

# Recognition of Geons by Parametric Deformable Contour Models

Maurizio Pilu\* and Robert B. Fisher

Department of Artificial Intelligence, University of Edinburgh,  
5 Forrest Hill, Edinburgh EH1 2QL - SCOTLAND (UK)

**Abstract.** This paper presents a novel approach to the detection and recognition of qualitative parts like geons from real 2D intensity images. Previous works relied on semi-local properties of either line drawings or good region segmentation. Here, in the framework of Model-Based Optimisation, whole geons or substantial sub-parts are recognised by fitting parametric deformable contour models to the edge image by means of a Maximum A Posteriori estimation performed by Adaptive Simulated Annealing, accounting for image clutter and limited occlusions. A number of experiments, carried out both on synthetic and real edge images, are presented.

## 1 Introduction

In Computer Vision, the task of detecting and recognising general 3D objects in static scenes is still very far from a solution.

One of the most relevant and early recognised approaches to overcome the limitations of traditional CAD-model based vision is the *recognition by parts* theory, pioneered by Marr and Binford [3], which provides both a paradigm and a computational model for computer vision as well as for human vision.

Early attempts to define part models notably included polyhedra, which turned out to be too simple, and generalised cylinders which turned out to be too general and not easily detectable from real images.

Building on Lowe's work, Biederman's Recognition-by-Components (RBC) theory [2] provided a link between studies on human perception and computational vision by proposing a novel part classification scheme based on four significant *non-accidental properties*, cross-section shape, symmetry, sweeping rule and shape of axis, which were then used as a perceptual basis for the generation of a set of 36 components that he called geometrical ions, or geons (examples of geons are shown in Fig. 1-right).

There have been very few works that aimed at the detection and recognition of geons. Bergevin [1] exploited precise properties of line drawings (such as T-junctions, corners, faces, etc.) and so did Hummel and Biederman [7] by mean of their elegant neural network implementation. In the approach by Dickinson *et al.* [4], aspects of geons were pre-compiled and stored in a hierarchical structure

---

\* Partially supported by SGS-THOMSON Microelectronics.

and recognition performed by a probabilistic graph matching technique. Later, Metaxas *et al.* [13] built upon that work and both improved the segmentation and integrated stereo information to fit superquadrics models. The main drawback was that a good initial segmentation was required.

This paper describes a new approach to the detection of qualitative parts such as geons in 2D images. The approach matches deformable contours of whole geons to the edge image of a scene. In contrast to other methods, ours is truly global, in the sense that the whole geon is holistically considered as a single entity to be extracted from the image. A new parametric model is introduced that allow us to represent compactly and efficiently the various geon outlines. Then, within the framework of Model-Based Optimisation (MBO), a cost function that represent the quality of fit between model and image is maximised in the parameter space by using Adaptive Simulated Annealing (ASA) and the best image-model match is found. The cost function expresses in Bayesian terms the probability that a certain number of edgels match the geon contour model and some others do not, which has also an analogy in information theoretical terms.

At this stage of the research, the initialisation of the optimisation procedure is performed semi-automatically by manually selecting out the image regions corresponding to parts and by computing the principal moments of the thresholded regions.

The structure of the paper is as follows. First we describe the construction of geon contour models. Next we present the design of the objective function and how we perform its optimisation by Adaptive Simulated Annealing; some experimental results are given that support the validity of this approach. We conclude by proposing possible future extensions.

## 2 Parametrically Deformable Contour Model of Geons

Within the framework of Model-Based Optimisation, the recognition of geons from 2D images needs a model that can describe in a compact way their contour. As the geon models are computed inside the innermost loop of the optimisation process, this must be done as efficiently as possible.

Following [14] and much recent work on part decomposition from range data (e.g. [16, 17, 15, 13]), we extend the use of superquadrics (SQ) to the 2D case by using their projected visible contour (outline) as a 2D model, thus creating a geon *parametrically deformable contour model* (henceforth PDCM).

