

A Heuristic Method for Generating 2D CSG Trees from Bitmaps

Sarah F. Gibson, Joe Marks, Danielle Feinberg, Manuel Sosa

TR95-19 December 1995

Abstract

This paper presents a method for generating compact and effective constructive-solid-geometry (CSG) object representations from 2D bitmap representations. The method consists of two parts: a stochastic search procedure for finding candidate CSG trees and a local optimization procedure that modifies the primitives of a given CSG tree to effect a better match. Results for several sample input bitmaps are presented and an application of this method to automatic geometric morphing between pairs of bitmaps is shown.

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

1. First printing, TR95-19, January 1996.

1 Introduction

1.1 Motivation

The generation of object models for graphics and animation applications typically requires a major investment of resources. While CAD tools for building geometric models exist, developing models of even simple objects often requires sophisticated design skills and a significant amount of time. On the other hand, new scanning technologies, including laser range scanners and industrial CT scanners, can produce volumetric models of existing objects reasonably quickly and inexpensively. Unfortunately, this volumetric format cannot be used by standard design or manufacturing tools. Hence, methods that convert volumetric data into a structured graphics format are needed in order to exploit these new scanning technologies.

There are several methods that can be used to convert volumetric or scanned-surface data into compact or structured surface models (*e.g.*, [2, 3, 7]). However, surface models do not support all modeling and editing tasks equally well. Some tasks are more appropriately handled by true solid modeling, as embodied in constructive-solid-geometry (CSG). There is therefore a need for techniques that convert scanned data to CSG models.

Shapiro and Vossler [11, 12, 13] have addressed a related problem. They present a practical system that converts limited boundary representations into CSG models. Their method uses a greedy heuristic to construct a compact CSG representation from a given set of half-spaces, those bounded by lines and circular arcs. However, their system does not take volumetric or scanned data as input: it requires that the candidate half-spaces be provided by the user. Furthermore, it is not clear how their method can be extended to incorporate additional geometric primitives. Other related problems have been addressed in computer vision [1, 4, 14] but they are not readily applied to the general scanned-data-to-CSG conversion problem.

1.2 New contributions

This paper addresses the direct conversion from scanned data to an effective CSG model by solving the 2D version of the problem: given a scanned 2D target bitmap, generate CSG trees that are compact, easily modifiable, and model the target shape well.

Our algorithm consists of two cooperating processes: 1) A global search procedure that looks in the space of all possible CSG trees for structures that match the target bitmap well; and 2) A local optimization procedure that modifies the primitives of a given CSG tree to better match the bitmap.

The major contributions of this work are:

- *A complete solution.* Our system takes raw scanned data — a bitmap — as input, and produces a compact CSG model using rectangular and circular primitives.
- *A flexible measure of model goodness.* Our evaluation metric for candidate CSG trees provides a quantitative measure of how well a tree models a target bitmap. Furthermore, our use of a stochastic search strategy permits this measure to be flexible; it can

be easily modified to reflect many factors including model accuracy, compactness, locality (proximity of primitives in the CSG tree should reflect proximity in the object), and appropriateness for design purposes.

- *Variety of solutions.* Unlike constructive heuristics, the proposed stochastic search strategy is capable of finding a variety of good solutions, some of which may be more useful than others for particular editing tasks.
- *Addition of new primitives.* Any shape that can be efficiently scan-converted can be added to the set of primitives elements at the cost of a larger search space.

In addition, as an example application of this work we show how bitmap-to-CSG conversion can be used to morph between pairs of bitmaps (§ 3.2).

