Digital Topology and Morphology

Segmenting 3D Objects into Geons

Kenong Wu and Martin D. Levine

Center for Intelligent Machines & Dept. of Electrical Engineering
McGill University, Montréal, Québec, Canada, H3A 2A7

Abstract. We propose a new approach to 3D object segmentation and description. Beginning with multiview range images of a 3D object, we segment the object into parts at deep surface concavities. Motivated by physics, we detect these concavities by locating surface points where the simulated electrical charge density achieves a local minimum. The individual parts are then described by parametric geons. The latter are defined as seven distinctive volumetric shapes characterized by constrained superellipsoids, with deformation parameters controlling the tapering and bending. We obtain a unique part model by fitting all parametric geons to each part and classifying the fitting residuals. The advantage of our classification approach is its ability to approximate the shape of an object, even one not composed of perfect geon-like parts. The resulting 3D shape description is a prerequisite for generic object recognition.

1 Introduction

Biederman has proposed a theory of Recognition-by-Components(RBC) [4] based on psychophysical experiments which involves a modest set of volumetric primitives, called *geons*, as qualitative descriptions of object parts. The theory requires that an arrangement of a few geons be recovered from an object by segmenting the object at deep surface concavities. Then, having the geon-based description of the object, it was postulated that it could be quickly recognized with great tolerance to variations in viewpoint and shape. Since 1985, significant computational efforts have been made to create geon-based descriptions of 3D objects [13, 11, 17, 3, 24, 29, 28, 20, 36, 23, 10].

The paradigm of geon-based representation consists of two interrelated stages, object decomposition and geon identification. Object decomposition is concerned with segmenting an object into parts, for each of which a geon model needs to be identified. Decomposition algorithms for geon-based representation can be categorized as being region(shape or primitive)-based or boundary-based approaches. Region-based approaches [11, 29, 12] first find image regions that correspond to object surface patches, and then group these patches into individual parts based on particular geon surface configurations. This scheme relies on the limited number of surface shape types of *perfect* geons. Therefore each surface corresponding to a geon can be inferred by (conceptually) examining all possible combinations of surface types. The disadvantage of this approach is that if an object is not composed of perfect geons or mistakes are made in surface segmentation, the surface type label might be in error. This would produce an

incorrect object decomposition. Boundary-based approaches [2, 24] are directly inspired by the theory of Hoffman and Richards [16] and locate part boundaries at deep surface concavities. Since they concentrate on part boundaries rather than part shape, they can segment an object regardless of whether it is composed of perfect geons. However, there is a limitation since the part boundary points must be explicitly defined. This may not always be the case [16]. We also note that the literature mentions other part segmentation algorithms which were not specifically intended for geon-based representations but could well serve this purpose [26, 25, 8, 14, 15, 22, 30].

The early work on geon identification focused on the recovery of geon models from complete line drawings which depicted objects consisting of *perfect* geon-like parts [3, 11, 17]. Generally, however, such “clean” or complete line drawings of objects cannot be obtained easily due to the color and texture of object surfaces and complex illumination configurations. Because of this, and also for practical reasons, some research has focussed recently on data obtained from laser rangefinders [24, 29]. In both cases, part descriptions are determined in a bottom-up fashion, inferring global properties by first aggregating local features. This approach is not very effective when object features do not fully satisfy the exact definitions of the geons. Clearly, any computer vision system which successfully recovers qualitative descriptions needs to address the problem of *shape approximation*. Two approaches for geon model recovery have been proposed to achieve this. Both essentially amount to using *quantitative* methods. Raja and Jain [28] identified a shape as one of 12 geon types. They fitted range data to a deformed superellipsoid model [26, 32] and then classified its shape by comparing the object’s superellipsoid parameters with a database of geon parameter models. However, these superellipsoid parameters are very sensitive to even minor shape variations. By comparison, we have defined only seven parametric geon models [35]. We characterize object shape by fitting all of the models to multi-view range data and classifying the fitting residuals [36, 37]. This method uses parametric geons as explicit shape constraints, thereby enabling shape verification by directly comparing the goodness of fit between the seven geons and the part shape. There are other reasons for preferring the quantification of parameters over purely qualitative models. For example, in the case where qualitative shape information is the same for two objects, discrimination must rely on quantitative information, such as relative size of object parts or the specific curvature of the part axis. Moreover, object pose is often required for certain tasks, such as manipulation, for example. Thus Dickinson et al. have used qualitative, followed by quantitative, methods to recover ten geon models from both intensity images [10] and range data [12].