However, since a direct computation of SQ outlines is extremely expensive, we have pragmatically created an efficient approximate model that is suitable for qualitative geon PDCMs. Starting with a cylinder centred on the  $z$  axis with superelliptical cross-section, centred on the origin (Fig. 1-left) we apply deformations and rotations and find the outline by simple geometric considerations. The initial superelliptical cylinder  $C$  of height  $2a_z$  and semi-axes  $a_x$  and  $a_y$  can be expressed as

$$\mathbf{C} = \begin{bmatrix} \mathbf{x}(\eta) \\ \mathbf{y}(\eta) \\ \mathbf{z} \end{bmatrix} = \begin{bmatrix} a_x \cos(\eta)^\epsilon \\ a_y \sin(\eta)^\epsilon \\ \mathbf{z} \end{bmatrix} \quad \begin{array}{l} -\pi \leq \eta \leq \pi \\ -a_z \leq \mathbf{z} \leq a_z \end{array}, \quad (1)$$

where  $0 \leq \epsilon \leq 1$  controls the degree of squareness of the cross-section from a virtual block for  $\epsilon \rightarrow 0$  to a cylinder for  $\epsilon \rightarrow 1$ .

Any curve lying on this cylinder can be variously deformed. We use three types of deformations, *tapering* ( $\mathcal{T}$ ), *bending* ( $\mathcal{B}$ ) and *swelling* ( $\mathcal{S}$ ) along the principal axis.

Let us indicate by  $\mathbf{x}$ ,  $\mathbf{y}$ ,  $\mathbf{z}$  and  $\mathbf{X}$ ,  $\mathbf{Y}$ ,  $\mathbf{Z}$  the vector of shape points before and after the deformations, respectively. Then we have:

$\mathcal{T}(\mathbf{C}, K_x, K_y) =$	$\mathcal{B}(\mathbf{C}, c) =$	$\mathcal{S}(\mathbf{C}, s) =$
$\begin{cases} X = (\frac{K_x}{a_x} z + 1)x \\ Y = (\frac{K_y}{a_y} z + 1)y \\ Z = z \end{cases}$	$\begin{cases} X = x + \text{sign}(c)(R'' - r) \\ Y = y \\ Z = \sin(\gamma)(\kappa^{-1} - R'') \end{cases}$	$\begin{cases} X = x + \text{sign}(x)(R' \cos \alpha - (R' - \sigma)) \\ Y = y + \text{sign}(y)(R' \cos \alpha - (R' - \sigma)) \\ Z = R' \sin \alpha \end{cases}$

where  $r = \text{sign}(c) \cos(\beta) \sqrt{x^2 + y^2}$ ,  $\beta = \arctan \frac{y}{x}$ ,  $R'' = \kappa^{-1} - \cos(\gamma)(\kappa^{-1} - r)$ ,  $\gamma = \frac{z}{\kappa^{-1} - r}$ ,  $\kappa^{-1} = \frac{a_z}{|c|}$ ,  $\sigma = a_x s$ ,  $R' = \frac{(a_z^2 - \sigma^2)}{2\sigma}$ ,  $\alpha = \arctan \frac{z}{(R' - \sigma)}$  and  $0 \leq s \leq 1$ ,  $-1 \leq c \leq 1$  and  $-1 \leq K_x = K_y \leq 1$  are, respectively, the swelling, bending and tapering control parameters.

The tapering and bending deformations have been derived from [16]; the latter has been slightly modified by normalising the bending control parameter to  $a_z$ , which has improved the stability of its estimation, and to allow bending on both sides. The swelling deformation, however, has been introduced here to represent the geons' "expanding and contracting" sweeping rule [2].


Once deformed, the shape is rotated in space (by  $\theta_{pan}$ ,  $\theta_{tilt}$ ), projected ( $P$ ) and roto-translated in the image plane (by  $t_x$  and  $t_z$  and  $\theta_{opt}$ ). The whole chain of transformations of the initial 3D shape  $\mathbf{C}$  to its full projection onto the image plane  $z$ - $x$   $\mathbf{C}' = \begin{bmatrix} \mathbf{x}' \\ \mathbf{z}' \end{bmatrix}$  is then:

$$\mathbf{C}' = T(R_y(P(R_z(R_x(\mathcal{B}(\mathcal{S}(\mathcal{T}(\mathbf{C}, K_x, K_x), s), c), \theta_{pan}), \theta_{tilt})), \theta_{opt}), t_x, t_z) \quad (2)$$

Now we are ready to describe the construction of the geons' PDCM. The knottiest problem is to determine the occluding contour. For doing this, we employed the following approximation.

