INTERNATIONAL JOURNAL OF RESEARCH & ENGINEERING

POSTAL APPLICATIONS

JOURNAL INTERNATIONAL DE RECHERCHE ET TECHNIQUE

APPLICATIONS POSTALES

INAUGURAL ISSUE

LINKING THE WORLD WITH TECHNOLOGY

BOTICS • VISION SYSTEMS • PATTERN RECOGNITION • NEURAL NETWORKS • COMPUTER ARCHITECTURE • INTELLIGENT SYST
GONOMICS • NETWORKING • ECONOMICS • ELECTRONIC COMMUNICATIONS • CODE PRINTING/READING • MANAGEMENT INFO

BOTIQUES • SYSTÈMES DE VISION • RECONNAISSANCE DE FORMES • RÉSEAUX DE NEURONES • ARCHITECTURE DU CALCULA
STÈME INTELLIGENT • TRAITEMENT DES OBJETS • ERGONOMIE • RESEAU • ECONOMIE • SYSTÈMES DE TRANSMISSION • IMPRE

TECHNOLOGIES DANS LE MONDE

PREMIER NUMERO

ACKNOWLEDGEMENTS

The Editors would like to recognize and express their appreciation to everyone involved in the publication of this journal. This includes USA Today, RICOH and San Francisco/Moscow Teleport for supporting an effort for an international transmission for the Journal; MIR Publications, USSR, for inputting the text of the papers into their publishing system; Radio i Svyaz, USSR, for their partnership effort with MIR Publications; and Hemisphere Publications for finalizing the journal information transmitted from the USSR. This team effort has allowed the inaugural issue of the IJRE to truly be an international publication.

Publication Editors

Dmitry Zaidin
MIR Publishers, USSR

Louis Fancher
Hemisphere Publishing Corporation
A Member of the Taylor & Francis Group

Terrence McCaffrey
U.S. Postal Service

Journal Coordinators

Kim Alcorn
Linda Mahoney
Evelena Carroll

The *International Journal of Research and Engineering, Postal Applications* (IJRE) (ISSN 1043-7134) is published semi-annually by the IJRE. For more information complete the enclosed information sheet or call (202) 268-3867.

Postmaster, please send address changes to the: IJRE; Office of Advanced Technology; 475 L'Enfant Plaza, SW; Washington, DC 20260-8118.

REMERCIEMENTS

Les rédacteurs en chef désirent exprimer leur reconnaissance à tous ceux qui ont participé à la publication de ce journal. Et parmi eux, à USA Today, à RICOH et à San Francisco/Moscou Teleport pour avoir collaboré à la transmission internationale pour le Journal ; aux Publications MIR, URSS pour avoir saisi le texte des articles dans leur système d'édition ; à Radio i Svyaz, URSS, pour l'aide qu'ils ont apportée aux publications MIR ; et aux Publications Hemisphere pour avoir mis au point les informations du journal en provenance de l'URSS. Ce travail d'équipe a permis au premier numéro de l'IJRE d'être une publication véritablement internationale.

Rédacteurs en chef des publications

Dmitry Zaidin
MIR Publishers, USSR

Louis Fancher
Hemisphere Publishing Corporation
A Member of the Taylor & Francis Group

Terrence McCaffrey
U.S. Postal Service

Coordinateurs du journal

Kim Alcorn
Linda Mahoney
Evelena Carroll

Le *Journal International de Recherche et Technique, Applications Postales* (IJRE) (ISSN 1043-7134) est une revue semestrielle publiée par l'IJRE. Pour obtenir des informations supplémentaires, veuillez compléter la carte de demande d'informations jointe ou téléphonez au (202) 268-3867.

Monsieur le Receveur des Postes, ayez l'amabilité d'envoyer les changements d'adresse à : IJRE ; Office of Advanced Technology ; 475 L'Enfant Plaza, SW ; Washington, DC 20260-8118.

TABLE OF CONTENTS

Department: Mail Sorting
Subject Area: Electro-Optics and Character Recognition

RECOVERY OF MAIL PIECE SHAPE FROM RANGE IMAGES USING 3-D DEFORMABLE MODELS

Franc Solina
Rozena Bajcsy
GRASP Laboratory,
Department of Computer and Information Science,
University of Pennsylvania,
Philadelphia, PA 19104, USA
Telephone (215) 898-6222

Abstract

Automatic sorting of all classes of mail must take into account different shapes and sizes of mail pieces. A system for classification of mail pieces according to their shape using range images is proposed. Volumetric models of single mail pieces characterized by position, orientation, size and shape parameters are recovered using least squares minimization of a model fitting function. The models used are superquadrics with global deformations. The recovered parameters can serve for different classification schemes which depend on the particular method used for the manipulation of the mail stream. A classification scheme that divides mail pieces into flats, rolls, boxes and irregular mail pieces is presented.