The key issue for geon-based representation is how to reliably produce *shape approximations* as object descriptions. This implies two steps: object decomposition and part identification. In this paper, we propose a novel approach to object decomposition into parts which is motivated by physics [38]. Assuming that the object to be segmented is a charged conductor, we simulate electrical charge over the object surface. Then we detect part boundaries by locating deep

surface concavities where local charge densities reach their minima. The object is then segmented into parts at these points. Our strategy is consistent with boundary-based approaches [16]. However, unlike previous work [14, 30], we do not compute surface curvature and also do not directly address the scale problem when computing surface properties. Other work has relied on surface curvature estimation. These approaches suffered because of the need to select a proper scale for the window size to compute surface curvature. Following the object decomposition stage we recover a parametric geon model for each part by simultaneously fitting all parametric geon models to a putative part and selecting the model with the minimum fitting residual. The advantage of this approach is that we can recover part models which *approximately* fit the data. Thus we can tolerate noisy data, object surfaces which are slightly rough, and parts whose shape deviates from perfect geons. By contrast, methods based on local feature aggregation would probably encounter difficulties using imperfect data.

2 Object Decomposition

2.1 Physics

When a charged conductor with an arbitrary shape is in electrostatic equilibrium, all charge resides unevenly on the outer surface of the conductor [7]. The charge density is very high at sharp convex edges and corners. Conversely, almost no charge accumulates at sharp concavities. Therefore, deep surface concavities which have been shown to delineate part boundaries [16] can be detected by significant charge density minima.

Electrical charge densities at sharp edges and corners have been carefully studied by Jackson [19]. An object edge or corner is defined as a C^1 discontinuity of an object surface. Fig. 1 depicts examples of these. By ignoring secondary global effects, Jackson has derived an approximate relationship governing the charge density ρ at an edge formed by two conducting planes, as shown in Fig. 2.

In Fig. 2 (a), β is the angle between two planes defining an edge and η is the distance from the edge to a point P , where the charge density is measured. Fig. 2 (b) shows ρ as a function of β and η . The larger β and the smaller η , the greater the charge density. We observe that all sections of the function are monotonic for constant η . This relationship (see [19]) also reveals a theoretical singular behavior of charge densities at edges (for $\eta = 0$) as follows:

$$\rho = \begin{cases} \infty & \text{if } \beta > \pi \\ \text{constant} & \text{if } \beta = \pi \\ 0 & \text{if } \beta < \pi \end{cases} \quad (1)$$

This means that the charge density is infinite, constant and zero when the angle defined by the two planes is convex, flat and concave, respectively. The singular behavior of charge densities at corners, which is similar to that at edges, has also been investigated [19].

Fig. 1. Examples of edges and corners.

The physical model we have used is the charge density distribution on a perfect conductor in free space, where there is no other conductor or charge. Let $\mathbf{r}$ define an observation point and $\mathbf{r}'$ a source point charge q on the surface, as shown in Fig. 3. It is customary to select the reference point of the electric potential ϕ at infinity. Thus ϕ at $\mathbf{r}$ caused by q at $\mathbf{r}'$ is equal to [5]:

$$\phi(\mathbf{r}) = \frac{q}{4\pi\epsilon_0} \frac{1}{|\mathbf{r} - \mathbf{r}'|} \quad (2)$$

Here ϵ_0 is a constant, known as the permittivity of free space. Suppose the charge is continuously distributed over the object surface S (see Fig. 3). Then the electrical potential at $\mathbf{r}$ is contributed by *all of the charge on S* . This can be expressed as follows:

$$\phi(\mathbf{r}) = \frac{1}{4\pi\epsilon_0} \int_S \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dS' \quad (3)$$

Here $dq = \rho(\mathbf{r}')dS'$, $\rho(\mathbf{r}')$ is the charge density at $\mathbf{r}'$ and S' represents a unit area on S .