We apply the transformations in Eqn. (2) to the two bases of the superelliptical cylinder and take the four outermost points  $P1'_a$ ,  $P1'_b$  and  $P2'_a$ ,  $P2'_b$  (small circles in Fig. 1-left-B) and find the two corresponding points in the original undeformed superellipses (small circles in Fig. 1-left-A). These two points are linked by two 3D straight lines  $L1$  and  $L2$ , as shown in Fig. 1-left-A, and successively deformed according to Eqn. (2); the resulting  $L1'$  and  $L2'$  (Fig. 1-left-B) will then be used as the two lateral parts of the occluding contour.

By checking the projection of the normals  $\mathbf{n}_a$  and  $\mathbf{n}_b$  to the superelliptical ends on the image plane, we can then determine whether each of the two ends



**Fig. 1.** Determination of occluding contour (left) and some examples of PDCM (right).

are visible or not: if visible, the whole superellipse contour will be added in the geon PDCM; otherwise only its outermost part between  $P1'_a$  ( $P1'_b$ ) and  $P2'_a$  ( $P2'_b$ ) will be included in the final PDCM.

In the case we have a geon with squared cross-section (say  $\epsilon < 0.5$ ), the central edge is determined by joining the two corners  $P3_a$  and  $P3_b$  (Fig. 1-left-C) from the undeformed superelliptical bases occurring at  $\eta = \pi/4$  in Eqn. (1) by a 3D straight line and then deforming it by Eqn. (2); the resulting 2D curve is shown in Fig. 1-left-D.

Some examples of geon PDCMs produced by this method can be seen in Fig. 1-right; it takes less than 1ms of a SPARC 10 machine to create a model instance, over two order of magnitudes faster than any other method that would use raster scan techniques or surface normals.

This superquadric-inspired model can, as in [16, 17], represent the 12 geon classes which have cross-section symmetry [2] with a sufficient level of accuracy. However, by virtue of the employed geometric construction of the outline, they do not work properly with large bending or under certain viewing directions (about  $|\theta_{tilt}| > \frac{\pi}{3}$ ), where the occluding contour construction we employed is no longer applicable; despite this, it is a good trade-off between accuracy and speed.

### 3 Objective Function Design

Within the framework of Model-Based Optimisation, the matching of a PDCM to a geon image involves the minimisation of an objective (or cost) function that expresses the quality of image-model fit.

Seeking the best interpretation of the available data by the model, we express this cost function in Bayesian terms.

Let  $\mathbf{x}_i = [a_x \ a_y \ a_z \ \epsilon \ K_x \ s \ c \ \theta_{pan} \ \theta_{tilt} \ \theta_{opt} \ P_x \ P_z]^T$  be the vector of the PDCM parameters and  $M_i = M(\mathbf{x}_i)$  be a geon PDCM instance built as in Sec. 2,

expressed in term of pixels by a set of  $(i, j)$  pairs. Furthermore, let  $I$  be the original image and  $E$  the corresponding binary edge image:  $E(i, j)$  is zero if no edge has been detected at  $I(i, j)$ .

From the Bayes rule, the MAP estimate of the best model-image fit  $M_{best} = M(\mathbf{x}_{best})$  occurs for

$$\mathbf{x}_{best} : \max_{\mathbf{x}_i \in \mathbf{X}} \{P(M_i | E)\} = \min_{\mathbf{x}_i \in \mathbf{X}} \left\{ -\log \left( \frac{P(E | M_i) P(M_i)}{\sum_j P(E | M_j) P(M_j)} \right) \right\}, \quad (3)$$

with the denominator constant over  $\mathbf{X}$ , which is the set of hypotheses produced by the optimisation procedure.

