Enhanced 3D Representation Using Multiple Models

Nigel Ayoung-Chee, Gregory Dudek and Frank P. Ferrie
Center for Intelligent Machines, McGill University
Montréal, Québec, Canada H3A 2A7
e-mail: {ayoung, dudek, ferrie}@cim.mcgill.ca

Abstract

This paper deals with generic 3D shape modelling for the purposes of object recognition. Difficulties with many existing methods are that they either capture insufficient detailed structure or fail to provide sufficiently abstract descriptions. The approach presented here attempts to address this problem by building a composite representation of the data in terms of a superquadric augmented with multi-scale patches.

This is illustrated experimentally using laser range data. The superquadric that results in the best possible fit is expressed in terms of its position, size, shape and pose parameters. The residual of the fit is then modelled at several scales using multiple surface patches with uniform mean and Gaussian curvature. A hierarchical ranking of these patches is used to describe the residual based on geometric properties. These geometric properties are ranked according to criteria expressing their stability and utility. The most stable patches are selected as the description of the residual. The resulting representation can then be used for both pose estimation and object recognition.

1 Introduction

In this paper, we consider the 3D modelling with particular emphasis on recognition, but keeping in mind that generic models should be useful for other tasks such as grasping. The recovery and representation of depth data from scenes is a key goal in computational vision. The representation of this data in a format that allows a wide range of tasks to be performed has proven to be an elusive goal. Common approaches to the problem include using volumetric models such as superquadrics or deformable solids. While superquadrics are effective for capturing the global geometry of an object they lack the ability to model small scale features. Their use in recognition tends to be limited with respect to certain canonical shapes [2]. While deformable solids are able to *model* a wider class of shapes, their many degrees of freedom make it difficult to extract stable representations of objects for recognition. In short, they do not provide a convenient abstraction mechanism for robust indexing.

An alternative involves modelling objects as a collection of surface patches or features. Typical examples involve the use of specialized edge information [21] or curvature extrema [10, 6]. For object representation and recognition, this often involves fitting patches or membranes to the raw surface data. While the resulting description is often rich enough to accurately describe the surface of an object, the fitting and recognition process can be computationally expensive. Such descriptions also do not capture the generic shape of objects which is needed for tasks such as grasping.

One solution is to use a model composed of two different representations such that their benefits complement each other. The problem that remains is how to combine these different levels of abstraction into a representation that suitable for recognition. In this paper we describe such a model that is comprised of a volumetric model and a model based on surface curvature.

1.1 Outline

The approach in this paper involves modelling objects described by range data. Although the emphasis in this paper is on recognition, we are also specifically interested in grasping and trajectory planning. The sensing context we use involves a laser rangefinder attached to the end of a robot arm.

Range data images $z(x, y)$ were made of the objects from a single viewpoint (directly above). A superquadric model is then fitted to the data using nonlinear least squares minimization (Levenberg-Marquardt fitting) applied to the superquadric inside-outside function [18]. This representation is used for determining the gross superquadric fit to the data and is used to define the equivalence class of the model. The residual error in the superquadric fit is then computed and provides a residual surface

for subsequent modelling. It is this surface that provides the association between the volumetric and surface modelling. The structure of this residual surface R is modelled at various scales with surface patches obtained by decomposing the residual surface into patches of uniform mean (H) and Gaussian (K) curvature. These surface patches are then ranked according to their stability and salience for recognition.

The most stable patches are used to supplement the superquadrics in the recognition process. The superquadrics are used to focus on the same class of objects with respect to geometry and generic shape of the object. The information from the surface patch modelling is then used to disambiguate between the possible solutions.

1.2 Background

Barr [4] proposed the use of superquadric primitives and angle preserving transformations to extend the geometric primitives, quadric surfaces and parametric patches, to allow complex solids and surfaces to be constructed. Subsequently, Pentland [16] and also Solina and Bajcsy [18] used superquadrics for the recovery of compact volumetric models for shape representation from 3 dimensional range points. Model recovery was formulated as a least squares minimization of a cost function whose value depended on distance of the data points from the model's surface and the global parameters of the model (size, position and orientation). However, the class of objects that could be represented accurately are limited to single-part convex shapes.

