Superquadrics for Segmenting and Modeling Range Data

Aleš Leonardis, Aleš Jaklič, and Franc Solina

Abstract—We present a novel approach to reliable and efficient recovery of part-descriptions in terms of superquadric models from range data. We show that superquadrics can *directly* be recovered from unsegmented data, thus avoiding any presegmentation steps (e.g., in terms of surfaces). The approach is based on the *recover-and-select* paradigm [10]. We present several experiments on real and synthetic range images, where we demonstrate the stability of the results with respect to viewpoint and noise.

Index Terms—Range image segmentation, recover-and-select paradigm, recovery of volumetric models, superquadrics.

1 Introduction and Motivation

THE significance of detecting geometric structures in images has long been realized in the vision community. One of the primary intentions has been to build primitives that would bridge the gap between low-level features and high-level symbolic structures useful for further processing. Perceptually, the world can be broken down into parts, and the goal of computer vision is to recover from images this part structure (segmentation) and the metric properties of individual parts (shape recovery).

Two types of volumetric models have emerged for such part-level modeling: generalized cylinders [2], [13], [15] and superquadrics [1], [16], [20]. Superquadric models have received significant attention in the vision and robotics community because of their compact representation and robust methods of recovery of individual models.

The motivation for this work is that volumetric part models have been successfully used mostly for shape recovery of individual parts but not for segmentation of images into parts. This is due to the difficulty of simultaneous classification of image elements (grouping) and of model parameter estimation which has been a major obstacle to successful applications that require reliable extraction of volumetric models from the data. For example, rigorous schemes for recovery of volumetric models have been developed, but most of them make the assumption that the segmentation problem has been solved by some other means [20], [19], [21], which obscures the real complexity of the task.

There have been several attempts to segment and recover volumetric models from the data [17], [18], [6], [7]. These approaches usually involve various procedures, mostly applied in a hierarchical fashion, ranging from the estimation of local surface properties, curvature, etc., to more complex, such as symmetry seeking, in order to partition the data into parts that can supposedly be represented with a single volumetric model. Such approaches, in fact, isolate segmentation stage from the representation stage and significant efforts are necessary to combine, usually surface type descriptions into volumetric models. The ability to even identify a set of surfaces as belonging to a common volume is not a trivial task without knowing at least connectedness of surfaces, and preferably surface closure. Moreover, a surface-level description may not be consistent with the volumetric description. Fig. 1 shows an example of an L-shaped object whose volumetric description cannot be obtained by a simple combination of recovered surfaces. Specifically, planar patches of the surface-level description in Fig. 1c cannot be partitioned in a simple way to correspond to volumetric parts in Fig. 1d, since a single surface patch No. 3 in Fig. 1c belongs to two different volumetric models (see Fig. 1d).

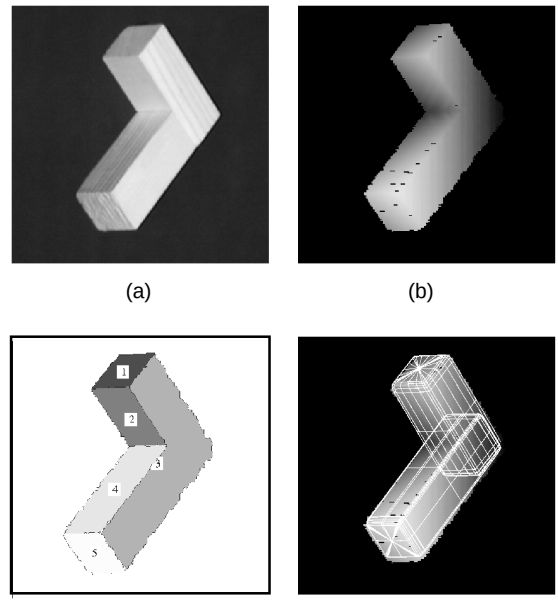

Fig. 1. (a) Intensity image, (b) range image, (c) planar patches of the surface-level description, (d) volumetric models.

In this paper, we demonstrate that a specific set of volumetric models, i.e., superquadrics can be *directly* recovered from range data, which is in contrast to common beliefs that the recovery of volumetric models is possible only after the data has been presegmented using extensive preprocessing. To achieve this goal, we have cast the problem of volumetric recovery [20] in the *recover-and-select* paradigm for the recovery of geometric parametric models from image data [10]. We have been motivated for this new application of the recover-and-select paradigm by the encouraging results that the paradigm achieved on a variety of domains: segmentation and recovery of surface models in range data [12] and curve models [11] in intensity images.

The paper is organized as follows: In Section 2, we shortly describe the geometric volumetric primitives, namely *superquadrics*. A brief outline of the *recover-and-select* paradigm is given in Section 3. In Section 4, we describe specific details that pertain to the recovery of superquadric models. Some of our experimental results are shown in Section 5. In conclusion, we summarize our paradigm and outline the work in progress.

2 Superquadric Models

The criteria for the selection of geometric primitives have been studied extensively by vision researchers (see [3] and the references therein). These criteria not only constrain the choice of the primitives, but impose certain conditions on the model-recovery procedure as well. However, all model-based approaches are restricted, since they cannot model everything possibly present in

• *The authors are with the Computer Vision Laboratory, Faculty of Computer and Information Science, University of Ljubljana, Tržaška c. 25, 1001 Ljubljana, Slovenia.*
E-mail: {Ales.Leonardis, Ales.Jaklic, Franc.Solina}@fri.uni-lj.si.