2 Technical Approach

2.1 Global search

The global search strategy is outlined in Figure 1. The search begins with a set of random candidate solutions; hill climbing is performed on the candidate solutions in parallel; and at regular intervals the search is refocused on the more promising solutions by throwing away less promising ones. The CSG primitives consist of rectangles and circles, whose parameters — $\langle(x, y), (w, h), \theta\rangle$, indicating a rectangle’s position, size, and orientation, and $\langle(x, y), r\rangle$ indicating a circle’s position and radius — are chosen randomly from the following uniform distributions: $x \in [\frac{W}{4}, \frac{3W}{4}]$, $y \in [\frac{H}{4}, \frac{3H}{4}]$, $r \in [1, \frac{\min(W, H)}{2}]$, $w \in [1, \frac{W}{2}]$, $h \in [1, \frac{H}{2}]$, $\theta \in [0, 360]$, where W and H are the width and height of the bitmap, respectively. The type of each primitive, *i.e.*, rectangle or circle, is chosen with equal probability. Likewise, CSG operators are chosen with equal probability from the set: difference, intersection, and union. The number of primitives in a tree is determined by the function $\text{norm}(a, s)$, which returns normally distributed random integers with average a and standard deviation s . Values for these and other global search parameters are listed in Table 2.

The key elements of the global search — those that most affect performance — are the perturbation operators used to generate new candidate solutions from existing ones. The operators and their probability of application are summarized in Table 1. The REP operator replaces an existing subtree containing e primitive objects with a new random subtree containing $\text{norm}(e, s')$ primitives. The CUT operator deletes a randomly chosen subtree, and the MOD operator changes the type of a randomly chosen primitive from a rectangle to a circle, or vice versa.

ADD is the most complicated operator: half the time it adds a newly generated subtree by unioning it with an existing subtree; otherwise it adds a new subtree by differencing it from an existing subtree. In both cases the new subtree is small, containing $\text{norm}(a', s')$ primitives. In contrast to the behavior of the initialization procedure and the REP operator, the ADD operator locates primitives at specific positions of the target bitmap. These positions are the local maxima or minima (depending on whether they make a net addition or subtraction to the resultant CSG object) of a Euclidean distance map [5] derived from the bitmap. The

```

Initialize  $S$  to be a set of random CSG trees,
  each comprising  $\text{norm}(a, s)$  primitive objects;
Locally optimize and evaluate each tree in  $S$ ;
 $i \leftarrow 0$ ;
 $j \leftarrow 0$ ;
while  $i < N$  and  $j < B$  do
{
  Perturb, locally optimize, and evaluate
  a randomly selected tree in  $S$ ;
  if perturbed tree is better then
  {
    Replace original tree with perturbed tree;
     $j \leftarrow 0$ ;
  }
  else  $j \leftarrow j + 1$ ;
  if  $i \bmod n = 0$  then
  {
    Sort  $S$  by tree score;
    Replace bottom half of  $S$  with top half;
  }
   $i \leftarrow i + 1$ ;
}
Return best tree in  $S$ ;

```

Figure 1: Global search by parallel hill climbing.

<i>Perturbation</i>	<i>Probability</i>
REP: Replace a random subtree with a newly initialized one.	0.3
ADD: Add a targeted subtree at a randomly chosen location.	0.3
CUT: Delete a randomly chosen subtree.	0.3
MOD: Change the type of a randomly selected primitive.	0.1

Table 1: Perturbation operators for global search.

primitive size parameters — width and height for rectangles, radius for circles — are given by the expression $\max(1, \text{norm}(d, d))$, where d is the value of the Euclidean distance map at the relevant local maximum or minimum and represents the distance to the nearest edge in the bitmap.

2.2 Local optimization

Once a CSG tree has been generated, it is passed to the local optimizer, which adjusts the parameters of each primitive to improve the match between the CSG tree and the target bitmap. In order to measure the goodness of this match, primitive elements of the CSG tree are scan-converted and combined using the node operators to generate a bitmap representation of the CSG tree. As illustrated in Figure 2, the mismatch between the target bitmap and the CSG bitmap is represented by the sum of the UNCOVERED region — pixels in the target bitmap that are not covered by the CSG tree — and the EXCESS region — pixels in the scan-converted CSG tree that do not lie in the target bitmap.

Figure 3 outlines the general approach taken in the local optimizer. Adjustments are made to each primitive in the CSG tree depending on the amount and distribution of the EXCESS and UNCOVERED regions local to the primitive. For example, the position of the primitive is shifted until it lies roughly centered between the local EXCESS and UNCOVERED regions. This is done by shifting the primitive in small steps towards the center of the local UNCOVERED region and away from the center of the local EXCESS region as illustrated in Figure 4a.