Ferrie et al. [9] approached the problem of multi-part objects in a purely bottom-up approach which segments the range data based on curvature consistency. Each part of the segmented range image is then modelled with superquadrics. This form of modelling has been implemented in an active recognition system in which the data acquisition is guided by feedback from the recognition strategy [20]. In this way, as much data as possible is collected to ensure an accurate model of the object. The only drawback to this system is that it is sometimes impossible to obtain information from more than one viewpoint. This lack of data increases the possibility of producing ambiguous models of different objects.

Metaxas and Terzopoulos [19] developed dynamic models based on global and local deformation properties from superquadric ellipsoids and membrane splines. However the parameters for local deformation are only used in the reconstruction process and only the gross shape information is used for recognition.

Pentland and Sclaroff [16] propose a method for recovering a 3D model based on finite element methods (FEM) and parametric solid modelling using implicit functions. Physically based methods are usually not useful for recognition due to the fact that the problem is under-constrained (due to number of degrees of freedom). To overcome this, some of the high-frequency modes are discarded to obtain an over constrained estimate of shape. The primary limitation of such an approach is the fact that it is impossible to uniquely determine an object's rotation state from a single viewpoint.

Several methods for describing surfaces using local properties or hierarchies have been proposed [14]. We are particularly interested in methods based on perceptually significant cues. The significance of curvature in feature detection and object identification has been recognized for some time [3, 12] is well documented in psychophysics. Besl and Jain [5] implemented a surface reconstruction strategy which involved segmenting the surface into regions of arbitrary shape and approximating the image data with bivariate functions. In their work, coarse segmentation is performed based on surface curvature. This is then refined using an iterative region growing method based on variable-order surface fitting. However experimental results focused on image segmentation and surface reconstruction and did not explore its applicability to recognition.

Various authors have used curvature information for the representation of curves or surfaces [13, 15]. An issue with curvature extraction is the selection of the correct scale. Most curvature measurement techniques assume that there is a unique curvature that can be measured at each point. Some work has considered the combination of scale information with curvature information using methods such as Gaussian scale-space or curvature-tuned smoothing. Mokhtarian [15] presents results that support the use of a multiple scale space representation for representing curves. However its use in the recognition process is not explored. Dudek and Tsotsos [8] use a multi-scale representation of curvature to perform recognition, but they do not take advantage of stable features to optimize the recognition process.

2 Theory

Our goal is to develop a general purpose model that is useful for a wide variety of tasks. This includes relatively simple tasks such as grasping or localization to recognition. At the same time we also want

to avoid constructing a model that too complex to manipulate efficiently.

Rather than try to satisfy these conflicting requirements with a single model, we have chosen to use two classes of primitive. A volumetric primitive for simpler tasks such as obtaining the overall shape of an object (for classifying the object), and a more detailed model based on a surface decomposition for complex tasks such as recognition. This allows us to scale the complexity of the model to the task.

To represent the geometric properties of objects superquadric models were used. Superquadrics provide basic information about pose and shape, this is often sufficient to perform tasks such as grasping and classifying canonical shapes. Also, the restricted degrees of freedom makes the parametrization of the model tractable. Rather than construct a separate representation of the original data using surface patches, the surface patches are used to model the error in the superquadric fit. In short, we describe the residual surface that expresses the aspects of the data not captured by the superquadric. This gives us a measure of the quality of the fit of the superquadric model to the original object and provides a relationship that is more meaningful than two independent modelling processes.

2.1 Formulation of Superquadrics

The form of superquadrics are the same as that used by Ferrie et al.[9]. The superquadric is described by 11 parameters: size (a_x, a_y, a_z), shape (ϵ_1, ϵ_2), position (t_x, t_y, t_z) and orientation (r_x, r_y, r_z). The surface is defined by the following 3D vector:

$$v = \begin{pmatrix} a_x \cos^{\epsilon 1}(\eta) \cos^{\epsilon 2}(\omega) \\ a_y \cos^{\epsilon 1}(\eta) \sin^{\epsilon 2}(\omega) \\ a_z \sin^{\epsilon 1}(\eta) \end{pmatrix} \quad (1)$$

$$-\pi/2 \leq \eta \leq \pi/2$$
$$-\pi \leq \omega \leq \pi$$

where η and ω correspond to latitude and longitude angles of vector v expressed in spherical coordinates. The vector v originates in the coordinate centre and sweeps out a closed surface in space when the two independent parameters, η and ω change in the given intervals.