Manuscript received 8 May 1996; revised 9 June 1997. Recommended for acceptance by P. Flynn.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 105705.

the input data. But what is important, is that the method signals, through the error-of-fit measure, when the models are inadequate to describe the data, so that a different type of model can be invoked.

Superquadrics, which are an extension of basic quadric surfaces and solids, satisfy most of the criteria. They have been considered as volumetric primitives for shape representation in computer graphics [1] and computer vision[1] [16], [20], [21], [14], [6]. The reason being that they are convenient part-level models that can further be deformed and glued together to model articulated objects.

A superquadric surface is defined by the following implicit equation

$$F(x,y,z) \equiv \left(\left(\frac{x}{a_1} \right)^{\frac{2}{\epsilon_2}} + \left(\frac{y}{a_2} \right)^{\frac{2}{\epsilon_2}} \right)^{\frac{\epsilon_2}{\epsilon_1}} + \left(\frac{z}{a_3} \right)^{\frac{2}{\epsilon_1}} = 1, \tag{1}$$

where a_1, a_2, a_3 define the superquadric size, and ϵ_1 and ϵ_2 define a smoothly changing family of shapes from rounded to square.

To recover a superquadric in a general position, the implicit function for general position is used as a base for computing the distance of a point from the superquadric

$$F(x, y, z; \mathbf{\Lambda}) = F(x, y, z; a_1, a_2, a_3, \epsilon_1, \epsilon_2, \phi, \theta, \psi, p_x, p_y, p_z), \tag{2}$$

where ϕ, θ, ψ define the orientation in space, and p_x, p_y, p_z define the position in space. We refer to the set of all-model parameters as $\mathbf{\Lambda} = \{a_1, a_2, \ldots, a_{11}\}$. Since, due to self occlusion, not all sides of an object are visible at the same time, an additional constraint is introduced which imposes the smallest superquadric that fits the given N range points in the least squares sense, leading to the equation

$$G(\Lambda) = \sum_{i=1}^{N} a_1 a_2 a_3 \left(F^{\epsilon_1}(x_i, y_i, z_i) - 1 \right)^2. \tag{3}$$

Note that the term $\left(F^{\epsilon_1}(x,y,z) - 1 \right)^2$ is related to the radial distance d of the point (x, y, z) from the superquadric as [22]

$$F^{\epsilon_1}(x,y,z) - 1 = \frac{d}{|\mathbf{r}_0|} \left(\frac{d}{|\mathbf{r}_0|} + 2 \right), \tag{4}$$

where $\mathbf{r}_0$ is the vector from the center of the superquadric to the closest point on the surface of the superquadric.

The estimation of model parameters $\mathbf{\Lambda}$ which involves nonlinear optimization is described in detail in [20].

3 Recover-and-Select Paradigm

Here we present briefly the concept of the recover-and-select paradigm. For details the reader is referred to [10], [12]. The work that is similar in the aspect of using multiple models and the MDL principle appeared in [5].

On a very general level, one can describe the recover-and-select paradigm as an attempt to solve the following two problems related to the recovery of parametric models from unsegmented data:

1) How do we find image elements that belong to a single parametric model *and*, simultaneously, determine the values of the parameters of the model?
2) How do we achieve a high degree of resistance to outliers?

With respect to the traditional region growing methods, it also attacks the following two problems:

1) How do we find an initial set of data points (a seed) that belongs to a single parametric model taking into account a high degree of noise present in real images?
2) How do we prevent erroneous initial (or intermediate) results to influence (hamper) the development of recovery of other models?

1. Although the term superquadrics is usually used, the models are, in most cases, restricted to superellipsoids only.

The *recover-and-select* paradigm consists of two intertwined stages: model-recovery and model-selection. At the model-recovery stage, a search for parametric models is initiated in regularly placed seeds in an image, and the individual models are allowed to grow which involves an iterative procedure simultaneously combining data classification and parameter estimation. The main advantage of this *classify-and-fit* approach is that the performance of the fitting is constantly monitored, i.e., the procedure dynamically analyzes data consistency allowing the rejection of outliers. The idea to *independently* build all models (as initialized by the statistically consistent seeds) is a crucial one: Due to uncertain initial estimates, any interaction between the partially recovered models would propagate the erroneous estimates through the entire procedure resulting in an overall faulty output.[2] The recovered models, which are treated only as hypotheses that could compose the final description, are then passed to the model-selection procedure. The selection procedure, defined on the MDL principle (minimum description length principle), leads to a quadratic Boolean problem, whose solution is sought with the WTA (winner-takes-all) technique [4]. It selects those models which produce the simplest description, i.e., the one that describes the data with the minimum number of models while keeping the deviations between data points and models low.

In order to achieve a computationally efficient procedure, the model-recovery and model-selection procedures are combined in an iterative way. The recovery of currently active models is interrupted by the model-selection procedure which selects a set of currently optimal models which are then passed back to the model-recovery procedure. This process is repeated until the remaining models are completely recovered.

The choice of specific models, in our case, superquadrics, imposes certain constraints on the recovery of these structures from images, which will be described next.

4 Superquadrics in Recover-and-Select Paradigm

The overall processing scheme is given in Fig. 2. Procedures for seed selection, superquadric fitting, decision making regarding the growth of superquadric models, search for new compatible points (growing), and model selection are given in the following subsections.