The following two subsections describe how we defined the model-conditional image and prior probabilities.

### 3.1 Model-Conditional Image Probability

In Eqn. (3),  $P(E | M_i)$  expresses the conditional probability of having particular image evidence in the presence of the model.

Earlier experiments using various techniques based on sum of distances between model and edges (as used by Lowe [12] and others) had problems of unwanted shrinking or expansions so we concluded that they cannot properly cope with the use of deformable models and sparse unsegmented data. Our approach is to find the model that best accounts for the image in terms of matched and unmatched edgels.

Let

$$E_m(M_i) = \{(k, l) : |(i, j) - (k, l)| \leq d, (i, j) \in M_i\}$$


be the  $d$ -neighbourhood of the model contour  $M_i$ , and  $E_b(M_i) = E - E_m(M_i)$  the rest of the edge image which is not covered by it (background); henceforth we drop the  $M_i$  arguments wherever there cannot be ambiguities.

Let  $N_{b1}$  and  $N_{b0}$  be the number of pixels in the background  $E_b$  that are edge ("1") or non-edge ("0"). On the other hand, in the model neighbourhood  $E_m$ , and  $N_{m1}$  the number of *edgels* whose direction is locally consistent with that of the model  $M_i$  at that point (we set a threshold of  $\pi/10$ ); furthermore let  $N_{m0}$  be the number of pixels in  $E_m$  that are *not* matching the model, either because no edge has been detected there or because of directional inconsistency.

The model conditional probability  $P(E | M_i)$  globally expresses the probability of the presence or absence of edge points or consistent edgels in  $E_b$  and  $E_m$ . Since  $E_b$  and  $E_m$  can be regarded, to a first approximation, as realizations of two independent binary ergodic processes, whose probabilities are given by the product of single local outcome probabilities  $p_{b1}$  and  $p_{m1}$ , respectively, we have:

$$P(E | M_i) = P(E_b | M_i) \cdot P(E_m | M_i) = [p_{b1}^{N_{b1}} (1 - p_{b1})^{N_{b0}}] \cdot [p_{m1}^{N_{m1}} (1 - p_{m1})^{N_{m0}}].$$

The value of  $p_{b1}$  is given by the ratio between edge locations and the number of pixels in the image (typical values: 0.02-0.06);  $p_{m1}$  ranges from 0.6 to 0.9,



**Fig. 2.** Model prior probabilities.

depending on the neighbourhood dimension  $d$  and how good the edge detection of object contour is expected to be. By taking the logarithms on both sides and expanding we obtain:

$$\log(P(E | M_i)) = \log(P(E_b | M_i)) + \log(P(E_m | M_i)) = \quad (4)$$

$$K + [N_{m1} \log(p_{m1}) + N_{m0} \log(1 - p_{m1})] - [N_{m1} \log(p_{b1}) + N_{m0} \log(1 - p_{b1})],$$

where  $K$  is a constant term that can be dropped in the MAP estimation.

The terms of this equation can be also reinterpreted as the number of bits “saved” by representing  $E$  by the model  $M_i$  bringing all into an information theoretical framework (see, e.g., [6, 11]).

### 3.2 Model A Priori Probability

Within a Bayesian framework we can assign an occurrence probability  $P(M_i)$  of a certain PDCM, called the *model a priori probability*. The reasons for introducing it are essentially three: *i*) some parameter configuration are unlikely to occur (such as a bent and swollen object); *ii*) Certain configurations of parameters arise from a weird view point that would make detection impossible; and *iii*) It biases the fitting to more perceptually likely shapes. These considerations are both practical and also correspond to sensible assumptions to reduce the quantitative shape ambiguities caused by the projection.