Superquadrics can model a large set of basic shapes like spheres, cylinders and parallelepipeds. Parallelepipeds are produced when ϵ_1 and ϵ_2 are $\ll 1$. When ϵ_1 and ϵ_2 are 1, the shape is more rounded. As ϵ_1 and ϵ_2 become greater than 2 the model becomes progressively more pinched. For this paper, $0.1 < \epsilon_1, \epsilon_2 < 1$.

2.2 Superquadric Equations

By eliminating parameters η and ω in eq. 1 we get the following implicit equation

$$\left(\left(\left(\frac{x}{a_x} \right)^{2/\epsilon_2} + \left(\frac{y}{a_y} \right)^{2/\epsilon_2} \right)^{\epsilon 2/\epsilon 1} + \left(\frac{z}{a_z} \right)^{2/\epsilon_1} \right) = 1 \quad (2)$$

Based on this equation, the following function known as the "inside-outside" function is obtained

$$F(x,y,z) = \left[\left(\left(\frac{x}{a_x} \right)^{2/\epsilon_2} + \left(\frac{y}{a_y} \right)^{2/\epsilon_2} \right)^{\epsilon 2/\epsilon 1} + \left(\frac{z}{a_z} \right)^{2/\epsilon_1} \right]^{\epsilon 1} \quad (3)$$

This function returns a result that is less than, greater than, or equal to 1 if the point (x, y, z) lies inside, outside or on the surface of the superquadric. For a superquadric in general position and orientation, the inside-outside function is expanded to 11 parameters

$$F(x,y,z) = F(a_x, a_y, a_z, \epsilon_1, \epsilon_2, t_x, t_y, t_z, r_x, r_y, r_z) \quad (4)$$

2.3 Fitting to Range Data

The fitting of a superquadric to range data is based on a least squares minimization of the superquadric inside-outside function. There are 2 stages to the fitting process, an initialization stage followed by a stage in which the initial estimate is refined. The initial estimate for the model uses the moments of inertia of the data points to estimate the orientation of 3 three orthogonal axes of an ellipsoid. Once the orientation is obtained the data is projected onto the axes and the maximum projected value on each axis is used as the axis length (a_x, a_y, a_z) for that axis. The Levenberg-Marquardt [17] method for nonlinear least squares minimization is then used to iteratively refine the initial estimate until the average error from one stage to the other is within some threshold.

2.4 Residual Surface

Because superquadrics can only represent a restricted class of shapes, the resulting representation often fails to capture small scale (relative to the size of the object) features such as bumps or concavities. In addition, incomplete information such as that obtained from a single viewpoint can lead

to gross errors in the estimation of the shape of the object.

To determine how well the object is modelled by a superquadric, the error in the fit is calculated. Computing the residual by a parallel (to the global coordinate system) projection from the range data to the surface of the superquadric was initially considered, but this resulted in unstable results for a superquadric in general position and orientation. Consider a vector (P_{ij}) projected from the range data to the superquadric. As Figure 1 shows, the magnitude of P_{ij} can vary depending on the orientation of the superquadric. In Figure 1(a), the residual for the projection shown is less that the residual shown in Figure 1(a) for the same point.

In extreme cases, this type of projection can result in discontinuities in the residual surface. Figure 2 illustrates an example in which the surface of the superellipsoid does not lie directly below part of the range data. Such discontinuities result in a loss of potentially useful information at these points. It is possible to determine when such conditions might occur, but it is not clear what should be done if they are detected.

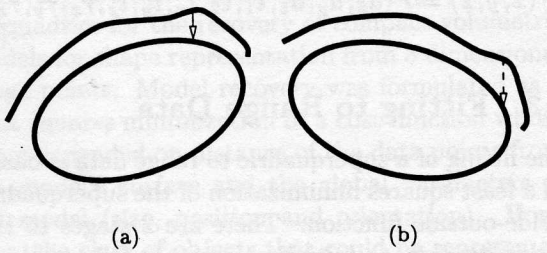

Figure 1: Illustration of how estimates of the residual can change with the orientation of the range data when they are calculated using the global coordinate system as the reference frame.