4.1 Seed Selection

Initial seeds are placed on the image in the grid-like pattern of windows. An initial seed encompasses a set of range data points R_i in a small window whose size is determined on the basis of scale and can be adaptively changed depending on the task. A superquadric model M_i is fitted to the data set. Next, a decision is made whether all the data points R_i belong to that single model M_i. This decision is based on the average error-of-fit measure

$$\overline{\xi_i} = \frac{1}{|R_i|} \sum_{\mathbf{x} \in R_i} d(\mathbf{x}, M_i), \tag{5}$$

which is compared to a threshold value max_average_model_error (see Section 4.3 on *decision making*). The decision is not critical due to the redundant nature of the paradigm. It just eliminates those seeds that were placed on data sets that cross part boundaries or contain outliers, and, thus, helps reduce the number of seeds at the start of processing.

4.2 Superquadric Fitting

The fitting function (3) can be regarded as an energy function on the space of model parameters. Minimization method can, in general, only guarantee convergence to a local minimum. The starting

2. In addition, the conception of the independent recovery of models offers a possibility for a fully parallelized system.

input: a range image
determine a set of seeds (initial sets R_i)
for all seeds **do**
  fit an SQ model M_i to R_i (estimate parameters)
  if ($\overline{\xi_i}$ < max_average_model_error)
    put the SQ model into the set of currently *active, not fully grown* models
  else
    the seed is rejected
  endif
end for
steps ← length of the square-shaped seed region in pixels / 2
while there are any *active, not fully grown* SQ models **do**
  for i = 1 **to** steps
    for all *active not fully grown* SQ models **do**
      extrapolate / find a set of compatible points N_i
      $N_i = \{\boldsymbol{x}: d(\boldsymbol{x}, M_i) < \text{max\_point\_distance and } \boldsymbol{x} \notin R_i$ and x is 8-connected to $R_i\}$
      if no new compatible points $N_i = \emptyset$
        the SQ model is *fully grown*
      else
        $R_{temp} \leftarrow R_i \cup N_i$
        fit an SQ model $M_{initial}$ to data set R_{temp} using initial estimate of parameters based on center of gravity and central moments
        fit an SQ model $M_{previous}$ to data set R_{temp} using parameters of M_i as an initial estimate
        set M_{temp} to one of the models $M_{initial}$, $M_{previous}$ that has the minimal error-of-fit
        if $\overline{\xi_{temp}}$ < max_average_model_error)
          $R_i \leftarrow R_{temp}$
          $M_i \leftarrow M_{temp}$
        else
          the SQ model M_i is *fully grown*
        end if
      end if
    end for
  end for
  perform selection among all *active* SQ models for the current optimal description
    (only the selected SQ models remain *active*)
  steps ← 2 * steps
end do
output: part-level (SQ) description of a range image

Fig. 2. Superquadric recovery in the *Recover-and-Select* paradigm.

position in the parameter space ($\boldsymbol{\Lambda}_E$) determines to which minimum will the minimization procedure converge. Initial values for both shape parameters, ϵ_1 and ϵ_2 are set to 1.0, which means that the initial model $\boldsymbol{\Lambda}_E$ is always an ellipsoid. Position in world coordinates is estimated by computing the center of gravity of all range points, and the orientation is estimated by computing the central moments with respect to the center of gravity. The z axis of initial model $\boldsymbol{\Lambda}_E$ is oriented along the longest side (axis of least inertia) of the object.

To prevent sudden changes in superquadric orientation due to selection of direction of the z axis for the initial estimate, we use parameters of model M_i recovered from data set R_i as an initial estimate for the model that is to be fitted to the data set $R_i \cup N_i$. On the other hand, such a procedure might force the parameters of the growing model to stay around the local minimum of the initial model recovered from the seed points. To prevent this, we also recover a model from $R_i \cup N_i$ using the initial estimate of superquadrics parameters as described above. We then use the model with the smaller error-of-fit as a final result for the model recovery.

4.3 Decision Making

A decision whether a model should grow further or not depends on the established similarity between the model and the data. If sufficient similarity is established, ultimately depending on the task at hand, we accept the currently estimated parameters together with the current dataset and proceed with the search for a set of compatible points N_i. The question is what could be used as an average error-of-fit measure. Due to its dependence on the superquadric size and shape parameters (a_1, a_2, a_3, ϵ_1, ϵ_2) the algebraic distance (inside-outside function—(1)) is not suitable. Instead we use radial Euclidean distance metric [22]

$$d(\mathbf{x}, \boldsymbol{\Lambda}) = |\mathbf{x}|\,|1 - F^{-\frac{\epsilon_1}{2}}(\mathbf{x}, \boldsymbol{\Lambda})|, \tag{6}$$

where F is given in (1). The sum of distances over all data points belonging to the model determines the total error-of-fit ξ_i of the entire model. This measure is passed to the selection procedure together with the set R_i of data points encompassed by the model.

4.4 Search for New Compatible Points

In accordance with the paradigm, an efficient search for more compatible points is performed in the vicinity of the present border points of the region corresponding to a particular model. The border points are determined on eight-connectedness, and then their distance to the corresponding model is evaluated. Only those points $\mathbf{x}$ that are close enough to the original model ($d(\mathbf{x}, \boldsymbol{\Lambda}) <$ *max_point_distance*) are included in the updated set of points. On this new set of points a new superquadric model-recovery procedure is started.