According to physics [7], all points on a charged conductor in electric equilibrium are at the same electric potential. If we restrict $\mathbf{r}$ in Equation (3) to the conductor surface, then $\phi(\mathbf{r})$ is constant. Accordingly, (3) may be rewritten as follows:

$$V = \int_S \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dS' \quad (4)$$

Here $V = 4\pi\epsilon_0\phi(\mathbf{r})$ is a constant.

We note that the object being considered here is located in a real 3D space. A diffusion-based method to compute surface properties of such objects has been reported in the literature [39]. Working on a voxel-based coordinate system, this approach simulates the propagation of a specific number of particles among the object voxels. At a certain stage of the diffusion process particle accumulations at sharp surface concavities and convexities become significant. Therefore sharp concave and convex surface points can be detected by diffusion. There are two

Fig. 2. Charge densities near edges. (a) An edge formed by two planes with an angle β . (b) The charge density at $P(\eta, \beta)$.

problems associated with this approach. First, it is difficult to decide when the diffusion process should be stopped. At diffusion equilibrium, the particle density is uniform everywhere within the object and therefore cannot indicate any specific geometrical property of the object surface. Thus the diffusion process must be stopped before equilibrium. However, if diffusion is stopped at an inappropriate time, the distinction between particle densities at concave and convex surface points will not be strong. In contrast, our approach deals with electrical equilibrium, where charge densities at concave and convex surface points are very distinct. Second, a diffusion-based approach must work in a voxel-based coordinate frame. This is very time-consuming if the object is large. However, our approach concentrates only on the surface of the object and does not need to perform computation within its interior. This reduction in dimensionality reduces the cost of both computational memory and time.

Fig. 3. The observation point $\mathbf{r}$ and the charge source point $\mathbf{r}'$ on the surface S of an ellipsoid. O is the origin of the coordinate system.

In the next section, we derive the algorithm for computing the charge density based on Equation (4).

2.2 Finite Element Solution

Our objective is to compute the charge density distribution on the surface of an object having an irregular shape. Since S in Equation (4) is an arbitrary surface, it is impossible to solve the equation analytically. However, we can obtain an approximate solution to Equation (4) by using finite element methods [31]. To do this, we tessellate the object surface using a triangular mesh which has N planar triangles, $T_k, k = 1, \dots, N$. Each triangle is assumed to possess a constant charge density, ρ_k , as shown in Fig. 4. A set of basis functions $f_k, k = 1, \dots, N$ is defined on this triangular mesh as follows:

$$f_k(\mathbf{r}') = \begin{cases} 1 & \text{if } \mathbf{r}' \in T_k \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

The basis function, f_k , is nonzero only when $\mathbf{r}'$ is on the triangle T_k , as shown in Fig. 4. Therefore, the charge density $\rho(\mathbf{r}')$ can be approximated by a piecewise constant charge density function as follows:

$$\rho(\mathbf{r}') \approx \sum_{k=1}^N \rho_k f_k(\mathbf{r}') \quad (6)$$

Substituting (6) into Equation (4), we have

$$V = \sum_{k=1}^N \rho_k \int_{T_k} \frac{1}{|\mathbf{r} - \mathbf{r}'|} dS' \quad (7)$$

Fig. 4. Triangular mesh on the surface of an ellipsoid.

Since the charge density is assumed to be constant on each T_k , we may take $\mathbf{r}_i$ as the observation point on each T_i and rewrite Equation (7) as:

$$V = \sum_{k=1}^N \rho_k \int_{T_k} \frac{1}{|\mathbf{r}_i - \mathbf{r}'|} dS' \quad i = 1, \dots, N. \quad (8)$$

Because of charge conservation, the sum of the charges on each triangle equals the total charge on the surface of the conductor. Let Q be the total charge on the conductor and S_k be the area of T_k . Then we have

$$Q = \sum_{k=1}^N \rho_k S_k \quad (9)$$

Assuming Q is known, and using (8) and (9), we obtain a set of linear equations with $N + 1$ unknowns, $\rho_1, \dots, \rho_N$ and V . Since the integral in (8) can be evaluated analytically [34], the charge density distribution ρ_k and the constant V can be obtained by solving a set of linear equations.