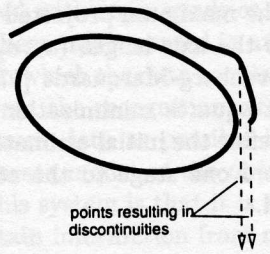

Figure 2: Illustration of how discontinuities in the residual surface can occur. Because of the orientation of the range data, a parallel projection from the range data does not intersect the surface of the superquadric.

To overcome these problems we decided to couple the residual surface to the local coordinate system of the superquadric. This was done by defining P_{ij} as a radial projection from the range data to the centre of the superquadric. Figure 3 illustrates how this is done.

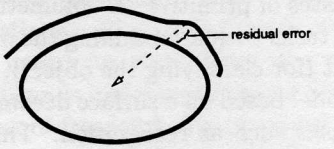

Figure 3: Illustration of the calculation of the residual surface using radial projection.

To use the information in the residual surface, the data needs to be abstracted into a model. Since the residual surface is a product of features that are not easily modelled by superquadrics, the model must be rich enough to capture the features that were overlooked by the first modelling process. For our purposes, surface patches with uniform *mean* (H) and *Gaussian* (K) curvatures were chosen. To address issues in selecting the correct scale, the patches are extracted at multiple scales and ranked based on their geometric properties [1].

2.5 Surface Patches: Calculation of Mean and Gaussian Curvatures

At a given point on a surface a curve is formed by the intersection of the surface and the normal plane in a given tangent direction. The curvature of this planar curve is called the *normal curvature* κ_n at that point. The maximum and minimum curvatures at a point define the principal curvatures κ_{max} and κ_{min}. The Gaussian curvature at a point is defined as the product $\kappa_{max}\kappa_{min}$ and the mean curvature is defined as $(\kappa_{max} + \kappa_{max})/2$.

To obtain a description via surface patches, an initial coarse segmentation of the surface based on H and K mapping is performed. The H and K curvature was chosen because of their invariance to rotation, translation and re-parametrization. The surface is re-sampled to a uniform grid and is multiply subsampled in the process of curvature calculation. Zero to multiple pixel skipping is applied along both the i and j directions and the sampled pixels are a linear average of the skipped pixels. Although the surface is subsampled in the curvature calculation process the full surface grid is preserved for subsequent patch merging and attribute calculation.

To calculate differential measurements it is necessary to filter the surface to smooth out local fluctuations in the surface due to quantization errors and

measurement noise. This was done using a triangular kernel as shown in figure 4. This linear filter was chosen for its simplicity and because it is localized in the frequency domain. Therefore it is less likely to produce features that were not part of the original image.

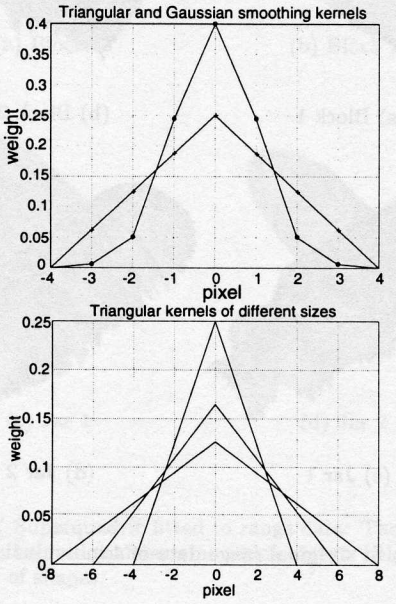

Figure 4: Filters used for data smoothing

2.6 Surface Estimation and Segmentation

The Mean and Gaussian curvatures are the basis for the segmentation of the surface into patches of constant curvature. A label or surface type is then assigned to every pixel in the image based on the signs of the Mean and Gaussian curvatures according to the following equation,