4.5 Model Selection

As described in the general outline of the method, the model recovery procedure is interrupted by model selection to eliminate superfluous models. An important decision is when to perform the selection. We choose to interrupt the recovery procedure after every n growing operations, and we name the sequence of growing operations between two consecutive interruptions a *growth iteration*. n is set to half the seed size in pixels at the beginning, and then incremented to $2n$ after every selection. For example, if seeds are squares of 4×4 pixels, the first growth iteration consists of two growing operations, which produces regions of size at most 8×8 pixels, after the first selection, the second growth iteration consists of four growing operations, leading to regions of size at most 16×16 pixels, etc. Since the initial seed regions are placed in a grid-like manner, such a procedure ensures sufficient overlap of the models, so that the number of models is reduced during the selection while it is still not too computationally expensive to let the models grow to such a size.

The objective function $F(\mathbf{m})$, which is to be maximized in order to produce the "best" description in terms of models, has the following form:

$$F(\mathbf{m}) = \mathbf{m}^T\mathbf{Q}\mathbf{m} = \mathbf{m}^T\begin{bmatrix} c_{11} & \cdots & c_{1N} \\ \vdots & & \vdots \\ c_{N1} & \cdots & c_{NN} \end{bmatrix}\mathbf{m}, \tag{7}$$

where $\mathbf{m}^T = [m_1, m_2, \ldots, m_N]$. m_i is a *presence variable* having the value one for the presence and zero for the absence of the model

M_i in the final description. The diagonal terms of the matrix **Q** express the cost-benefit value of a particular model M_i,

$$c_{ii} = \mathrm{K}_1 |R_i| - \mathrm{K}_2 \xi_i - \mathrm{K}_3 |P_i|. \quad (8)$$

K_1, K_2, K_3 are weights which can be determined on a purely information-theoretical basis (in terms of bits), or they can be adjusted in order to take into account the signal-to-noise ratio [12]. The $|R_i|$ is the number of points in region R_i, ξ_i the total error-of-fit measure, and $|P_i|$ the number of model parameters, which is equal to 11, and constant, since we use just one type of models. The off-diagonal terms handle the interaction between the overlapping models

$$c_{ij} = \frac{-\mathrm{K}_1 |R_i \cap R_j| + \mathrm{K}_2 \xi_{ij}}{2}. \quad (9)$$

$|R_i \cap R_j|$ is the number of points that are explained by both models. ξ_{ij} corrects the diagonal error terms in case both models are selected

$$\xi_{ij} = \max\left(\sum_{\mathbf{x} \in R_i \cap R_j} d(\mathbf{x}, M_i), \sum_{\mathbf{x} \in R_i \cap R_j} d(\mathbf{x}, M_j) \right). \quad (10)$$

The error terms $d(\mathbf{x}, M_i)$ and $d(\mathbf{x}, M_j)$ are calculated in the area of intersection and correspond to deviations from the ith and jth models, respectively. For details the reader is referred to [10], [12]. In Appendix A we propose a fast greedy algorithm to solve the Quadratic Boolean Problem.

5 Experimental Details and Results

The method has been tested on a variety of real and synthetic range images. Fig. 3 shows a sequence of growing and selection phases on a real-range image. Individual images in the sequence are marked with a letter as follows:

a) input-range image,
b) initial seeds,
c) models after the first growth iteration,
d) remaining models after the first selection,
e) models after the second growth iteration,
f) remaining models after the second selection,
g) models after the third and final growth iteration,
h) remaining models after the third and final selection.

To reduce the computational burden, mostly dictated by the numerical minimization during recovery of superquadrics, the images were subsampled at factor four. The initial size of the seeds was 16×16 pixels, that is 4×4 pixels after the subsampling. The thresholds in the model-recovery procedure were determined on the basis of the analysis of the noise properties of the acquisition process and on the selected subsampling and were kept constant during the segmentation. The actual values used were:

```
max_point_distance = 6.0
max_average_model_error = 2.0
K1 = 1.0, K2 = 0.25, K3 = 0.5
```

To prevent numerical degeneracies in minimization of the superquadric fitting function for subsets of points from planar regions, parameters a_1, a_2, a_3 were limited from below at 1.0 during the minimization by a projection method. In most cases, 15 iterations of Levenberg-Marquardt method were sufficient for convergence. On the average, it took less than 10 minutes to obtain the final results on HP 715/100 workstation running HP-UX or on Intel Pentium 133 MHz running Linux. However, processing time is not critical, since individual models could be recovered in parallel. Besides, the computation of the superquadric fitting function and its derivatives is independent for each range point and can also be parallelized in a straightforward way.

Figs. 4a, c, and e demonstrate the inherent limitation on precision of the basic recover-and-select segmentation paradigm. Points from parts that touch or penetrate the part, which is being recovered, are sufficiently close to the recovered model to get included into the dataset, since the points do not cause the average error-of-fit of the new model to increase over the threshold. This effect of misclassifying data points influences the recovered model, which may, in turn, accumulate even more data points belonging to other parts. To improve segmentation results we applied a simple post-processing method. For each recovered model M_i, we constructed