1. Introduction

Automatic sorting of all mail pieces is a difficult problem because mail pieces differ widely in size and shape. In general, one has to know the location, orientation, size and shape of a mail piece to initiate the right handling procedure. Computer vision as a method for locating and describing objects without direct physical contact seems to be the right approach to do that in a fast and reliable manner.

Computer vision has been successfully applied in many industrial applications. The methods used in the majority of these industrial vision systems, however, cannot be applied to the problem of mail piece classification. Most of the so-called "model-based" object recognition vision systems rely on a set of precise and rigid models of all objects that the system is capable to recognize. On the basis of some predefined features recognized in the images, hypotheses are selected whose models are projected onto the image, to find if they match with the rest of the image features (for an example of such a method see [Bolles and Horaud, 1986]). This approach is possible only in tightly controlled environments where the shape of objects is well defined in advance and the number of different objects is small. This is clearly not the case with mail pieces, which come in a variety of sizes and shapes.

When precise object models are not available, traditional computer vision advocates a stepwise reduction of data [Marr, 1982]. First, low level shape models such as edges, corners and surface patches are computed locally. Due to the small granularity of these models, a large number of such models is required even for description of simple scenes. To facilitate any reasoning about the scene, these local models must be merged into larger entities, generalized cylinders being the most popular model, that correspond to individual parts or objects. In the case of mail pieces, this merging of local models is difficult and error-prone because mail pieces, like natural scenes, do not conform to perfect geometrical shapes due to rounded edges, distorted corners, bulging sides and wrinkled surfaces. It is difficult to reliably distill the most appropriate large scale model out of the large set of noisy low level models. We believe that such bottom-up approach for classifying mail pieces from range images according to shape would not be successful.

At the Second United States Postal Service Advanced Technology conference [Solina and Bajcsy, 1986] we proposed to directly apply volumetric models of larger granularity—superquadric models with deformations—for mail piece interpretation. Additional information which is needed for patching up the missing information and rejecting the erroneous local information can be supplied by implicit constraints of compact volumetric models provided they have the right granularity. Superquadrics provide a shape vocabulary well suited for describing real world shape on the level of parts that correspond to the human notion of parts [Pentland, 1986]. One superquadric model is in general sufficient for modelling a single mail piece. The parameters of these superquadric models can be recovered directly from range images [Solina, 1987]. The recovery can be intuitively explained in terms of extrinsic and intrinsic forces. Extrinsic or image forces mold the models whose possible shapes and arrangement of model parameters are governed by internal forces. Since superquadrics can model most common geometrical shapes, such as parallelepipeds, cylinders, ellipses and shapes in between, they are appropriate models for most mail pieces.

The rest of the paper is organized as follows: Section 2 is a short introduction to superquadrics, Section 3 describes model recovery from range images, and Section 4 is on classification. Discussion in Section 5 compares the advantages and deficiencies of the method and points to future extensions, especially to the use of such shape recovery for segmentation. Although we have used range images from various sources, all of the range images used in this paper, however, were obtained with a laser imager built at University of Pennsylvania [Tsikos, 1987].

2. Superquadrics

Superquadrics are an extension of basic quadric surfaces and solids. Superquadrics have been considered as primitives for shape representation in computer graphics [Barr, 1981] and computer vision [Pentland, 1986, Bajcsy and Solina, 1987, Boult and Gross, 1987]. Superquadrics can be compared to lumps of clay that can be further deformed and glued together into realistic looking models as is nicely demonstrated by Pentland's *Supersketch* graphics system [Pentland, 1986].

A superquadric surface is defined by the equation

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

When both ϵ_1 and ϵ_2 are 1, the surface defined is an ellipsoid or, if a_1, a_2, a_3 are all equal, a sphere. When $\epsilon_1 \ll 1$ and $\epsilon_2 = 1$, the superquadric surface is shaped like a cylinder. Parallelepipeds are produced when both $\epsilon_1 \ll 1$ and $\epsilon_2 \ll 1$. Modelling capabilities of superquadrics can be enhanced by deforming them in different ways, such as tapering and bending [Solina, 1987]. Some examples of superquadric models are in Figure 1.