$$T(i,j)=1+3(1+sgn_{\epsilon_H}(H(i,j)))+(1-sgn_{\epsilon_K})k(i,j) \quad (5)$$

$$\text{where } sgn_\epsilon(x) = \begin{cases} +1 & \text{if } x > \gamma \\ 0 & \text{if } |x| \leq \gamma \\ +1 & \text{if } x < \gamma \end{cases}$$

and γ is a threshold. Coarse initial segmentation is obtained by finding connected regions in the sense of four connectedness. The *coherence* property of piecewise smooth surfaces makes pixels with the same label cluster together. After the range data has been segmented into regions of uniform curvature, a refinement to the initial coarse segmentation is performed by merging segments that appear to be part of the same surface primitive but are labelled as distinct due to noise, quantization error, or other factors. The resulting regions of constant curvature are approximated with biquadric polynomial surfaces defined as:

$$f_l(x,y) = \sum_{i=0}^{m}\sum_{j=0}^{n} a_{ij}x^i y^j \quad \text{where } (i+j) \leq 2 \quad (6)$$

Initial coarse segmentation patches of size **N** pixels or higher are used as seed patches for refinement of the segmentation. Quadric surface patches represented by Equation 6 are fit to each patch and the fit error is calculated. Patches with a low residual fit are selected for "growing". The patches are extrapolated to its surrounding regions and merged with a neighbouring patch if the residual of the extrapolation fit is less than a threshold.

2.7 Selection of Best Patches

Once the object has been segmented the geometric size and shape of each patch is determined by fitting ellipses to the patches using a *best eigenvector fit*[7]. In addition, the following properties are calculated

- $p_{type}(S_i)$: $surface\ type$
- $p_{size}(S_i)$: $size$, the area of a silhouette determined by the projection of the 3-D patch points on a plane fitting the patch.
- $p_{compactness}(S_i)$: $compactness(4\pi a/l^2)$, a measure of how close to a circle the silhouette is, where A is the size of the patch and l is the perimeter.
- $p_{max-radius}(S_i)$: $maximum\ radius$, variance along the major axis of ellipse.
- $p_{min-radius}(S_i)$: $minimum\ radius$, variance along the minor axis of the ellipse.
- $p_{elongation}(S_i)$: $elongation$, $(p_{max-radius}(S_i)/p_{min-radius}(S_i))$
- $p_{fit}(S_i)$: $goodness-of-fit$ of surface patch to data.

These properties are then used in a linear function to rank the utility and stability of patches. The most stable patches are used to describe a curved object. The exact function is described in Equation 7

3 Experimental results

In this section we intend to demonstrate that the superquadric model is sufficient for classifying objects based on their overall geometry. We will also show that a meaningful curvature-based description can

be extracted from the fit residual, and that surface patches provide sufficient information to distinguish between objects in the same class.

For this paper, range data was used to test our approach. The size of the range images are 256x256 pixels and the objects were scanned from a single viewpoint using a laser range-finder attached to a Puma robot arm. The objects were separated from the background by clipping the range data corresponding to the supporting platform. Figure 5 shows the range data used for the experiments. Two classes of objects were used. One group (Figures 5(a),5(b)) is composed of 2 blocks with the same dimensions and generic shape but with different shapes cut out of their surface. The second group is a pair of jars (Figures 5(c),5(d)). These also have the same basic shape but one has an additional series of grooves (Figure 5(d)).

Superquadrics were then fitted to the data as described in section 2.1. Figure 6 shows the range data superimposed over the fitted models.

The resulting superquadric models for the blocks (Figures 6(a),6(b)) have the same general shape but differ from the superquadric models in Figures 6(c) and 6(d) that were fitted to the jars. Table 1 lists the parameters (P_{sq}) for the superquadrics that were fitted to the data. The results show that the parameters, particularly the size and shape parameter, are similar enough that small perturbations in the data can make it difficult to distinguish between the two pairs of models [2].