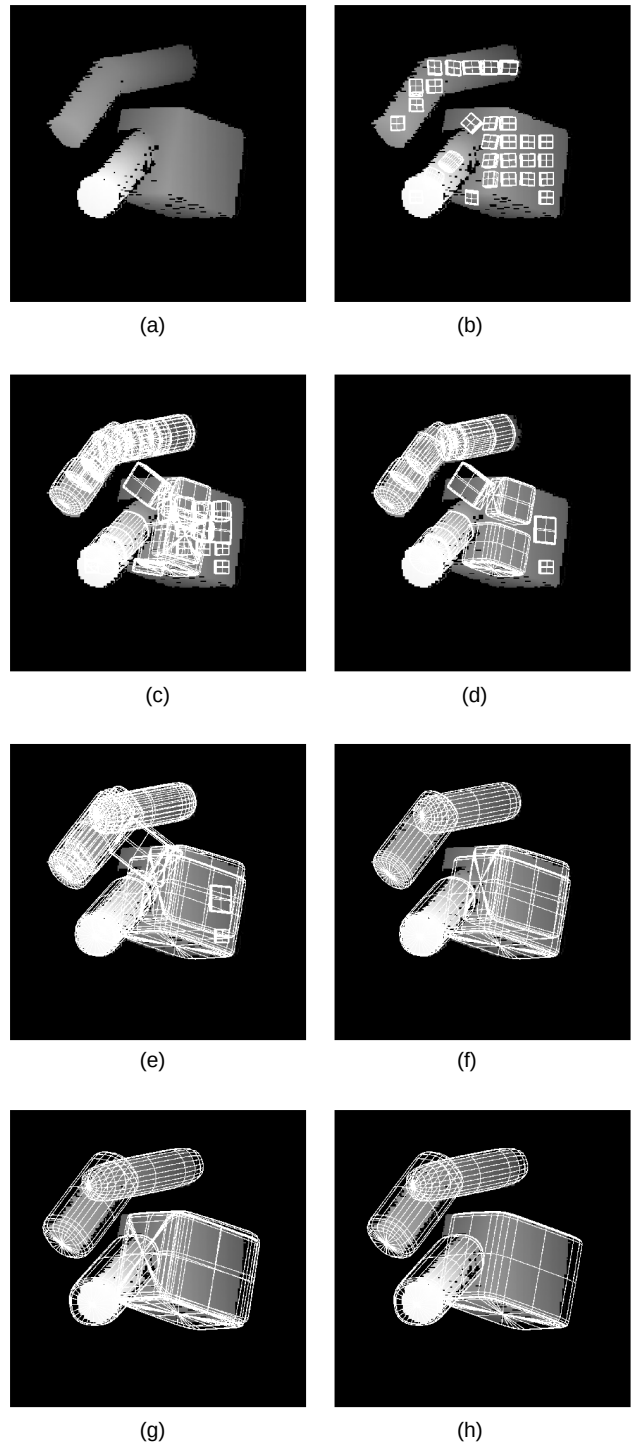

Fig. 3. A typical sequence of growing and selection phases of the segmentation algorithm.

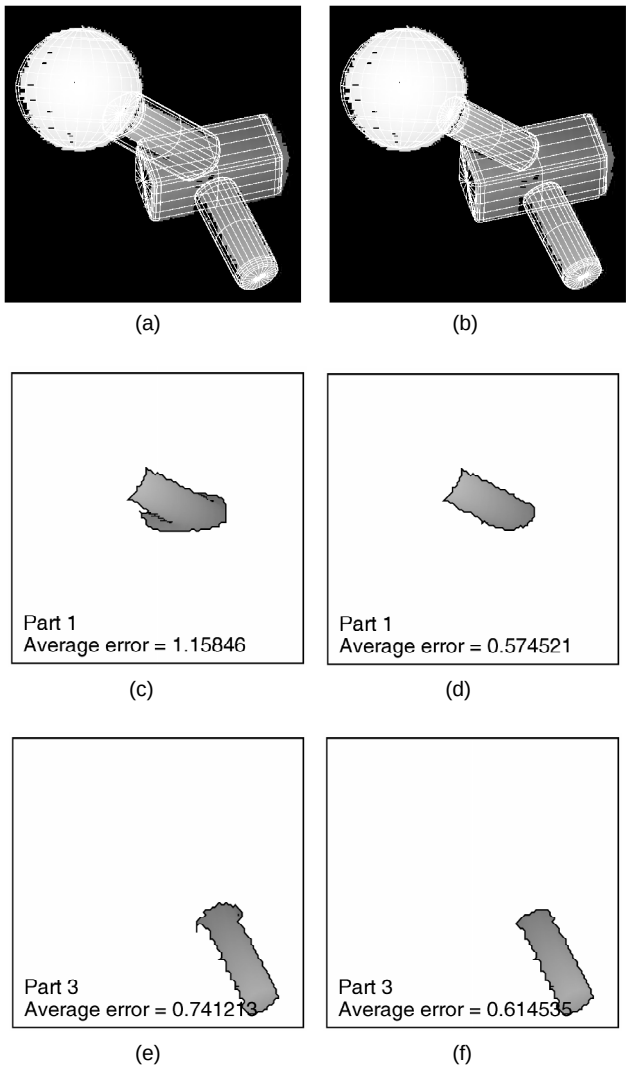

Fig. 4. Segmentation results before postprocessing (a), (c), (e) and after postprocessing (b), (d), (f).

a set $R_i' = R_i \setminus \bigcup_{j=1, j\neq i}^{N} R_j$, thus eliminating data points of the region R_i that are also claimed by other models. From the set R_i' we recovered a new model M_i'. The remaining data points from the set $\bigcup_{j=1}^{N} R_j \setminus \bigcup_{j=1}^{N} R_j'$ were then assigned to the closest model if they were eight-connected to the region corresponding to the closest model. From this set of regions a new set of models was recovered and this procedure of assigning data points to the closest model and subsequent model recovery was repeated until there were no points left that could be assigned to a model. As the Figs. 4b, d, and f show, the postprocessing results in better segmentation and consequently in models with smaller average error. For further details on the rationale behind the postprocessing method and experimental evidence on the effect of postprocessing on error distributions of individual parts, see [9].