We have defined a sensible heuristic to represent these loose constraints. The probability of each PDCM is expressed by overlapping (multiplying) marginal densities of parameter values or combination of them, tacitly assuming statistical inter-independence. The parameters we took into considerations are  $a_x$ ,  $a_y$  and  $a_z$ , swelling, bending and the viewpoint angles; the tapering and roundness parameters have not been taken into account and thus are given uniform probabilities. Below we describe these probabilities, whose non-normalised p.d.f. are also displayed in Fig. 2.

$$- P(\theta_{tilt}, \theta_{pan}) = N(\theta_{tilt} - \frac{\pi}{4}, \sigma_{\theta_{tilt}}) \cdot N(\theta_{tilt} + \frac{\pi}{4}, \sigma_{\theta_{tilt}}) \cdot N(\theta_{pan} - \frac{\pi}{4}, \sigma_{\theta_{pan}})$$

In our perception [2] there is a bias towards objects or surfaces in canonical stable positions.  $P(\theta_{tilt}, \theta_{pan})$  has therefore probability density like the one shown in Fig. 2-A for  $\sigma_{\theta_{tilt}} = \sigma_{\theta_{pan}} = \pi/6$ , which has two Gaussian bells

in correspondence to the most likely values of  $\theta_{tilt} = \pi/4$ ,  $\theta_{tilt} = -\pi/4$  and  $\theta_{pan} = \pi/4$ .<sup>2</sup>

-  $P(a_x, a_y) = N(|a_x - a_y|, \sigma_{a_x, a_y})$

This favours more compact cross-sections rather than weird rotation angles to account for the ambiguity caused by the projection and basically gives higher probability for  $a_x$  close to  $a_y$ , as shown in Fig. 2-B for  $\sigma_{a_x, a_y} = 20$ .

-  $P(a_z | a_x, a_y) = \begin{cases} N(\tau \frac{|a_x + a_y|}{2} - a_z, \sigma_{a_z}) & \text{if } a_z \leq \tau \frac{|a_x + a_y|}{2} \\ \frac{1}{\sqrt{2\pi} \cdot \sigma_{a_z}} & \text{otherwise} \end{cases}$

This favours elongated geons by giving lower and lower probability if its length is smaller than the average of cross-section dimensions by a factor  $\tau$ , as shown in Fig. 2-C for  $\tau = 1.5$  and  $\sigma_{a_z} = 20$

-  $P(c, s) = N(|c| \cdot s, \sigma_{c, s})$ .

This expresses the perceptual incompatibility between high swelling and bending; the p.d.f. is shown in Fig. 2-D for  $\sigma_{c, s} = 0.3$ .

Now that we have all the non-normalised p.d.f. and given the assumption of prior inter-independence between parameters, we just multiply them together to obtain the (non-normalised) *a priori* p.d.f. of the model:

$$\log(P(M_i)) = H + \log(P(\theta_{tilt}, \theta_{pan})) + \log(P(a_x, a_y)) + \log(P(a_z | a_x, a_y)) + \log(P(c, s)) \quad (5)$$

The normalisation constant  $H$  is unnecessary because it does not affect the MAP estimate.

We wish to remark that there are other possible ways of defining the model prior probability and that we could also incorporate more detailed specific domain-dependent knowledge about the scene structure.

## 4 Model Fitting Procedure

The model fitting is obtained, from (3), by the minimisation of

$$-\log(P(M_i | E)) = -\log(P(E | M_i)) - \log(P(M_i)), \quad (6)$$

where the two terms are given by Equations (5) and (4). This minimisation is, however, rather difficult to achieve, since it is extremely irregular and presents many shallow and/or narrow minima. By trying to minimise Eqn. (6) alone, we also found that sometime the optimisation got stuck in local minima because of the step-like nature of the model-conditional probability of Eqn. (4) (remember we used a binary “belonging to the model” criteria). For overcoming this problem we have added to Eqn. (6) a small smoothing term representing the average minimal distance between contour model and image edge points (we used a minimal distance transform computed off-line). The smoothing term does not

<sup>2</sup>  $\theta_{tilt}$  and  $\theta_{pan}$  affect the topology of the PDCM and these values can be regarded also as giving the high disambiguation distance between visual events [9].