P_{sq}	Block 1	Block 2	Jar 1	Jar 2
a_x	84.4594	79.8201	133.1806	120.2001
a_y	50.2412	63.4172	51.9548	41.4134
a_z	11.6777	12.7192	46.3033	37.9726
ϵ_1	0.1	0.2657	0.9820	0.9697
ϵ_2	0.1	0.1136	1.0819	0.8463

Table 1: Parameters for fitted superquadrics

The residual fit of the superquadric to the range data was then calculated as described in Section 2.4. Projections were made from the range points to the centre of the superquadric. The residual fit of the superquadric R was defined as the distance from the range points to the surface of the superquadric along these projections. Figure 7 shows the residuals of the superquadric fit. In Figures 7(a) and 7(b) the "bottom" faces of the superquadrics were fitted to the bottom of the cavities in the block models, therefore the residual is low everywhere except at the edges of the cavity.

The surface R was then modelled at different

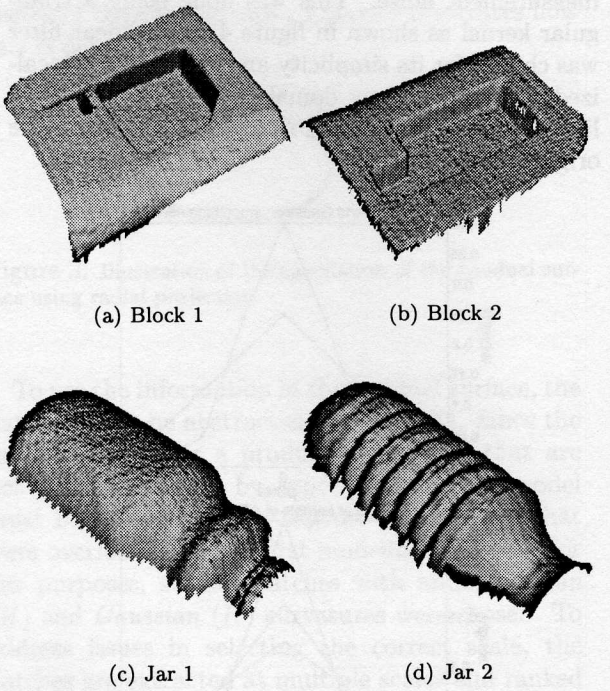

(a) Block 1 (b) Block 2

(c) Jar 1 (d) Jar 2

Figure 5: Original range data of 4 individual models

scales with surface patches using linear filters from 4 (small scale) to 10 (large scale) pixels. The patches with a size greater that α and a mean squared error lower than ϵ were then ranked in terms of their utility and stability based on a weighted sum of the characteristics outlined in section 2.7. Specifically, this *reliability* measure is defined as follows:

$$\begin{aligned} reliability(S_i) = \; & \omega_1 p_{scale}(S_i) + \omega_2 (p_{size}(S_i)/A) \\ & + \omega_3 p_{compact}(S_i) \\ & + \omega_4 p_{goodness-of-fit(S_i)} \end{aligned} \quad (7)$$

where ω_i is a weighting vector, and α, ϵ and ω_i are experimentally determined and maintained fixed. This measure favours large, compact, low fit error patches at high scale for describing an object.

Figure 8 illustrates the result of extracting surface patches at multiple scales (small to large) from the residual surface in Figure 7(d). At a small scale, the smaller features such as the ridges and grooves encircling the bottle are apparent. The larger scales extract the more stable features and less noise but they also fail to detect the finer details. By combining these multi-scalar descriptions, we can obtain a detailed description of the object. Figure 8 also demonstrates that certain features are apparent across several scales. In this particular example the neck of the jar is extracted at all scales. Features

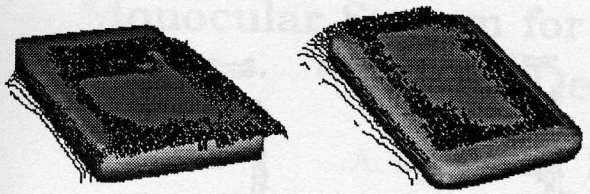

(a) Block 1 (b) Block 2

(c) Jar 1 (d) Jar 2

Figure 6: Superquadric fitted to range data: The resulting models can be used to determine which objects belong to the same class of shapes.