To test the stability of segmentation with respect to different views of the same object, we generated a set of synthetic range images. Figs. 5a, b, c, d show the recovered models. The number of recovered models is stable with respect to changing viewpoint. Instead of a direct numerical comparison of superquadric parameters and parameters for position and orientation of superquadrics, which is not trivial, since the mapping from superquadric size and shape parameters to shapes is not one-to-one, we

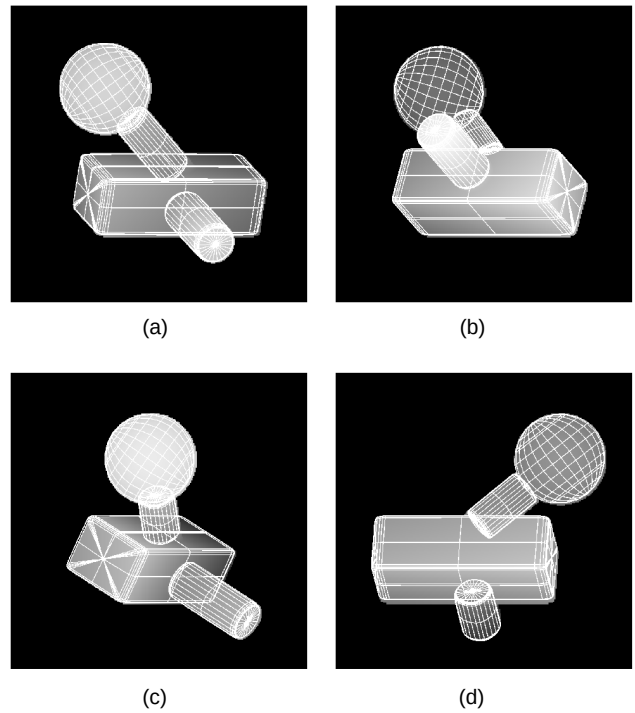

Fig. 5. Segmentation of different range images of the same object (a), (b), (c), (d).

chose a simple application as a test of the stability of the recovered representations. From the recovered superquadrics, we compute an estimate of rigid transformation between two range images based on centers of gravity and moments of inertia of the recovered compositions of superquadrics in each range image. The computation of volume and moments of inertia of individual superquadrics is straightforward with closed-form expressions derived in [9]. The results presented in Figs. 6a and b show that the recovered descriptions are sufficiently stable for pose estimation, despite the fact that completely different sides of objects are visible in the corresponding images. The recovered transformation is in good accordance with known ground truth and, as expected, superior in precision in comparison to the same method applied to raw range data [9].

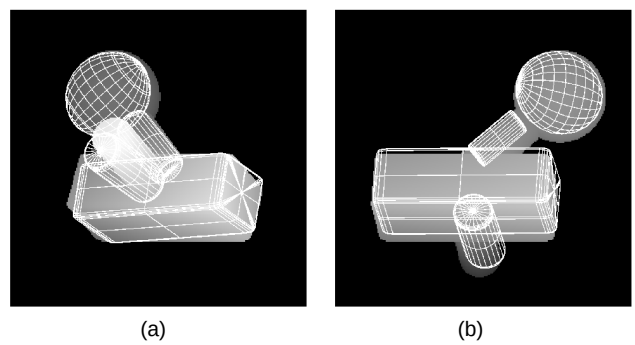

Fig. 6. Recovery of a rigid transformation based on volumetric models: (a) The recovered model from image in Fig. 5a overlaid on the image in Fig. 5b using recovered rigid transformation from the corresponding representations. (b) The recovered model from Fig. 5b overlaid on the image in Fig. 5d.

The procedure is robust with respect to noise (outliers) as demonstrated in Fig. 7. A slightly rotated set of range image data points in Fig. 7a clearly shows outlier-points produced by the

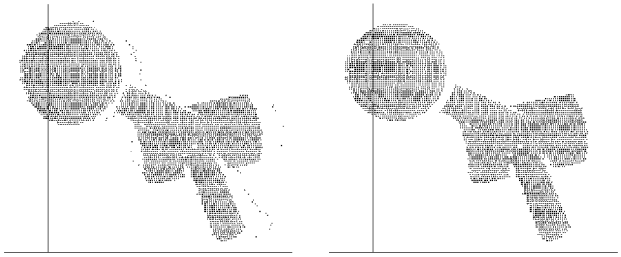

(a) (b)

Fig. 7. The constant monitoring of error-of-fit measure and individual point distance from the current model during growing phases in the segmentation procedure filters out outliers. (a) Raw range image data. (b) Data from the final regions.

range scanner at depth discontinuities. The outliers are not included in the regions produced during the segmentation, as shown in Fig. 7b.

During the course of experimentation, we observed that the method degrades gracefully if the assumptions which are made by the choice of primitives are not met, i.e., when the data cannot be globally modeled by superquadrics. For example, a geometric object, like a croissant in Fig. 8, is described by numerous superquadric models with average error-of-fit close to the *max_average_model_error*, which do not result in a simple description, indicating that different kinds of models should be invoked. On the other hand, the range image of a human-like model is relatively well represented.

Experiments were performed within an object-oriented framework for image segmentation called Segmentor [8]. The framework provides the necessary tools for application of the recover-and-select segmentation paradigm to specific images and models. It is written in C++ and can be obtained in source code from http://razor.fri.uni-lj.si/~alesj.

