstruct complex sketches of 100 or more strokes in near interactive time. Further research will center on extensions of the algorithm to irregular objects and curved lines.

References

Berger, M. 2002. Pen-based programs debut with tablet pc. *PC World* 7.

Clowes, M. B. 1971. On seeing things. *Artificial Intelligence* 2(1):79–112.

Cormen, T. H.; Leiserson, C. E.; and Rivest, R. L. 1999. *Introduction to Algorithms*. MIT Press.

Davis, R. 2002. Position statement and overview: Sketch recognition at MIT. In *AAAI Sketch Understanding Symposium*.

Fukui, Y. 1998. Input method of boundary solid by sketching. *Computer aided design* 20(8):434–440.

Grimstead, I. J., and Martin, R. R. 1995. Creating solid mdodels from single 2D sketches. In *Solid Modeling '95*, 323–337.

Heath, M. T. 2002. *Scientific Computing, An Introductory Survey, Second Edition*. New York, NY: McGraw-Hill.

Huffman, D. A. 1971. Impossible objects as nonsense sentences. In Meltzer, B., and Mitchie, D., eds., *Machine Intelligence*. Edinburgh: Edinburgh University Press.

Igarashi, T.; Matsuoka, S.; and Tanaka, H. 1999. Teddy: sketching interface for 3d freeform design. In *Proceedings of SIGGRAPH*, 409–416.

Lamb, D., and Bandopdhay, A. 1990. Interpreting a 3D object from a rough 2D line drawing. In *Proceedings of Visualization '90*, 59–66.

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

Lipson, H., and Shpitlani, M. 2000. Conceptual design and analysis by sketching. *Journal of AI in Design and Manufacturing (AIEDEM)* 14:391–401.

Lipson, H., and Shpitlani, M. 2002. Correlation-based reconstruction of a 3D object from a single freehand sketch. In *AAAI Spring Sumposium on Sketch Understanding*, 99–104.

Mackworth, A. K. 1973. Interpreting pictures of polyhedral scenes. *Artificial Intelligence* 4:121–137.

Stahovich, T. F.; Davis, R.; and Schrobe, J. 1998. Generating multiple new designs from a sketch. *Artificial Intelligence* 104:211–264.

Sugihara, K. 1986. *The interpretation of line drawings*. MIT Press.