The function in equation (1) is called the inside-outside function because it determines where a given point (x, y, z) lies relative to the superquadric surface. If $F(x, y, z) = 1$, point (x, y, z) lies on the surface of the superquadric. If $F(x, y, z) > 1$, the corresponding point lies outside and if $F(x, y, z) < 1$, the corresponding point lies inside the superquadric.

The inside-outside function (1) defines the superquadric surface in an object centered coordinate system (x_S, y_S, z_S). 3-D points in range images, on the other hand, are expressed in an image coordinate system. To recover a superquadric in general position, an inside-outside function for general position is used where the relation between the image coordinate system and the object centered coordinate system is described with a homogeneous transform **T**. We express the elements of the rotational part of transformation matrix **T** with Euler angles (ϕ, θ, ψ) [Paul, 1981]. The inside-outside function for superquadrics in general position is then

$$F(x, y, z) = F(x, y, z; a_1, a_2, a_3, \epsilon_1, \epsilon_2, \phi, \theta, \psi, p_x, p_y, p_z). \quad (2)$$

This expanded inside-outside function has 11 parameters; a_1, a_2, a_3 define the superquadric size; ϵ_1 and ϵ_2 are shape parameters; 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 $\Lambda = \{a_1, a_2, ..., a_{11}\}$.

3. Recovery of Superquadrics

For shape recovery of single mail pieces we assume that just a single mail piece is present in the range image at a given moment. Such singulation of mail pieces can be achieved by mechanical means. The range points that represent the supporting surface in the image can be removed from the set of all range points by fitting a plane to the supporting surface and removing all points on or close to that plane. We can assume that remaining range points lie on the surface of the single mail piece. Suppose we have N 3-D surface points (x_W, y_W, z_W) which we want to model with a superquadric. We want to vary the 11 parameters a_j, $j = 1, ..., 11$ in equation 2 to get such values for a_j's that most of the 3-D points will lie on, or close to the model's surface. There will probably not exist a set of parameters Λ that perfectly fits the data. Finding the model Λ for which the distance from points to the model's surface is minimal is a least-squares minimization problem. Since due to self occlusion, not all sides of an object are visible at the same time, we have to introduce an additional constraint. Among all possible solutions we want to find the *smallest* superquadric that fits the given range points in the least squares sense. We define the following function which has a minimum corresponding to the smallest superquadric that fits a set of 3-D points *and* a function value for surface points which is known before minimization

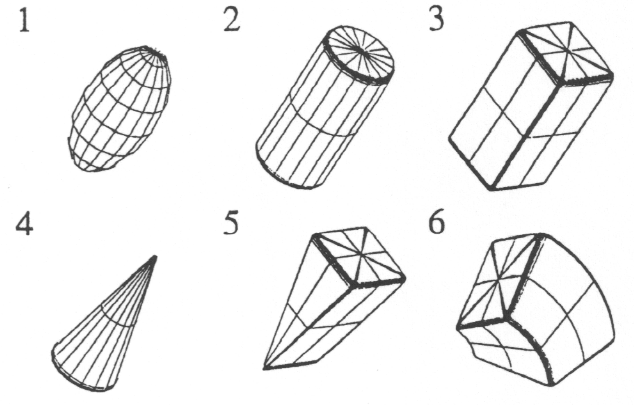

FIGURE 1

$$R = \sqrt{a_1 a_2 a_3}\,(F - 1). \quad (3)$$

Since, for a point (x_W, y_W, z_W) on the surface of a superquadric

$$R(x_W, y_W, z_W; a_1, ..., a_{11}) = 0, \quad (4)$$

we have to find

$$G = \min \sum_{i=1}^{N} [R(x_{W_i}, y_{W_i}, z_{W_i}; a_1, ..., a_{11})]^2. \quad (5)$$

Since R is a nonlinear function of 11 parameters a_j, $j = 1, ..., 11$, minimization must proceed iteratively. Given a trial set of values of model parameters Λ_k, we evaluate equation (3) for all N points and employ a procedure to improve the trial solution. The procedure is then repeated with a set of new trial values Λ_{k+1} until the sum of least squares (5) stops decreasing, or the changes are statistically meaningless. In most cases 15 iteration are more than sufficient. We use the Levenberg-Marquardt method for non-linear least squares minimization [Press et al, 1986] since first derivatives $\delta R/\delta a_i$ for $i = 1, ..., 11$ can be computed analytically.