2.3 Implementation

Our data are obtained using a laser rangefinder to scan objects from multiple views. Transformations between successive views are computed by a view registration algorithm [6]. Surface triangulation is performed by a method developed by DeCarlo and Metaxas [9]¹. The observation point $\mathbf{r}_i$ on each triangular patch is selected at the centroid of the triangle. The set of linear equations is solved by a standard LU decomposition method [27].

¹ Triangular mesh data used in this paper have been kindly provided by D. DeCarlo and D. Metaxas at the University of Pennsylvania.

Fig. 5. Direct Connection Graph (DCG). (a) A triangular mesh. (b) DCG of the triangular mesh in (a). (c) Subgraphs of (b) after boundary node deletion. Here triangular patches 1 and 2 are assumed to be located on the part boundary.

2.4 Part Segmentation

After obtaining the charge densities on the object surface, we segment an object into parts by detecting and then deleting points on the part boundaries. The method is based on a *Direct Connection Graph* (DCG) defined on the triangular mesh, as shown in Fig. 5. Here the nodes represent the triangular patches in the mesh and the branches represent the connections between direct neighbors. By direct neighbors we imply that two triangles share two vertices. For example in Fig. 5 (a), triangles 1 and 2 are direct neighbors while 2 and 3 are not. Thus, the DCG provides a convenient coordinate system on the object surface.

We have assumed that a part boundary forms a *closed* 3D curve, explicitly defined by deep surface concavities. This ensures that the decomposition algorithm will be able to segment a part from the rest of the object. The assumption also provides a stopping criterion for the boundary tracing procedure.

The algorithm examines the charge density on all triangles to find a starting triangle for tracing each boundary. A starting triangle must satisfy the following conditions:

1. It must be a *concave* extremum; that is, its charge density must be a local minimum.
2. It must be located at a *deep* concavity. Thus the charge density on the triangle must be lower than a preselected threshold².
3. It and its neighbors must not have been visited before. This ensures that the same boundary will not be traced again.

Beginning at the starting triangle, the algorithm proceeds to the neighbor with the lowest charge density. During the tracing procedure, all triangles detected on the boundary are marked. The marked ones will not be checked again and

² This threshold determines when an object should not be decomposed further. If the charge density at a starting triangle is greater than this threshold, we assume that all boundary points have been found. The selection of the threshold depends on a priori knowledge of the surface concavity and there is no universal rule to determine it. Currently we choose 1.05 times the lowest charge density on the object surface as the threshold.

eventually will be deleted from the DCG. The process continues until it returns to the starting triangle. As a result of the above assumption, this means that all triangles on this part boundary have been visited. Next the algorithm finds a new starting triangle and traces another boundary. It recursively repeats the same tracing procedure, and then stops when the charge density at a starting triangle is higher than the preselected threshold. After all triangles on part boundaries have been found, the nodes of the DCG representing these triangles are deleted. Thus the original DCG is now divided into a set of disconnected subgraphs, as shown in Fig. 5. Physically the object has been broken into parts. Each subgraph represents one part of the object which is ready for geon identification.

3 Geon Identification

This section is an overview of our work on parametric geon recovery. Details can be found in [35, 36, 37].

Motivated by the art of sculpture, we have defined seven parametric geons (see Fig. 6). These shapes are commonly used by sculptors as basic volumetric shapes. They are regular, simple, symmetrical, distinctive and can be expressed by compact analytic functions. Defined by

$$f(\mathbf{x}, \mathbf{a}_i) = 0, \quad i = 1, \dots, 7 \quad (10)$$

parametric geons are derived from superellipsoids [1] by specifying appropriate shape parameters and applying tapering and bending deformations [32]. Here $\mathbf{x} \in R^3$, $\mathbf{a}_i$ is the nine- to eleven-dimensional vector of model parameters, and i is the model shape type. Parametric geons provide both qualitative shape and quantitative size and deformation information. Compared to the geon shapes proposed by Biederman [4], certain qualitative properties have been simplified for the sake of simplicity. The detailed derivation of the implicit and normal vector equations of parametric geons can be found in [35].