The dimensions of the primitive are adjusted until the amount of local EXCESS is roughly equal to the amount of local UNCOVERED. Each dimension of the object (radius for circles, height and width for rectangles) is adjusted independently. As illustrated in Figure 4b, if the amount of local UNCOVERED exceeds the amount of local EXCESS, the primitive dimension is increased. Otherwise if the local EXCESS exceeds the local UNCOVERED, the primitive dimension is reduced.

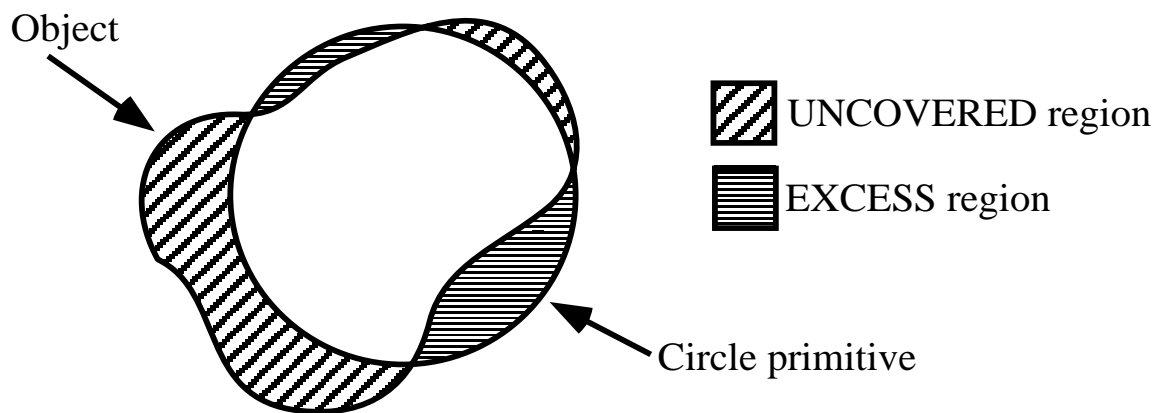


Figure 2: Excess and uncovered regions.

Adjusting the rotation of the primitive using the local EXCESS and UNCOVERED regions is somewhat more complicated. The general approach involves rotating the primitive until its principal axis is aligned with the principal axis of local pixels of the target bitmap. This is accomplished by rotating the primitive away from the principal axis of the local EXCESS region and towards the principal axis of the local UNCOVERED region. The principal axis of a region is determined by the principal eigenvector of the region's inertial tensor matrix. Although this method is valid for both 2D and 3D systems, it is computationally expensive and can be subject to instability. Hence, in the 2D system, a simpler method is used: the primitive is rotated by small steps in a positive direction while the score improves; if the score does not initially improve, the primitive is rotated in a negative direction.

Once the parameters of each primitive have been locally optimized, primitives that do not contribute to the final CSG model are pruned from the tree. The tree score is then computed as the sum of the number of EXCESS pixels, the number of UNCOVERED pixels, and a scalar weight, k , times the number of primitives in the CSG tree. (Including the number of primitives in the score promotes compact trees.) This score is then returned to the global search.

3 Results

3.1 Bitmaps and their CSG trees

Figures 5, 8, and 10 contain hand-drawn bitmaps that served as some of our test input. The algorithm described in § 2 was run five times for each bitmap using different random-number seeds. The CSG trees shown below were selected from the five final trees generated for each bitmap. For each tree we report the total number of different CSG trees that were evaluated during its run before arriving at the result shown: this number always consists of 500 initial random trees, plus some variable number of trees produced by application of the


```

 $P \leftarrow$  Number of primitives in CSG tree;
Generate a bitmap representation of the CSG tree;
Determine EXCESS and UNCOVERED regions;
Calculate the score;
 $j \leftarrow 0$ ;
while  $j < M$  and score is improving do
{
  for  $i \leftarrow 0$  to  $P - 1$  do
  {
    Adjust position of primitive  $i$ ;
    Adjust dimensions of primitive  $i$ ;
    Adjust orientation of primitive  $i$ ;
  }
  Recalculate the score;
   $j \leftarrow j + 1$ ;
}
Prune the CSG tree;
Return the optimized tree and its score;

```

Figure 3: Local optimization of primitive parameters.

perturbation operators.¹

As with many heuristic search methods, this approach requires setting many parameters. The parameter values that were used in the examples reported here were derived by experimentation and are listed in Tables 1 and 2.