Only very rough initial estimates of object's true position, orientation, and size suffice to assure convergence to a local minimum that corresponds to the actual shape. This is important since these parameters can be estimated only from the range points on the visible side of the object and hence the estimates cannot be very accurate to begin with. Initial values for both shape parameters, ϵ_1 and ϵ_2 can always be 1, which means that the initial model Λ_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. We orient the initial model Λ_E so that the axis z of the object-centered coordinate system lies along the longest side (axis of least inertia) of the object. This is because bending and tapering deformations normally affect objects along their longest side. Estimates for model's size are simply the extent of range points along the new coordinate axis. Figures 2, 3, 4, and 5 show examples of model recovery for each of the four proposed classes of mail: a box, a roll, a flat and an irregular postal piece. The poor fit of the first undeformed model in Figure 5 can be improved by applying global deformations of superquadric models.

Deformed superquadrics can be recovered using the same technique as for the recovery of non-deformed superquadrics. The only difference is that some additional parameters describing deformations must also be recovered. Deformations such as simplified tapering, bending and twisting require just a few additional parameters [Barr, 1984]. Shape deformation is a function D which explicitly modifies the global coordinates of points in space

$$\mathbf{X} = D(\mathbf{x}), \qquad (6)$$

where $\mathbf{x}$ are the points of the undeformed solid and $\mathbf{X}$ are the corresponding points after deformation. Both $\mathbf{x}$ and $\mathbf{X}$ are expressed in the object centered coordinate system. Any translation or rotation is performed after the deformation. A tapered and bent model can be described schematically as

$$\mathit{Trans}(\mathit{Rot}(\mathit{Bend}(\mathit{Taper}(\mathbf{x})))). \qquad (7)$$

The structure of this model seems to have perceptual significance. When forming a model, one has to proceed from inside out, first taking an undeformed model, reshaping it and putting it in the right general position in world coordinates. When recovering a model, the operation must go in the reverse direction—one has to recover first the position and orientation of the modeled part in space, before its shape can be recovered. The inside-outside

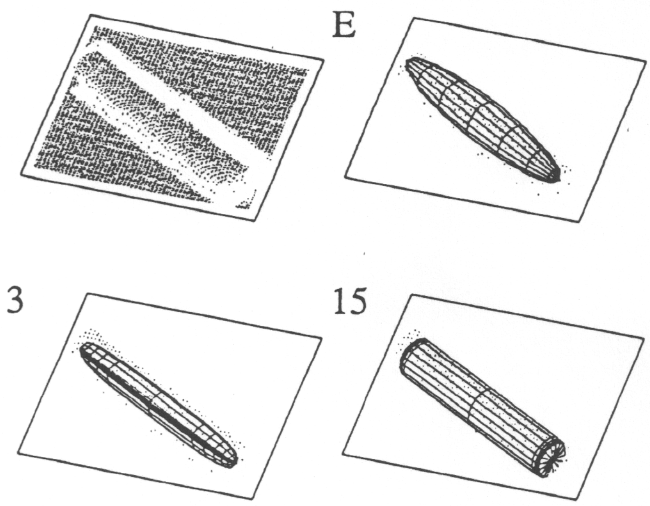

FIGURE 3

function of the deformable model that we use has 4 additional parameters; two parameters for tapering along axis z for independent tapering in direction of axis x and y, and two bending parameters, α—an angle to define the orientation of the bending plane around axis z and β—the actual bending angle

$$F(x, y, z) = F(X, Y, Z; a_1, \ldots, a_{11}, K_x, K_y, \alpha, \beta). \qquad (8)$$

The fitting function (3) can be regarded as an energy function on the space of model parameters. Minimization methods can, in general, only guarantee convergence to a local minimum. The starting position in the parameter space (Λ_E) determines to which minimum the minimization procedure will converge. We have to assure that the minimization procedure does not get stuck in a shallow local minimum, but finds a deep minimum instead. Solutions corresponding to shallow local minima are avoided by adding noise to the value of the fitting function of the accepted model at each iteration during model recovery. When the noise contaminated value is com-