	$a_x$	$a_y$	$a_z$	$\epsilon$	$K_x$	$s$	$c$	$\theta_{pan}$	$\theta_{tilt}$	$\theta_{opt}$	$P_x$	$P_z$
Lower	-5%	-5%	-5%	.01	-1.0	0.0	-0.8	$(\bar{\theta}_{pan} - 0.3)$	$(\bar{\theta}_{tilt} - 0.3)$	$(\bar{\theta}_{opt} - 0.4)$	$\frac{-N}{10}\%$	$\frac{-N}{10}\%$
Upper	+5%	+5%	+5%	.99	1.0	1.0	0.8	$(\bar{\theta}_{pan} + 0.3)$	$(\bar{\theta}_{tilt} + 0.3)$	$(\bar{\theta}_{opt} + 0.4)$	$\frac{+N}{10}\%$	$\frac{+N}{10}\%$
$\Delta$	1.0	1.0	1.0	0.2	0.2	0.2	0.1	0.05	0.05	0.05	1	1

**Table 1.** Ranges and gradient steps used in the optimisation procedure.

affect the MAP estimate but just helps convergence in cases where image and model are much displaced and the numerical computation of the gradient become meaningless due to the low number of edge points falling inside the model neighbourhood. This term can then be seen as “telling the optimiser where to go” in absence of other information.

In early stages of the work we used a Levenberg-Marquandt method with added random perturbations as used in [16] and other works, but this approach had difficult convergence. Our choice fell then to Simulated Annealing [10], which is a powerful optimisation tool that efficiently combines gradient descent and controlled random perturbations to perform the minimisation of non-convex functions. We used a recent publicly available version of Simulated Annealing called Adaptive Simulated Annealing (ASA) [8].

The initialisation and set-up of the ASA algorithm is discussed below.

**INITIALISATION** The initialisation part is concerned with estimating initial coarse part hypotheses (sometime called *frames*) that comprise position, orientation of the major axis and dimensions. This initialisation need not be precise and the degree of allowed inaccuracy depends upon the power of the optimisation procedure. Since no method is currently available for segmenting qualitative parts from 2D real cluttered images, we have used a semi-automatic procedure. Regions corresponding to parts are first selected out manually (see Fig. 4) and an adaptive thresholding technique is used to yield blob-like regions; orientation and size are found by computing the first and second moments and principal directions of the blobs, which are assigned to  $\bar{P}_x$ ,  $\bar{P}_z$ ,  $\bar{a}_z$ ,  $\bar{a}_x = \bar{a}_y$  and  $\bar{\theta}_{opt}$ , respectively. The values of  $\bar{a}_x$  and  $\bar{a}_y$  are set to be equal because we do not have prior information on the aspect-ratio of the part cross-section. Moreover tapering, bending and swelling are all initially set to zero,  $\bar{\theta}_{pan} = 0$  and  $\bar{\theta}_{tilt}$  experimentally set to  $\pi/6$ .

**OPTIMISATION SET-UP** The bounds of the hyper-rectangular search space and the step for the computation of the gradient, needed by the ASA optimiser greatly affect the nature of the convergence. Table 1 summarises the values we found reasonably good through our experiments. (The values expressed in percentage are relative to their respective initial values and  $N$  is the image resolution in pixels)

The two parameters specifying the annealing schedule, the Temperature Ratio Scale and Temperature Annealing Scale [8] have also been experimentally set to  $5 \cdot 10^{-5}$  and 50, respectively, and the Cost Precision to 0.0001. Finally, the number of iterations has been set to 2000, which we found to be a good



trade-off between speed and good convergence. The neighborhood distance  $d$  was set to 6 pixels for 512x512 images and to 2 pixels for 128x128 images. Each ASA running time ranged from 5 to 15 secs on a SPARC 10 machine, according to image resolution.

## 5 Experimental Results

In this section we present two sets of experiments that show the validity of the approach.