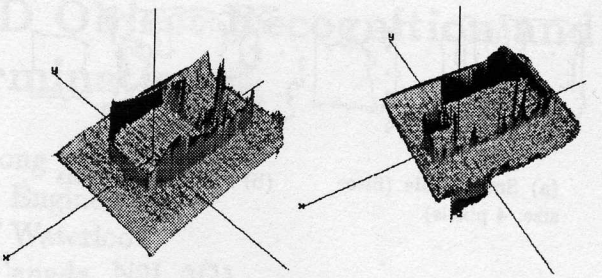

(a) Block 1 (b) Block 2

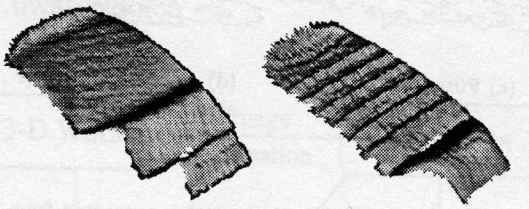

(c) Jar 1 (d) Jar 2

Figure 7: Residual surface of superquadric fit: The error is very small across the surface except around the edges of the concavities. By defining the outline of the concavities they help to differentiate one model from the other.

such as these are considered suitable for performing recognition.

Figure 9 shows an example of the significant patches that were extracted from the residuals shown in Figure 7 using a medium scale linear filter. At the representative scale the size, orientation and relative locations of the patches extracted differ enough to allow the objects to be distinguished from each other.

4 Discussion

As figure 6 shows, the superquadric models capture the canonical shape of the object. However it does not provide us with enough information for tasks such as recognition. The patches shown in figure 9 extract enough information to differentiate one object from the other. Key characteristics of the patches that could be used for recognition are the size, orientation and location of the patches *relative* to each other since our goal is a stable description that is viewpoint independent.

We propose a recognition scheme in which superquadrics would be used to determine the class of shapes to which an object belongs. For the case where this is not sufficient to uniquely identify the object, the surface description can then be used to refine the recognition process. This step would initially be performed at a coarse scale and progress to finer scales until the identity of the object is determined. Such a comparison can be performed by an algorithm such as an Interpretation Tree search [11]. In fact, our selection of significant surface patches is similar to the preprocessing used by Grimson and Lozano-Pérez to reduce the depth of the tree search.

The synthesis of superquadric and surface-based representations can be exploited in various ways. We have proposed using the superquadric as a primary indexing primitive. This, of course, presupposes that superquadric estimation can be accomplished robustly. An additional consideration is that the superquadric can be used for grasping or collision avoidance while a surface description may be used for tasks such as recognition.

References

[1] Wassim Alami and Gregory Dudek. Multiscale object representation using surface patches. In *Proceedings of SPIE - The International Society for Optical Engineering*, volume 2353, pages 108–119, Bellingham, WA USA, 1994.

[2] T. Arbel, F. P. Ferrie, and P. Whaite. Recognizing volumetric objects in the presence of uncertainty. In *International Conference on Pattern Recognition*, volume 1, pages 470–476, Piscataway, NJ, USA, 1994. IEEE.

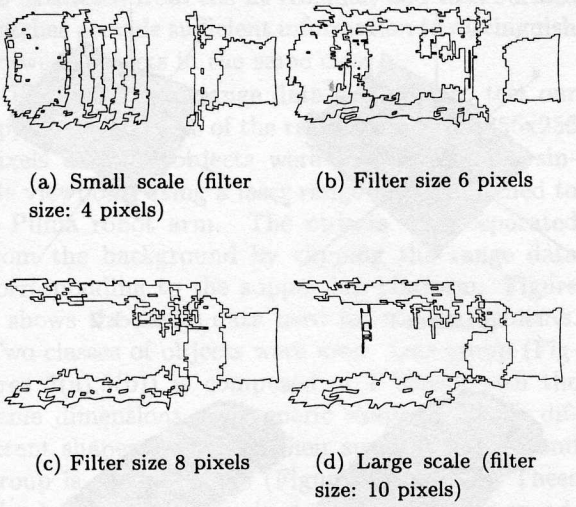

(a) Small scale (filter size: 4 pixels)

(b) Filter size 6 pixels

(c) Filter size 8 pixels

(d) Large scale (filter size: 10 pixels)

Figure 8: Examples of surface patches extracted from the residual of Jar 2 at various scales