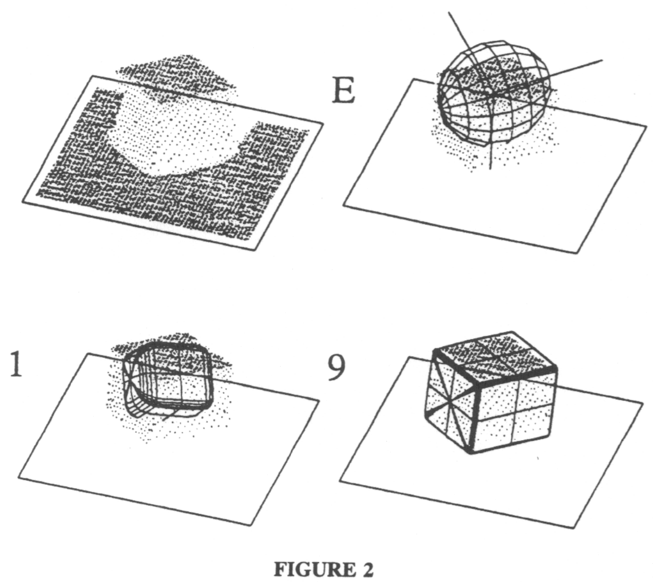

FIGURE 2

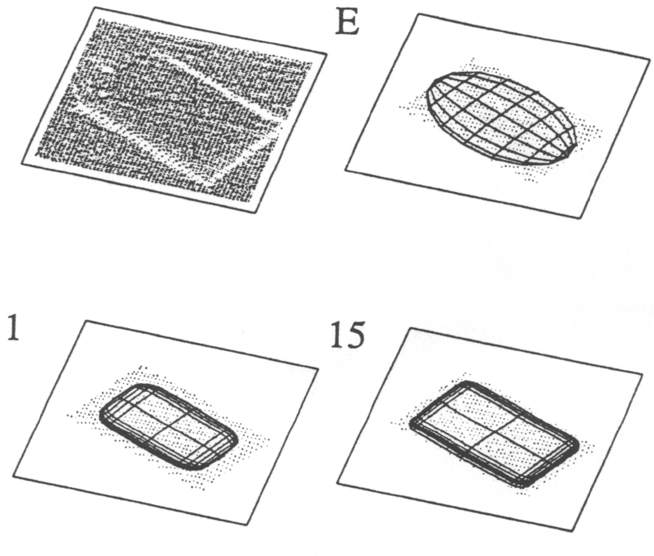

FIGURE 4

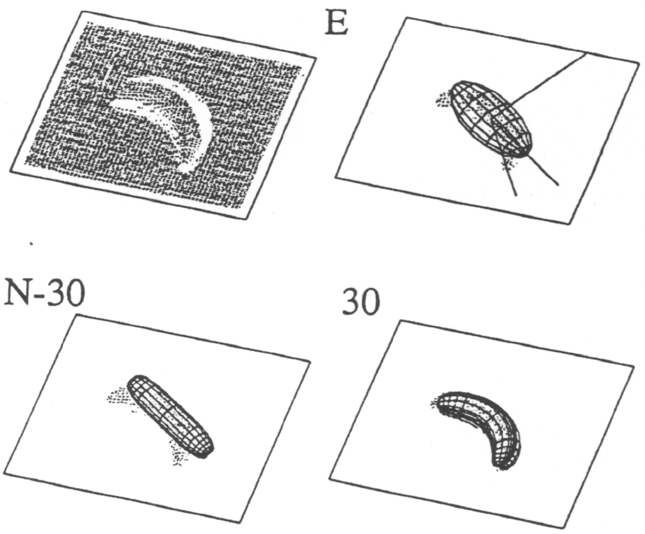

FIGURE 5

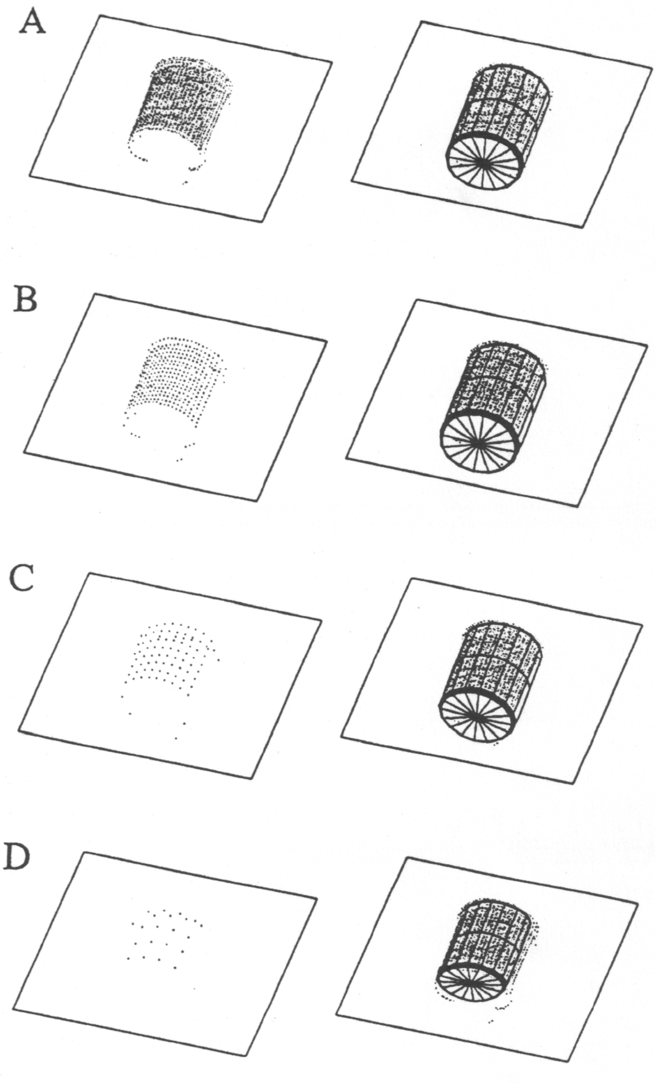

FIGURE 6

pared with the fit of the model in the current iteration, a new model might be selected although it does not fit the data as well as the old one. This stochastic technique of introducing "jitter" into the fitting procedure resembles simulated annealing—see Figure 7.

The most time consuming part of model recovery is the evaluation of the fitting function and of all of its partial derivatives for every input range point during each iteration. Since the sum of least squares (5) is a monotonically increasing function it pays off to monitor the partial sum after each addition. As soon as the sum is larger than the sum of least squares of the accepted and noise contaminated model, it makes no sense to continue. The model cannot be accepted.

A substantial speed-up can be achieved by subsampling the original range map and using a series of coarse to fine grids during minimization. The models recovered from coarser range maps can still be a good representation of the imaged object (Figure 6). During iterative model recovery, the fitting function typically drops very fast until it reaches a plateau. Further iterations gain no substantial improvements in fit (Figure 7). Fast and efficient