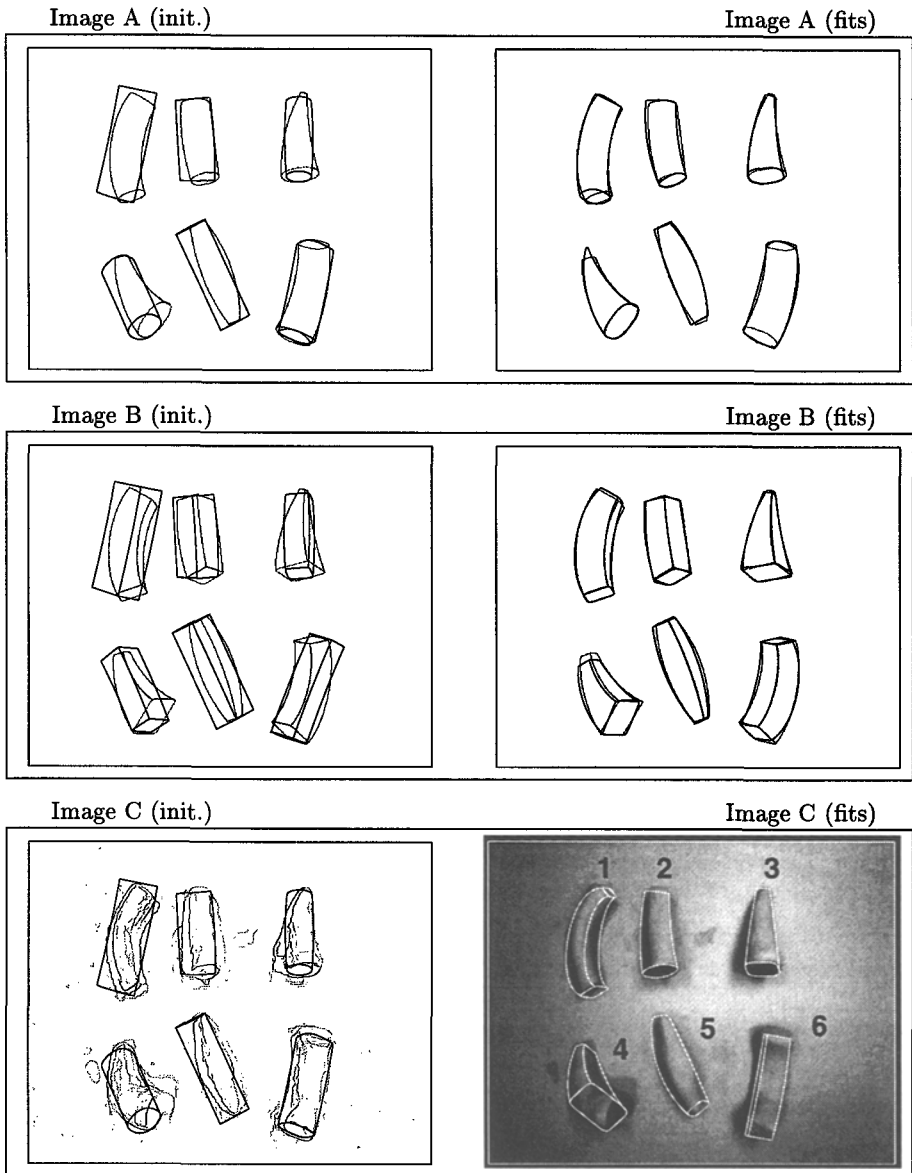
In the first set, shown in Fig. 3, a 512x512 image of six plasticine isolated geons was taken (Image C) and, after the application of a Canny Edge detector, no post-processing was carried out (Image C-left). From it, we created two synthetic edge images with the same geons but with roundish (Image A) and squared (Image B) cross-sections. The initializations (unique for each geon across Images A,B and C and overlapped to the edge images in the left figures) are rather crude; for the synthetic images the right aspect topology was imposed to each instance whereas it was set free for Image C. It can be seen that with synthetic images the results (displayed in the right column) are rather good in both examples. The results with real images are reasonably good too, considering that we intentionally left a high cluttering level; geon #1,#4 and #6, however, have been slightly misfit due to intolerably high noise within the geon body that, in particular, caused a change in topology.