6 Conclusions

We have successfully combined the two existing methods, namely recovery of superquadric models [20] and the *recover-and-select* paradigm [10]. Thus, we showed that a direct segmentation into part-level volumetric models is possible.

The idea to initiate a redundant set of models in the image not only enables us to cope with inherently unreliable initial estimates, but it also makes possible to successfully deal with *nonlinear* parameter estimation procedures where the convergence to a unique solution is not guaranteed. Starting to develop a redundant set of models from multiple seeds means potentially different initial estimates from which the solution is sought. The sensitivity of the method to the models that get stuck in local minima is thus reduced, since such models get rejected during the model-selection procedure if there are other models that better describe the same data.

We also plan to explore the possibility to use the proposed method for simultaneous recovery of multiple geometric modalities (surfaces and volumes) in different image modalities (range and intensity images). The final description of a scene would then result from a selection procedure employing the MDL principle.

Appendix A
Fast Greedy Algorithm

To simplify the description and analysis of the fast greedy algorithm, we will use alternative notation $f(S)$ for the objective function $F(\mathbf{m})$ in (7)

$$F(\mathbf{m}) = f(S) = \sum_{i \in S} \sum_{j \in S} c_{ij} \qquad S = \{i;\, m_i = 1\}. \tag{11}$$

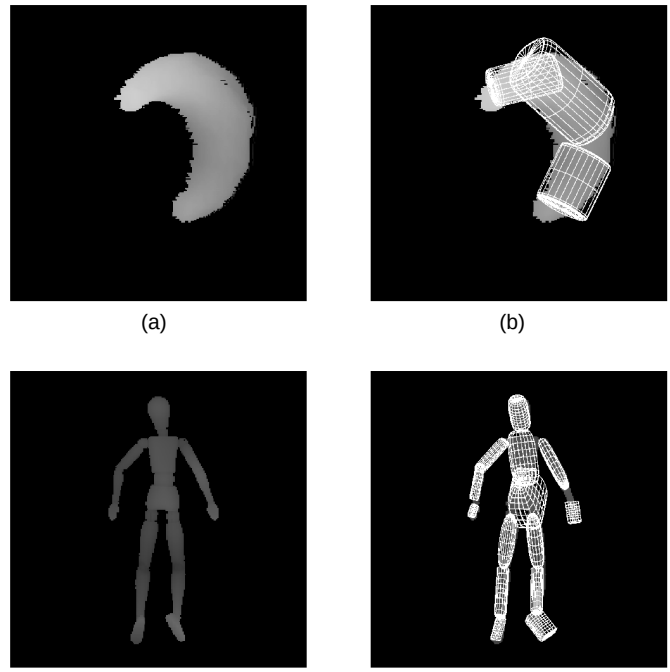

(a) (b) (c) (d)

Fig. 8. Segmentation of objects that cannot be perfectly modeled by superquadrics.

Although the set S contains natural numbers as elements, we will refer to the elements as models, having in mind a natural bijection between the indices and the models.

The greedy algorithm starts with an empty set of selected models S and a set of candidates C that contains all the models. At each iteration of the selection procedure a single model y from the set of candidates C is selected such that it maximizes the $f(S \cup \{x\})$ over all elements x of the set C

$$f(S \cup \{y\}) = \max_{x \in C} f(S \cup \{x\}) \tag{12}$$

and that the quality of the overall description increases by including the model y in the set S

$$f(S \cup \{y\}) > f(S). \tag{13}$$

The set S is then replaced with $S \cup \{y\}$ and the set C with the set $C \setminus \{y\}$. The iteration proceeds until there is no model in set C satisfying conditions (12) and (13) or the set C is empty.

To find a model $x \in C$ satisfying conditions (12) and (13), or to find out that there is no such model, it is sufficient to calculate the change of $f(S)$ if the model x is included into the set of models S

$$\begin{aligned} \Delta f(S, x) &= f(S \cup \{x\}) - f(S) \\ &= c_{xx} + \sum_{i \in S} (c_{xi} + c_{ix}) \\ &= c_{xx} + 2 \sum_{i \in S} c_{ix}\,. \end{aligned} \tag{14}$$

The change can be calculated in $O(|S|)$ time, which leads to $O(N^3)$ worst case time complexity for the algorithm. Further reduction of the time complexity to calculate $\Delta f(S, x)$ can be achieved by observing that the set S is built incrementally and so can be $\Delta f(S, x)$. Suppose that the value of $\Delta f(S, x)$ is known and that we would like to know its relationship to $\Delta f(S \cup \{y\}, x)$. From the definition it follows:

$$\begin{aligned}\Delta f(S \cup \{y\}, x) &= c_{xx} + 2 \sum_{i \in S \cup \{y\}} c_{ix} \\ &= c_{xx} + 2 \sum_{i \in S} c_{ix} + 2c_{xy} \\ &= \Delta f(S, x) + 2c_{xy}.\end{aligned} \tag{15}$$

Thus, the $\Delta f(S \cup \{y\}, x)$ can be calculated in constant time assuming that the $\Delta f(S, x)$ is known.

At the beginning of iteration, the set S is empty so the $\Delta f(\varnothing, x) = c_{xx}$. After the set C is examined and an element y is found, the set C is replaced by $C \setminus \{y\}$ and the $\Delta f(S, x)$ values in the array are incremented by corresponding $2c_{xy}$ terms. Alternatively, the diagonal elements of matrix **Q** can be incremented by $2c_{xy}$ and the matrix dimension reduced by one. We can also interpret the algorithm's incremental calculation of $\Delta f(S, x)$ as a reduction of the problem of size N in time $2N - 1$, that is spent during the search for model y and to update the $\Delta f(S, x)$ values, to the problem of size $N - 1$. This leads to a nonhomogenuous recurrence equation for worst-case time complexity

$$T(N) = 2N - 1 + T(N - 1), \tag{16}$$

with solution in $O(N^2)$.