computation can be done on a hierarchy of coarser grids. We implemented a multi-resolution model recovery scheme which starts on a very coarse range map. Once no improvement in fit is made, the minimization continues on a denser range map until the finest or the original range map is reached (Figure 7). The multi-resolution method is faster because it takes less time for computation in each iteration. The number of iterations, however, may not be smaller. The number of iteration might be even larger because, during multi-resolution recovery, the models for sparser range maps converge to somewhat different sets of parameters than the resulting set of parameters on the finest grid. An implementation of the recovery procedure on a fine grained parallel architecture would be straightforward since the evaluation of the fitting function and its partial derivatives is locally independent.

Recovery of models shown in this paper, where the number of range points for each model is on the order of several hundred, takes about 20 seconds of CPU time on a VAX 785

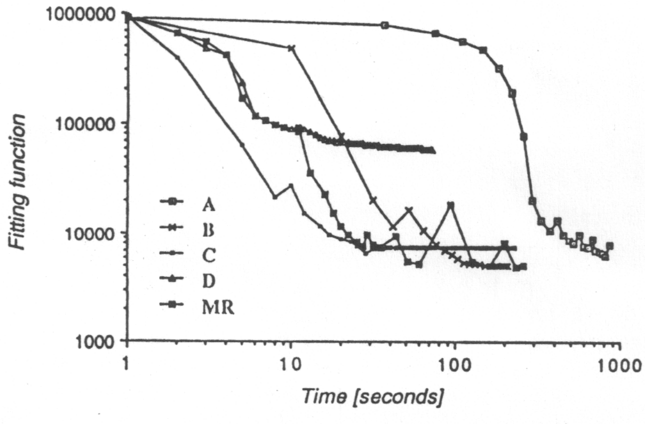

FIGURE 7

computer. We tested the consistency of the recovery method by taking range images of the same object in different positions and orientations. The recovered models compare favorably. For details and issues concerning consistency, stability and ambiguity of model recovery see [Solna, 1987].