In the second set, shown in Fig. 4, two real 128x128 images of an handset, a mug and a banana were used, and their edge images produced by a Canny edge detector and simple filtering (top); initializations are performed as outlined in Sec. 4 and are displayed in the middle row. In the case of the handset, three good estimates were produced although the edge image is cluttered and incomplete. Except for the mouth piece, which is slightly over-swollen, the essential qualitative features of the geons (such as cross-section and curvature) are extracted. The experiment with the banana is successful too, but, of course, because of the long shading edge running along its body, the fitting yielded squared cross-section. This spurious effect can be overcome only by integrating further information. The mug example was, as expected, a failure because both the initialization and edge image were excessively poor, undistinguishable even to the human eye.

## 6 Conclusion and Future Work


In this paper we have presented a novel method for the fitting of generic geon-like parts to 2D edge image. We developed a new efficient approximated model of superquadrics contour that, in the framework of Model-Based Optimisation, is fitted to the image by a MAP estimation procedure which seeks the best interpretation of the available data in terms of the model. Some experiments have been described that show the validity of the approach.

Single parts are currently semi-automatically selected out from the image, since no method is available for automatically doing it from real images, but we




**Fig. 3.** First set of experiments designed to assess the convergence of the method (see text for details).


Edge Image



Initialisation



Final estimates



**Fig. 4.** Second set of experiments with semi-automatic initialization (see text for details).

are actively investigating into the problem. Furthermore, early experiments using different aspect topologies [5], have shown a substantial increase in robustness. We are also planning to integrate other information, such as shading, to further increase stability and convergence.

*Acknowledgements:* Thanks to A. Fitzgibbon for useful suggestions.

## References

1. R. Bergevin. Primal access recognition of visual objects. Technical Report TR-CIM-90-5, Mc Gill University, Canada, February 1990.
2. I. Biederman. Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94:115–147, 1987.
3. T.O. Binford. Visual perception by computer. In *Proceedings of the IEEE System Science and Cybernetic Conference*, Miami, December 1971.
4. S.J. Dickinson, A.P. Pentland, and A. Rosenfeld. 3-D Shape Recovery Using Distributed Aspect Matching. *IEEE PAMI*, 14(2):130–154, 1992.
5. D. Eggert, L. Stark, and K. Bowyer. Aspect graphs and their use in object recognition. *Annals of Mathematics and Artificial Intelligence*, 13:347–375, 1995.
6. P. Fua and A.J. Hanson. Objective functions for feature discrimination: Applications to semiautomated and automated feature extraction. In *DARPA Image Understanding Workshop*, pages 676–694, 1989.
7. J.E. Hummel and I. Biederman. Dynamic binding in a neural net model for shape recognition. *Psychological Review*, 99:480–517, 1992.
8. L. Ingber. *Adaptive Simulated Annealing*. Lester Ingber Research, Mc Lean, VA, 1993. [ftp.alumni.caltech.edu./pub/ingber/ASA.tar.zip].
9. J.R. Kender and D.G. Freudenstain. What is a "degenerate" view? In *DARPA Image Understanding Workshops*, pages 589–598, 1987.
10. S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.
11. Y.G. Leclerc. Constructing simple stable description for image partitioning. *International Journal of Computer Vision*, 3:73–102, 1989.
12. D. Lowe. Fitting parametrized 3D models to images. *PAMI*, 13(5):441–450, May 1991.
13. D. Metaxas, S.J. Dickinson, R.C., Munck-Fairwood, and L. Du. Integration of quantitative and qualitative techniques for deformable model fitting from orthographic, perspective and stereo projection. In *Fourth International Conference on Computer Vision*, pages 364–371, 1993.
14. A.P. Pentland. Perceptual organization and the representation of natural form. *Artificial Intelligence*, 28:293–331, 1986.
15. N.S. Raja and A.K. Jain. Obtaining generic parts from range data using a multi-view representation. In *Appl. Artif. Intell. X: Machine Vision and Robotics, Proc. SPIE 1708*, pages 602–613, Orlando, FL, April 1992.
16. F. Solina and R. Bajcsy. Recovery of parametric models from range images: The case of superquadrics with global deformations. *IEEE PAMI*, 12(2):131–147, February 1990.
17. K. Wu and M.D. Levine. Recovering of parametric geons from multiview range data. In *IEEE Conference on Computer Vision and Pattern Recognition*, Seattle, WA, 1994.