Figures 6 and 7 show two CSG trees for the bitmap depicted in Figure 5. The former was found after evaluating 2,096 trees; the latter after evaluating 2,787 trees. The tree in Figure 6 contains a minimal number of primitives and is preferable for editing operations that will preserve the object's symmetry. The tree in Figure 7 contains one more primitive, but is preferable for asymmetric editing. These examples illustrate the utility of an algorithm that is capable of producing a variety of near-optimal results for the scan-and-edit applications we have in mind.

The bitmap in Figure 8 is one of those used for the morphing tests described in § 3.2. A corresponding CSG tree, the product of 2,824 tree evaluations, is shown in Figure 9. After 400 iterations of the search procedure, the likelihood that each perturbation operator (Table 1) would yield an improvement was fairly high: MOD, 0.49; REP, 0.33; ADD, 0.48; and CUT, 0.26. However, by the time the best tree was found, these numbers had declined significantly: MOD, 0.15; REP, 0.12; ADD, 0.10; and CUT, 0.09. This behavior was

¹An alternative measure of computational effort is the number of primitives that were scan-converted. A maximum of 400,000 primitive scan-conversions is performed in the local optimization and evaluation of a 10-primitive CSG tree; the average number is somewhat less. Multiplying the maximum number by the number of trees evaluated gives an upper bound on the number of primitive scan-conversions performed.




Figure 4: Position and dimension adjustments.




Figure 5: An example bitmap.




Figure 8: Bitmap for the letter 'P'.

typical of the other tests too.

Finally, the bitmap shown in Figure 10 was inspired by the familiar icon on the back of the U.S. \$1 bill. The results for this bitmap are shown in Figures 11 and 12. The former required 3,074 CSG-tree evaluations; the latter, 4,485. The two candidate trees further illustrate the diversity of solutions that can be generated by the stochastic search method.

3.2 Application to geometric morphing

To demonstrate the capability and flexibility of the approach described here, the algorithm was modified to perform geometric morphing between pairs of bitmaps. The modifications are straightforward:

- Two target bitmaps are used.
- Each candidate solution consists of two CSG trees that are structurally identical, but that can have different primitive objects.
- Each tree is locally optimized with respect to its corresponding target bitmap, and the score given to a candidate solution is the average score for its two trees.
- Structural identity between the two trees is maintained by having the initial trees be identical copies, and by applying the same REP, ADD, and CUT operations to both trees. However, MOD operations and local optimization are applied independently to each tree, and no pruning is done by the local optimization procedure.

The solution returned by the modified algorithm consists of two trees that are structurally identical, but each of which is a good match for one of the two target bitmaps by virtue of