3.1 The Objective Function

The major step in parametric geon identification is to fit these parametric models to object data by searching a specific parameter set $\mathbf{a}$ such that a two-term objective function

$$\mathcal{E}(\mathbf{a}) = d_1(\mathbf{a}) + \lambda \gamma d_2(\mathbf{a}) \quad (11)$$

is minimized. This objective function reflects the similarity of size and shape between the object and fitted models. The first term of the objective function is given by

$$d_1 = \frac{1}{N} \sum_{i=1}^N |e(D_i, \mathbf{a})| \quad (12)$$

Here N is the number of data points denoted by $\{D_i \in R^3, i = 1, \dots, N\}$.

Fig. 6. The seven parametric geons.

For the three regular primitives (ellipsoid, cylinder and cuboid), $e(D_i, \mathbf{a})$ is defined as the Euclidean distance from a data point to the model surface along a line passing through the origin of the model and the data point [33]. Since tapering or bending significantly complicates the implicit primitive equations, we cannot obtain a closed-form solution for $e(D_i, \mathbf{a})$, as can be done for the regular primitives. Thus, in order to reduce the amount of computation, we have derived an approximate distance measure for the tapered and curved models [36].

We define the second term ($\lambda \gamma d_2$) of the objective function by measuring a squared difference between the surface normal vectors $\mathbf{n}_d$ of objects and the surface normal vectors $\mathbf{n}_m$ of models at each corresponding position. Thus (d_2) is given by:

$$d_2 = \frac{1}{N} \sum_{i=1}^N e_n(i) \quad (13)$$

Here $e_n(i) = \|\mathbf{n}_d(i) - \mathbf{n}_m(i)\|^2$.

In (11) γ is taken as the average of the model size parameters. This tends to adapt the second term to the size of the object. λ is a weighting factor which controls the contribution of the second term to the objective function. In this paper, we choose $\lambda = 5$ according to a heuristic based on the shape difference between each pair of parametric geons [36]. A similar objective function [40] has been reported which used L_2 norms for both the first and second terms. In contrast, we use the L_1 norm for the first term in order to speed up the fitting process [36].

3.2 Very Fast Simulated Re-annealing

Fitting parametric geons to range data involves a search for a particular set of parametric geon parameters which minimize the objective function in (11).

Fig. 7. Recovery of parametric geons. (a) triangular mesh of the object; (b) Charge density over object surface. (c) Two separated parts; (d) object description in terms of parametric geons.

This function has a few deep local minima caused by an inappropriate orientation of the model, and many shallow local minima, caused by noise and minor changes in object shape. In order to find the model parameters corresponding to the global minimum of the objective function, we employ a stochastic optimization algorithm, Very Fast Simulated Re-annealing (VFSR) [18]. This algorithm permits an annealing schedule which decreases exponentially in annealing time and is dramatically faster than traditional simulated (Boltzmann) annealing [21] whose annealing schedule decreases logarithmically.

The search space can be limited somewhat by appropriately choosing constraints. A constraint on the model centroid and size can be estimated from the 3D range data. Constraints on the rotation, tapering and bending parameters can be specified by theoretical considerations [36].

After parametric geons are fitted to the 3D data, the best model for the object is selected according to the minimum fitting residual.

4 Experiments

4.1 Simulated Charge Distribution

We have applied our method to a simple vase in order to illustrate parametric geon recovery. Fig. 7 (a) shows the triangular mesh for this 3D object. Fig. 7 (b) shows the computed charge density distribution over its surface. The gray levels indicate charge densities, which are normalized in the range between 0 (darkest intensity) and 250 (white). It can be clearly seen that the lowest charge densities are located at surface concavities, which are at the intersection of the spherical and cylindrical portions of the vase. Conversely, since the edge on the top of the vase is sharply convex, the charge density at these points reaches a maximum.

4.2 Object Decomposition

Given that the charge density distribution over the object surface shown in Fig. 7 (b), our algorithm finds the part boundary shown in black. It then segments the object at these boundary points, breaking the object into two parts as illustrated in Fig. 7 (c).