4. Classification

The results of model recovery procedure are the positions and orientations of models, as well as their size and shape parameters. The parameter space is continuous, but some sets of parameters correspond to common geometric primitives such as parallelepipeds and cylinders. For those classes or categories, the within-category parameter differences look smaller than between-category parameter differences even when they are of the same size. By mapping symbols on the continuous superquadric parameter space it is possible to define distinctive classes of objects. For manual handling of mail pieces, four classes of mail pieces evolved, reflecting natural breaks in the structure of mail shapes. The four classes are flats, boxes, tubes and irregular mail pieces as a class of all mail pieces that do not belong to any of the first three classes. We could define several different classification schemes based on the recovered model parameters. Although the nature of the manipulation equipment for automatic mail handling must be taken into account, it makes sense to consider the above mentioned classification, since it already reflects some sensible criteria for material handling. We can do more than just classify mail pieces into groups according to shape. From the two shape parameters and size of the model, the actual radius of curvature on edges can be computed, to evaluate the sharpness or roundness of edges—parameters that are important for automatic manipulation. The least-squares minimization method for model recovery provides us with a measure of how well the model represents the actual object. This least-squares residuum G from equation (5) is a measure of goodness of fit which plays an important role in classification. The recovered model might be shaped as a parallelepiped but the goodness of fit can be very poor—indicating that the actual object is irregular (a film mailer, for example). The recovered model in this case is only a rough approximation for the actual shape but sufficient, for example, to grasp the object. Poor mail piece representation can be also an indication that global deformations should be applied for better shape description. Mail pieces that are normally lumped together into the class of irregular postal pieces can be better described with global deformations such as tapering and bending of basic superquadric models as shown in Figure 5.

Based on the size of the model in all three axes (a_0, a_1, a_2), the two shape parameters (ϵ_1, ϵ_2) and goodness of fit G, we designed the classification rules in Table 1.

TABLE 1

input $a_1, a_2, a_3, \epsilon_1, \epsilon_2, G,$
    $T_{FLAT}, W_{FLAT}, S_{BOX}, D_{ROLL}, L_{ROLL}, RES$

if $G > RES$
    then *mail piece is an IRREGULAR MAIL PIECE*

else if $[(a_1 < T_{FLAT}$ and $a_2, a_3 > W_{FLAT})$ or
    $(a_2 < T_{FLAT}$ and $a_1, a_3 > W_{FLAT})]$
    then *mail piece is a FLAT*

else if $[a_1 > S_{BOX}$ and $a_2 > S_{BOX}$ and $a_3 > S_{BOX}$ and
    $\epsilon_1 < 0.5$ and $\epsilon_2 < 0.5]$
    then *mail piece is a BOX*

else if $[a_1 > D_{ROLL}$ and $a_2 > D_{ROLL}$ and $a_3 > L_{ROLL}$ and
    $\epsilon_1 < 0.5$ and $\epsilon_2 > 0.5]$
    then *mail piece is a ROLL*

else *mail piece is an IRREGULAR MAIL PIECE*

5. Discussion

The proposed shape vocabulary is intended for rough description of objects, suitable for shape classification of mail pieces. Objects whose occluded side is not symmetrical to the visible side might not be represented adequately. Although basic deformations of tapering and bending are often sufficient, they do not cover all possible cases. A larger number of different deformations could be used, but that would require a larger number of deformation parameters. The model that we use seems to be a good compromise, not too complicated to recover, but adequate for grasping and handling of objects.

Nonuniform range data density and a large number of singular views in range images was another problem that we faced. Nonuniform range point density causes regions with higher density to have more influence on the shape of the recovered model than parts with lower density. However, the model recovery method is quite robust in this regard—note that no range data is available from occluded parts to begin with. A note of caution is still in order. When the shape recovery system is presented with a singular view such as, for example, when only one face of a cube is seen from the given view points, a very thin parallelepiped which fits to that face is recovered. Unfortunately, images taken with a passive range imager that uses triangulation have more singular views than normally associated with intensity images. The larger the distance (angle) between the source of illumination and the camera of the range imager, the better the accuracy of range measurements, but the more likely is a view of an object singular, since to get a range point from a surface patch, that patch must be illuminated by the laser and seen by the camera at the same time. One way of eliminating singular views is to combine several range images. Currently, we solve the problem of singular views by projecting the visible points onto the supporting surface and use them together with the rest of range points for fitting a model to them. This resolves the lack of data points due to the singular view but can distort some objects, such as cylinders lying on their side (the hidden side of a cylinder is flat instead of cylindrical). The problem is currently solved by recovering a model for both cases; first, we fit a model only to the set of visible points, and then again to the set of visible points and their projections onto the supporting surface. If the goodness of fit is about the same in both cases, the model with the *larger volume* is selected. Otherwise, the model with the *better goodness of fit* is taken.