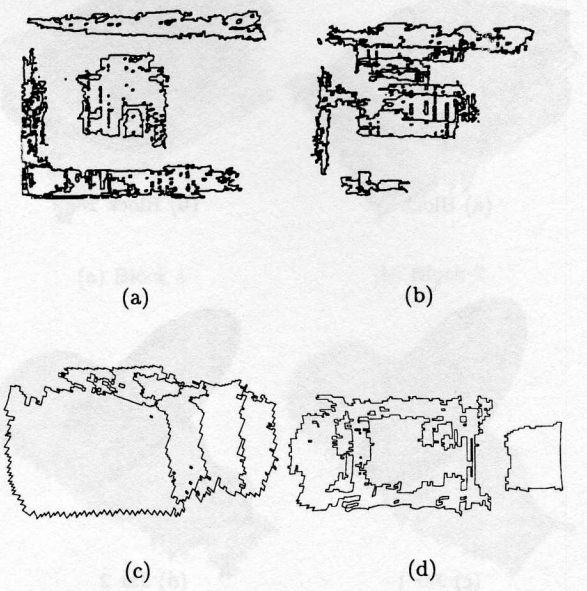

(a)

(b)

(c)

(d)

Figure 9: Significant surface patches extracted from residual using medium scale operator for (a) Block 1, (b) Block 2, (c) Jar 1 and (d) Jar2

[3] F. Attneave. Some informational aspects of visual perception. *Psychological Review*, (61):183–193, 1954.

[4] Alan H. Barr. Superquadrics and angle preserving transformations. *IEEE Computer Graphics and Applications*, 1:11–23, Jan. 1981.

[5] P. J. Besl and R. C. Jain. Segmentation through variable-order surface fitting. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 10(2):167–192, March 1988.

[6] Paul Besl. Invariant surface characteristics for three-dimensional object recognition in range images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 33(1):33–80, January 1986.

[7] Richard O. Duda and Peter E. Hart. *Pattern Classification and Scene Analysis*. John WIley & Sons, New York, N.Y., 1973.

[8] G. Dudek and J. K. Tsotsos. Shape representation and recognition from curvature. In *Proceedings of the 1991 Conference on Computer Vision and Pattern Recognition*, pages 35–41, Maui, Hawaii, June 1991. IEEE.

[9] F. P. Ferrie, J. Lagarde, and P. Whaite. Darboux frames, snakes and super-quadrics: Geometry from the bottom up. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 15(8):771–783, Aug. 1993.

[10] Dmitry B. Goldgof, Thomas S. Huang, and Hua Lee. Curvature-based approach to terrain recognition. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11(11):1213–1217, Nov 1989.

[11] W. Eric L. Grimson and Tomás Lozano-Pérez. Localizing overlapping parts by searching the interpretation tree. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 9(4):469–482, July 1987.

[12] D. G. Lowe. *Perceptual organization and visual recognition*. Kluwer Academic Publishers, Boston, Mass., 1985.

[13] David H. Marimont. A representation for image curves. *AAAI*, pages 237–242, 1984.

[14] Art Matheny and Dmitry B. Goldgof. The use of three and four dimensional surface harmonics for rigid and nonrigid shape recovery and representation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 17(10):967–981, Oct. 1995.

[15] F Mokhtarian. Evolution properties of space curves. In *Interlnational Conf. on Computer Vision*, pages 100–105, Tarpon Springs, Fla., Dec. 1988. IEEE 2nd International Conf. on Computer Vision.

[16] A. P. Pentland and S. Sclaroff. Closed-form solutions for physically based shape modelling and recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):715–729, July 1991.

[17] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, New York, N.Y., 1991.

[18] F. Solina and R. Bajcsy. Recovery of parametric models from range images: The case for superquadrics with global deformations. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(2):131–147, Feb. 1990.

[19] D. Terzopoulos and D. Metaxas. Dynamic 3d models with local and global deformations: Deformable superquadrics. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):703–714, July 1991.

[20] Peter Whaite and Frank P. Ferrie. From uncertainty to visual exploration. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(10):1038–1049, Oct 1991.

[21] D. Wilkes and J. Tsotsos. Active object recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 136–141, Urbana, Illinois, June 1992. IEEE Computer Society Press.