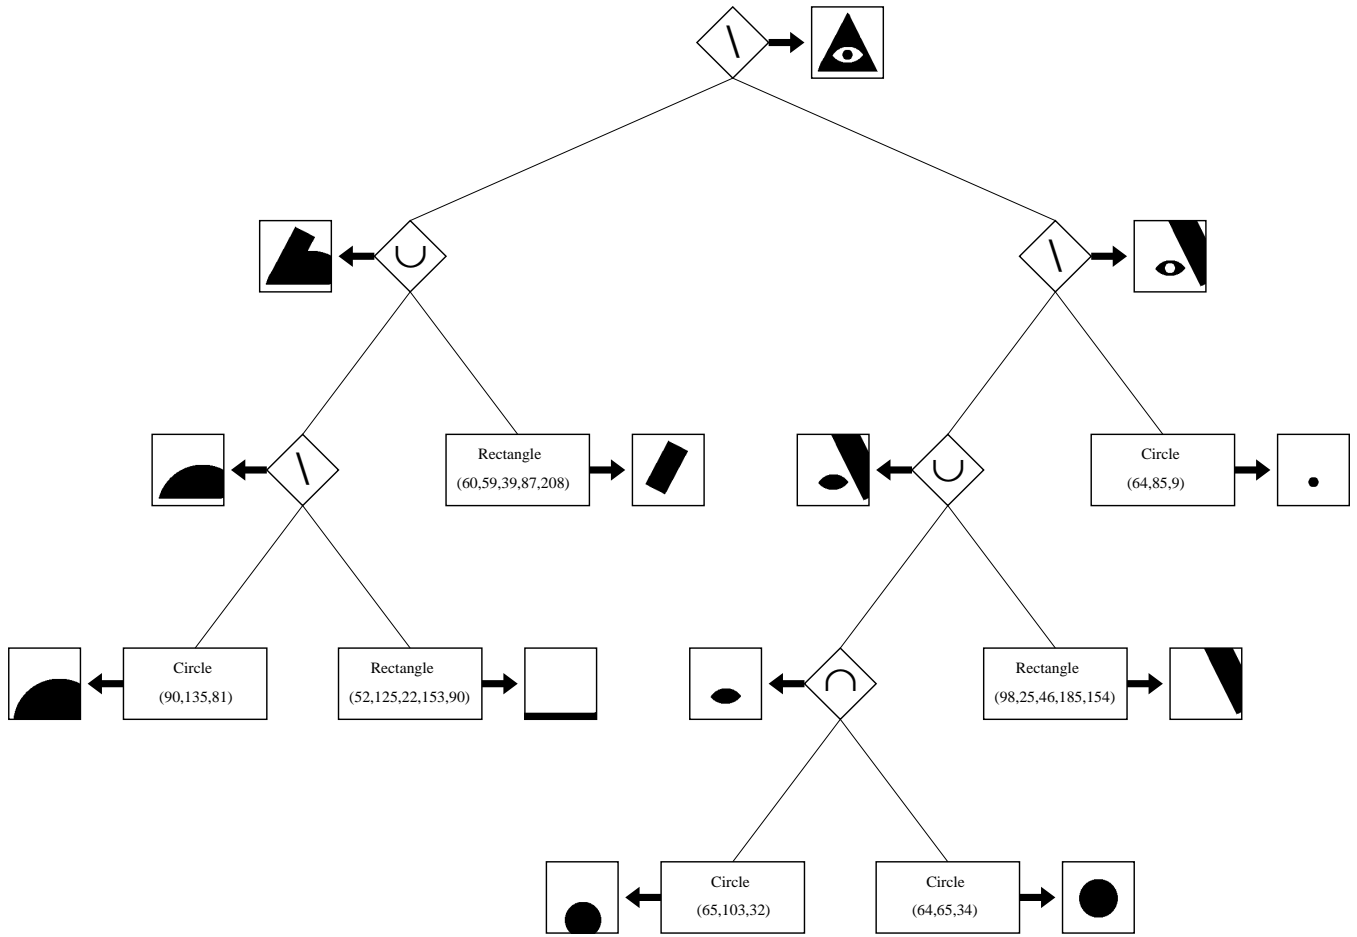


Figure 11: CSG tree for the eye-in-diamond icon.

possible differences between the primitive objects in each tree. Morphing is accomplished by straightforward interpolation of the corresponding primitives in both trees.²

Sample morphs are shown on the accompanying videotape. The algorithm parameters used were identical to those reported in Table 2, with one exception: in the morphing runs, for which simpler shapes with smaller expected tree sizes were used, the average initial tree size, a , was set to 5.0. Each pair of bitmaps shown depicts two topologically different letters. This approach is different from previous work on 2D geometric morphing [8, 9, 10] in both its ability to morph between topologically different shapes and in the style of morphs that it produces.

4 Conclusions and Future Work

In this paper we have presented a complete system for 2D bitmap-to-CSG conversion. Before attempting to solve the general problem of converting 3D point data to a useful, equivalent CSG representation, a better grasp of the 2D problem is still required. We plan to explore the following extensions to our current approach:

- *A greater variety of primitives.* One advantage of our system is its potential for accommodating different primitive objects. If a primitive can be scan-converted efficiently, it can be incorporated straightforwardly. For example, triangles and ellipses are obvious candidates for inclusion. However, the added benefit of having extra primitives will be offset somewhat by the increase in search-space dimensionality.
- *Faster scan-conversion.* The dominant computational cost in our current system is due to scan-conversion of primitive shapes. Taking advantage of scan-conversion hardware on high-end graphics workstations will increase the capacity for further experimentation and the ability to explore bigger search spaces.
- *More intelligent searching.* The use of a Euclidean distance map to generate targeted subtrees in the ADD operator (§ 2.1) caused a significant improvement in performance. We anticipate similar improvements from additional heuristics for proposing and positioning primitive objects. The analysis of Shapiro and Vossler [11, 12, 13] is a possible source of relevant ideas. However, care must be taken when developing such heuristics: the benefits of a more focused search can be outweighed by a decrease in the variety of solutions found.
- *Different objective criteria.* Compactness alone is not a sufficient criterion for identifying CSG models that are easy to modify and edit. Locality is doubtless an important consideration, and additional criteria may be derived from a study of typical model-editing tasks [6].

²Interpolating between two orientations (which is relevant for rectangles only) can be done either clockwise or counterclockwise. For the morphs on the video we chose the directions of rotation for effect.

5 Acknowledgments

Thanks to Bill Freeman, Ted Gibson, and Stuart Shieber for helpful discussions, to Ben Bernard for initial exploratory work, and to Denny Bromley for audio-visual assistance.

References

- [1] S. J. Dickinson, A. P. Pentland, and A. Rosenfeld. 3-D shape recovery using distributed aspect matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), Feb. 1992.
- [2] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle. Piecewise smooth surface reconstruction. In *Proc. of SIGGRAPH 94*, pages 295–302, Orlando, Florida, July 1994. ACM SIGGRAPH, New York. In *Computer Graphics Annual Conference Series*, 1994.
- [3] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Mesh optimization. In *Proc. of SIGGRAPH 93*, pages 19–26, Anaheim, California, Aug. 1993. ACM SIGGRAPH, New York. In *Computer Graphics Annual Conference Series*, 1993.
- [4] Y. G. Leclerc. Constructing simple stable descriptions for image partitioning. *International Journal of Computer Vision*, 3(1):73–102, 1989.
- [5] J. C. Russ. *The Image Processing Handbook*. CRC Press, Boca Raton, 1992.
- [6] E. Saund and T. P. Moran. Perceptual organization in an interactive sketch editing application. In *Fifth International Conference on Computer Vision*, pages 597–604, Cambridge, Massachusetts, June 1995. IEEE Computer Society Press.
- [7] W. J. Schroeder, J. A. Zarge, and W. E. Lorensen. Decimation of triangle meshes. In *Proc. of SIGGRAPH 92*, pages 65–70, Chicago, Illinois, July 1992. ACM SIGGRAPH, New York. In *Computer Graphics* 26(2), July 1992.
- [8] T. W. Sederberg, P. Gao, G. Wang, and H. Mu. 2-D shape blending: An intrinsic solution to the vertex path problem. In *Proc. of SIGGRAPH 93*, pages 15–18, Anaheim, California, Aug. 1993. ACM SIGGRAPH, New York. In *Computer Graphics Annual Conference Series*, 1993.
- [9] T. W. Sederberg and E. Greenwood. A physically based approach to 2-D shape blending. In *Proc. of SIGGRAPH 92*, pages 25–34, Chicago, Illinois, July 1992. ACM SIGGRAPH, New York. In *Computer Graphics* 26(2), July 1992.
- [10] M. Shapira and A. Rappoport. Shape blending using the star-skeleton representation. *IEEE Computer Graphics and Applications*, pages 44–50, Mar. 1995.
- [11] V. Shapiro and D. L. Vossler. Construction and optimization of CSG representations. *Computer-Aided Design*, 23(1):4–20, Jan./Feb. 1991.

- [12] V. Shapiro and D. L. Vossler. Efficient CSG representation of two-dimensional solids. *Transactions of the ASME: Journal of Mechanical Design*, 113(3):292–305, Sept. 1991.
- [13] V. Shapiro and D. L. Vossler. Separation for boundary to CSG conversion. *ACM Transactions on Graphics*, 12(1):35–55, Jan. 1993.
- [14] L. R. Williams. Perceptual organization of occluding contours. In *Third International Conference on Computer Vision*, pages 133–137, Osaka, Japan, December 1990. IEEE Computer Society Press.