4.3 Parametric Geon Recovery

After the object data are segmented into parts, our algorithm fits all of the parametric geon models to the data of each part. The model for a particular part is selected based on the minimum fitting residual. An ellipsoid and a cylinder were found as models of the lower and upper parts of the vase, respectively. The parametric geon description of the vase is shown in Fig. 7 (d).

5 Conclusions

We have presented a new approach to 3D object segmentation and representation. The general strategy is to segment an object into parts at deep surface concavities and approximate the shape of each part in terms of finite volumetric primitives. We compute and analyse a simulated charge density distribution over the object surface to find the deep surface concavities. The object is broken into parts at these points. Unlike previous curvature-based approaches, this method computes local surface information without the frustration of having to choose a particular scale for analyzing local data. Although our algorithm works in a full 3D space, it does not deal with the interior of objects as do voxel-based approaches. This saves both memory and time by restricting the computations to the object surface. We recover a qualitative part description by comparing the data with all seven geon models. A particular one is selected based on a measured similarity function which accounts for both shape and size. Using an experimental example we have illustrated how our approach derives a parametric geon-based description from multiview range images. This kind of description is a prerequisite for generic object recognition.

Acknowledgements

We wish to thank Professors Lester Ingber, Demetri Metaxas, and Jonathan Webb, Gerard Blais, Douglas DeCarlo and Gilbert Soucy for their kind help. MDL would like to thank the Canadian Institute for Advanced Research and PRECARN for its support. This work was partially supported by a Natural Sciences and Engineering Research Council of Canada Strategic Grant and an FCAR Grant from the Province of Quebec.

References

1. A. H. Barr. Superquadrics and angle-preserving transformations. *IEEE Computer Graphics Applications*, 1:11–23, 1981.
2. R. Bergevin and M. D. Levine. Part decomposition of objects from single view line drawings. *CVGIP: IMAGE UNDERSTANDING*, 55(1):73–83, January 1992.
3. R. Bergevin and M. D. Levine. Generic object recognition: Building and matching coarse descriptions from line drawings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(1):19–36, January 1993.

4. I. Biederman. Human image understanding: Recent research and a theory. *Computer Vision, Graphics, and Image Processing*, 32:29–73, 1985.
5. J. Van Bladel. *Electromagnetic Fields*. Hemisphere Publishing Corporation, revised printing edition, 1985.
6. G. Blais and M. D. Levine. Registering multiview range data to create 3D computer objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1994, Accepted for publication.
7. F. J. Bueche. *Introduction to Physics for Scientists and Engineers*. McGraw-Hill Book Company, New York, 3rd edition, 1980.
8. T. Darrell, S. Sclaroff, and A. Pentland. Segmentation by minimal description. In *Proceedings of Third International Conference on Computer Vision*, pages 112–116, Osaka, Japan, December 1990. IEEE.
9. D. DeCarlo and D. Metaxas. Adaptive shape evolution using blending. In *IEEE Proceedings of International Conference on Computer Vision (to appear)*, 1995.
10. S. J. Dickinson and D. Metaxas. Integrating qualitative and quantitative shape recovery. *International Journal of Computer Vision*, 13(3):311–330, 1994.
11. S. J. Dickinson, A. P. Pentland, and A. Rosenfeld. 3D shape recovery using distributed aspect matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):174–198, February 1992.
12. S. J. Dickson, D. Metaxas, and A. Pentland. Constrained recovery of deformable models from range data. In *Proceedings of International Workshop on Visual Form*, Capri, Italy, May 1994.
13. R. Fairwood. Recognition of generic components using logic-program relations of images contours. *Image and Vision Computing*, 9:113–122, 1991.
14. F. Ferrie and M. D. Levine. Deriving coarse 3D models of objects. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 345–353, Ann Arbor, Michigan, June 1988.
15. A. Gupta and R. Bajcsy. Volumetric segmentation of range images of 3D objects using superquadric models. *CVGIP: Image Understanding*, 58(3):302–326, November 1993.
16. D. Hoffman and W. Richards. Parts of recognition. *Cognition*, 18:65–96, 1984.
17. J. Hummel and I. Biederman. Dynamic binding in a neural network for shape recognition. *Psychological Review*, 99(3):480–517, July 1992.
18. L. Ingber. Very fast simulated re-annealing. *Mathematical and Computer Modelling*, 12(8):967–973, 1989.
19. J. D. Jackson. *Classical electrodynamics*. Wiley, New York, 1975. Secs. 2.11 and 3.4.
20. A. Jacot-Descombes and T. Pun. A probabilistic approach to 3-D inference of geons from a 2-D view. In K. W. Bowyer, editor, *Applications of Artificial Intelligence X: Machine Vision and Robotics, Proceedings of SPIE*, volume 1708, pages 579–588, Orlando, Florida, April 1992.
21. S. Kirkpatrick, Jr. C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.
22. A. Lejeune and F. Ferrie. Partitioning range images using curvature and scale. In *Proceedings of 1993 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 800–801, New York City, NY, June 1993. IEEE Computer Society Press.
23. R. C. Munck-Fairwood and L. Du. Shape using volumetric primitives. *Image & Vision Computing*, 11(6):364–371, July 1993.