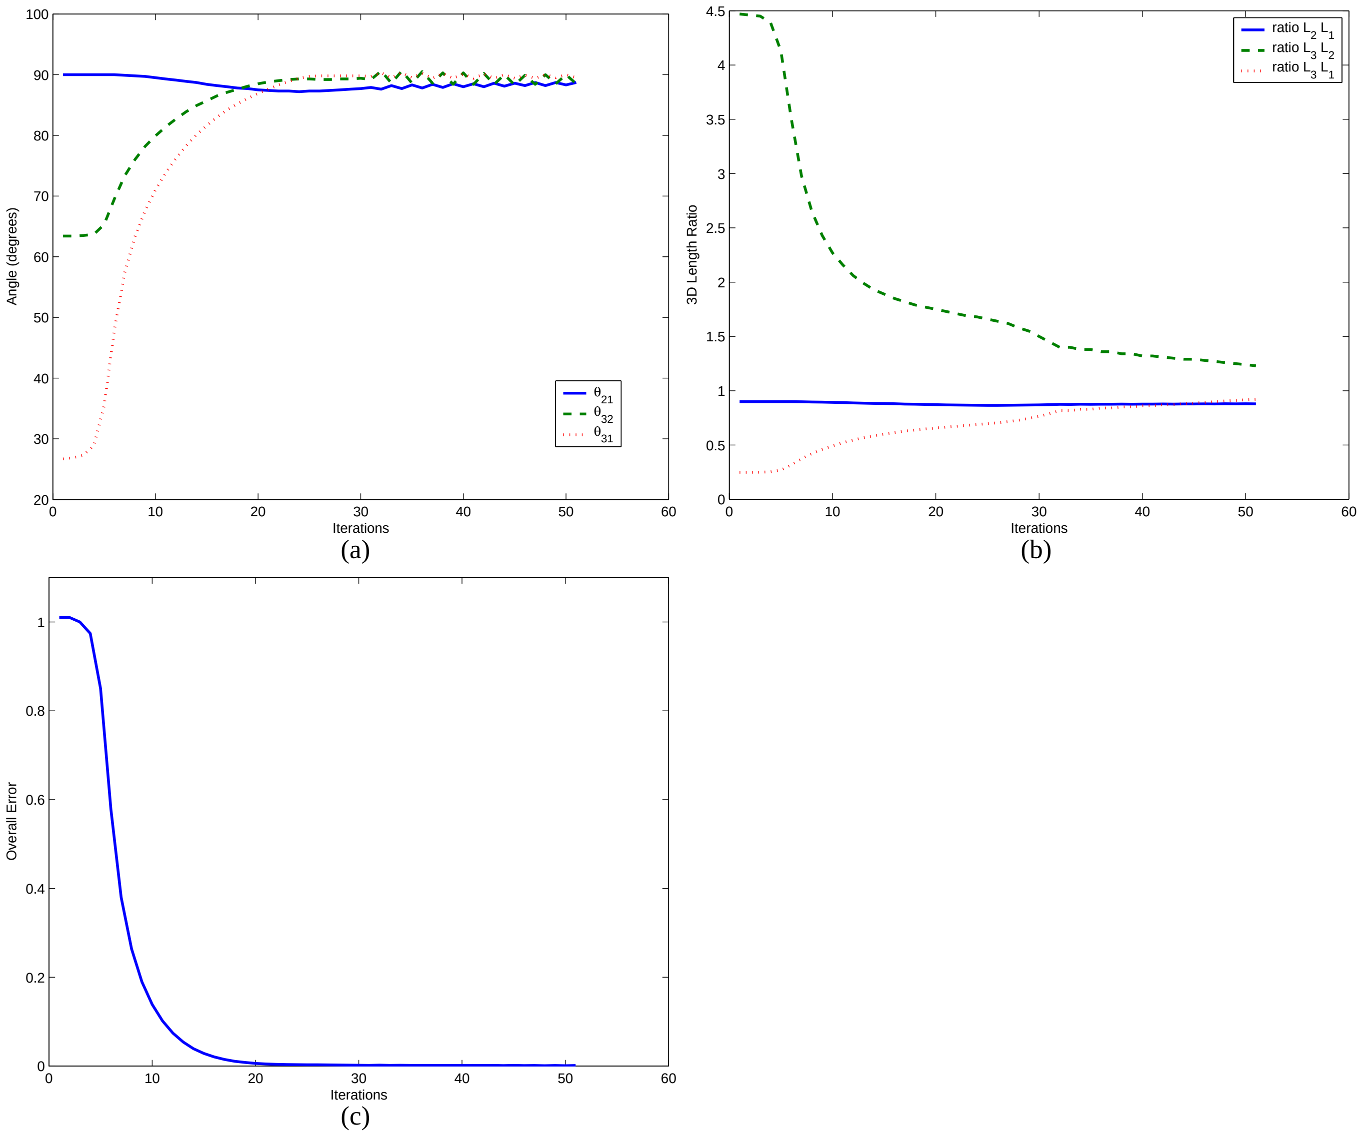

Figure 3: Reconstruction of the main axis system. (a) Change of angle value between any two axes of main axis system during iteration; (b) Change of length ratio values; (c) Deviation error value during iteration. Successful convergence is shown.

Wang, W., and Grinstein, G. 1989. A polyhedral object's CSG-Rep reconstruction from a single 2D line drawing. In *Proceedings of 1989 SPIE Intelligent Robots and Computer Vision III: Algorithms and Techniques*, volume 1192, 230–238.

Wei, X. 1987. *Computer Vision Method for 3D Quantitative Reconstruction from a Single Line Drawing*. Ph.D. Dissertation, Department of Mathematics, Beijing University, China. (in Chinese; for a review in English see Wang and Grinstein, 1993).

Zeleznik, R.; Herndon, K.; and Hughes, J. 1996. SKETCH: An interface for sketching 3d scenes. In *Proceedings of SIGGRAPH*, 163–170.