It is readily apparent from the algorithm that the elements c_{ij} of matrix **Q** can be calculated on a need basis, leading to substantial computational time savings when the number of finally selected models is much lower than N, since we do not have to construct the whole matrix **Q** of size N^2.

```
set function incremental_greedy(matrix Q)
  // input: symmetric matrix Q of dimensions N × N
  // with elements c_ij
  // output: set S of selected models, variable f contains
  // f(S) on return
  S ← ∅
  C ← {i; 1 ≤ i ≤ N}
  f ← 0
  while C is not empty
    find y ∈ C such that c_yy = max_{x∈C} c_xx
    if (c_yy > 0)
      S ← S ∪ {y}
      C ← C \ {y}
      f ← f + c_yy
      for x ∈ C
        c_xx ← c_xx + 2 c_xy
    else
      return S
  end
return S
```

Fig. 9. Fast greedy algorithm of worst case time complexity $O(N^2)$

ACKNOWLEDGMENTS

This work was supported by the Ministry of Science and Technology of Republic of Slovenia (Projects J2-6187, J2-8829), the European Union Copernicus Program (Grant 1068 RECCAD), and by the U.S. – Slovene Joint Board (Project #95-158).

REFERENCES

[1] A.H. Barr, "Superquadrics and Angle-Preserving Transformations," *IEEE Computer Graphics and Applications*, vol. 1, no. 1, pp. 11-23, 1981.

[2] T.O. Binford, "Visual Perception by a Computer," *Proc. IEEE Conf. Systems Science and Cybernetics*, Miami, Fla., Dec. 1971.

[3] M. Brady, J. Ponce, A. Yuille, and H. Asada, "Describing Surfaces," *Computer Vision, Graphics, and Image Processing*, vol. 32, no. 1, pp. 1-28, 1985.

[4] A. Cichocki and R. Unbehauen, *Neural Networks for Optimization and Signal Processing*. New York: Wiley, 1993.

[5] T. Darrell and A.P. Pentland, "Cooperative Robust Estimation Using Layers of Support," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, pp. 474-487, 1995.

[6] F.P. Ferrie, J. Lagarde, and P. Whaite, "Darboux Frames, Snakes, and Super-Quadrics: Geometry From the Bottom Up," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 8, pp. 771-784, 1993.

[7] A. Gupta and R. Bajcsy, "Volumetric Segmentation of Range Images of 3-D Objects Using Superquadric Models," *Computer Vision, Graphics, and Image Processing—Image Understanding*, vol. 58, no. 3, pp. 302-326, 1993.

[8] A. Jaklič, "Segmentor: An Object-Oriented Framework for Image Segmentation," Technical Report LRV-96-2, Computer Vision Laboratory, Faculty of Computer and Information Science, Univ. of Ljubljana, 1996.

[9] A. Jaklič, *Construction of CAD Models From Range Images*. PhD thesis, Faculty of Computer and Information Science, Univ. of Ljubljana, 1997.

[10] A. Leonardis, *Image Analysis Using Parametric Models*. PhD thesis, Faculty of Electrical Eng. and Computer Science, Univ. of Ljubljana, 1993.

[11] A. Leonardis and R. Bajcsy, "Finding Parametric Curves in an Image," *Proc. Second European Conf. Computer Vision*, pp. 653-657, 1992.

[12] A. Leonardis, A. Gupta, and R. Bajcsy, "Segmentation of Range Images as the Search for Geometric Parametric Models," *Int'l J. Computer Vision*, vol. 14, pp. 253-277, 1995.

[13] D. Marr, *Vision*. W.H. Freeman & Co., 1982.

[14] D. Metaxas and D. Terzopoulos, "Shape and Nonrigid Motion Estimation Through Physics-Based Synthesis," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 6, pp. 580-591, 1993.

[15] R. Mohan and R. Nevatia, "Using Perceptual Organization to Extract 3D Structures," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 11, pp. 1,121-1,139, Nov. 1989.

[16] A.P. Pentland, "Perceptual Organization and the Representation of Natural Form," *Artificial Intelligence*, vol. 28, no. 2, pp. 293-331, 1986.

[17] A.P. Pentland, "Recognition by Parts," *Proc. First Int'l Conf. Computer Vision*, pp. 612-620, London, 1987.

[18] A.P. Pentland, "Automatic Extraction of Deformable Part Models," *Int'l J. Computer Vision*, vol. 4, pp. 107-126, 1990.

[19] A.P. Pentland and S. Sclaroff, "Closed-Form Solutions for Physically Based Shape Modeling and Recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 715-729, July 1991.

[20] F. Solina and R. Bajcsy, "Recovery of Parametric Models From Range Images: The Case for Superquadrics With Global Deformations," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 2, pp. 131-147, 1990.

[21] D. Terzopoulos and D. Metaxas, "Dynamic 3D Models With Local and Global Deformations: Deformable Superquadrics," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 703-714, July 1991.

[22] P. Whaite and F.P. Ferrie, "From Uncertainty to Visual Exploration," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 10, pp. 1,038-1,049, Oct. 1991.