24. Q. L. Nguyen and M. D. Levine. 3D object representation in range images using geons. In *11th International Conference on Pattern Recognition*, Hague, Netherlands, August 1992.
25. A. Pentland. Part segmentation for object recognition. *Neural Computation*, 1:82–91, 1989.
26. A. P. Pentland. Recognition by parts. In *The First International Conference on Computer Vision*, pages 8–11, London, June 1987.
27. W. H. Press, S. A. Teukolsky, W. T. Vetterling, et al. *Numerical Recipes in C*. Cambridge University Press, Cambridge, UK, 2nd edition, 1992.
28. N. S. Raja and A. K. Jain. Recognizing geons from superquadrics fitted to range data. *Image and Vision Computing*, 10(3):179–190, April 1992.
29. N. S. Raja and A. K. Jain. Obtaining generic parts from range data using a multi-view representation. *CVGIP: Image Understanding*, 60(1):44–64, July 1994.
30. H. Rom and G. Medioni. Part decomposition and description of 3D shapes. In *Proceedings of the 12th International Conference on Pattern Recognition*, volume I, pages 629–632, Jerusalem, Israel, October 1994. IEEE computer Society, IEEE Computer Society Press.
31. P. P. Silvester and R. L. Ferrari. *Finite Elements for Electrical Engineering*. Cambridge University Press, Cambridge, 2nd edition, 1990.
32. F. Solina and R. Bajcsy. Recovery of parametric models from range images: the case for superquadrics with global deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(2):131–147, 1990.
33. P. Whaite and F. P. Ferrie. From uncertainty to visual exploration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(10):1038–1049, October 1991.
34. D. Wilton, S. M. Rao, A. W. Glisson, et al. Potential integrals for uniform and linear source distributions on polygonal and polyhedral domains. *IEEE Transactions on Antennas and Propagation*, AP-32(3):276–281, 1984.
35. K. Wu and M. D. Levine. Parametric geons: A discrete set of shapes with parameterized attributes. In *SPIE International Symposium on Optical Engineering and Photonics in Aerospace Sensing: Visual Information Processing III*, volume 2239, pages 14–26, Orlando, FL, April 1994. The International Society for Optical Engineering.
36. K. Wu and M. D. Levine. Recovering parametric geons from multiview range data. In *IEEE Conference on Computer Vision & Pattern Recognition*, pages 159–166, Seattle, June 1994. IEEE Computer Society.
37. K. Wu and M. D. Levine. Shape approximation: From multiview range images to parametric geons. In *Proceedings of 12th International Conference on Pattern Recognition*, volume 1, pages 622–625, Jerusalem, Israel, October 1994. IEEE Computer Society Press.
38. K. Wu and M. D. Levine. 2D shape segmentation: A new approach. Technical Report TR-CIM-95-01, Centre for Intelligent Machines, McGill University, Montreal, Quebec, Canada, February 1995.
39. Y. Yacoob and L. S. Davis. Labeling of human face components from range data. *CVGIP: Image Understanding*, 60(2):168–178, September 1994.
40. N. Yokoya, M. Kaneta, and K. Yamamoto. Recovery of superquadric primitives from a range image using simulated annealing. In *Proceedings of International joint Conference on Pattern Recognition*, volume 1, pages 168–172, 1992.