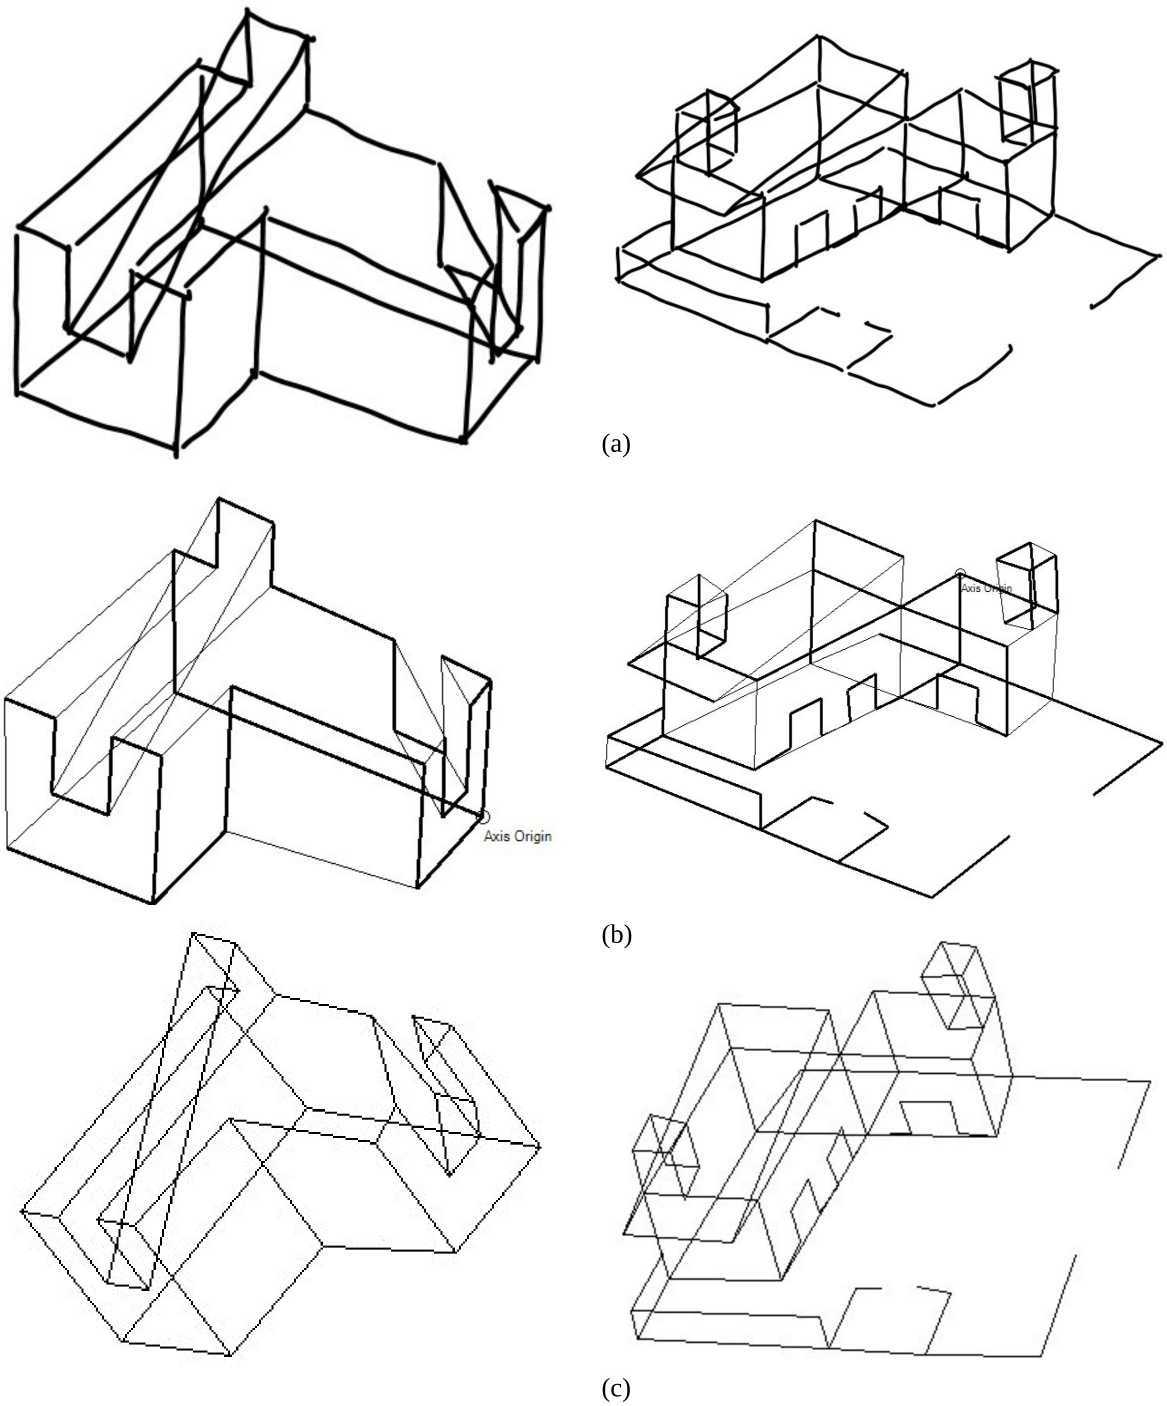

Figure 4: Reconstruction of two sketches. (a) Free hand sketches(b) Extraction of main axis systems and Maximum Spanning Trees (MST); A small circle shows the origin of the main axis system located on the most prominent vertex junction. The MST is shown in bold strokes. (c) Rotated view s of the original shapes