Another way of resolving singular views is by taking into account the structure of the surrounding scene. Objects rest on some support, they can touch each other but normally they do not penetrate each other. Witkin, Fleicher and Barr developed an elegant method for describing geometrical constraints between parts or objects in terms of energy [Witkin et al, 1987]. If this paradigm would be used for shape recovery of a single mail piece, a sum of energy terms would have to be minimized, one of them being the fitting function for the object's model, while the other terms would define constraints on different geometrical relations, such as that the object's model must touch the supporting surface in at least three non-colinear points.

In this paper we concentrated only on shape recovery of single mail pieces. When several possibly overlapping mail pieces are present in the scene, the scene must be segmented—each mail piece should be represented with a single model. Segmentation, however, depends on the shape of individual parts. But to recover the shape of parts, one must know which range points belong to the same part. Segmentation and shape recovery are hence intimately linked. Because of this interdependence, we believe that the segmentation and shape recovery of individual parts should be done simultaneously. Segmentation as an integral part of part-level shape recovery is possible [Solina, 1987; Pentland, 1987] but must be made more reliable and robust before practical applications can be considered.

Acknowledgement

This work was supported by United States Postal Service contract 104230-87-H-0001/M-0195 and DARPA grant N00014-88-K-0630.

References

BAJCSY, R. AND SOLINA, F. 1987. "Three dimensional shape representation revisited," in *Proceedings First International Computer Vision Conference*, London, England, 231-241.

BARR, A. H. 1981. "Superquadrics and angle-preserving transformations." *IEEE Computer Graphics and Applications*, vol. 1, pp. 11-23.

BARR, A. H. 1984 "Global and local deformations of solid primitives." *Computer Graphics*, vol. 3, pp. 21-30.

BOLLES, R. C. AND HORAUD, P. 1986. "3DPO: A three-dimensional part orientation system." *International Journal of Robotics Research*, vol. 5, no. 3, pp. 3-26.

BOULT, T. E. AND GROSS, A. D. 1987. "Recovery of superquadrics from depth information," in *Proceedings Spatial Reasoning and Multi-Sensor Fusion Workshop*, St. Charles, IL, pp. 128-137.

MARR, D. 1982. *Vision*. (Freeman, San Francisco).

PAUL, R. 1981. *Robot manipulators*. (MIT Press, Cambridge, MA).

PENTLAND, A. P. 1986. "Perceptual organization and the representation of natural form." *Artificial Intelligence*, vol. 28, no. 3, pp. 293-331.

PENTLAND, A. P. 1987. "Recognition by parts," in *Proceedings First International Computer Vision Conference*, London, England, pp. 612-620.

PRESS, W. H. et al. 1986. *Numerical recipes*. (Cambridge University Press, Cambridge).

SOLINA, F. AND BAJCSY, R. 1986. "Modelling of mail pieces with superquadrics," in *Proceedings Second USPS Advanced Technology Conference*, Washington, DC, pp. 472-481.

SOLINA, F. 1987. "Shape recovery and segmentation with deformable part models." Ph.D. dissertation, University of Pennsylvania, Philadelphia, USA.

TSIKOS, G. 1987. "Segmentation of 3-d scenes using multimodal interaction between machine vision and programmable mechanical scene manipulation." Ph.D. dissertation, University of Pennsylvania, Philadelphia, USA.

WITKIN, A. P. et al. 1987. "Energy constraints on parameterized models," in *Proceedings SIGGRAPH-87*, Anaheim, CA.

Franc Solina

Franc Solina received his Dipl. Ing. and M.S. degrees in Electrical Engineering at University of Ljubljana, Yugoslavia in 1979 and 1982, respectively. He received his Ph.D. degree in Computer Science at University of Pennsylvania, Philadelphia, USA in 1987 and is now an assistant professor of computer science at University of Ljubljana. His main research interests are models for shape representation and segmentation in vision. Franc Solina is a member of American Association of Artificial Intelligence.

Ruzena Bajcsy

Ruzena Bajcsy received her first Ph.D degree in Electrical Engineering at the Slovak Technical University in Bratislava, Czechoslovakia in 1967 and her second Ph.D. degree in Computer Science at Stanford University in 1972. She is a professor of computer science at University of Pennsylvania where she heads the GRASP Laboratory and is currently also the chairman of the Computer and Information Science Department. Her main research interests are active perception, sensory integration and visual segmentation. Ruzena Bajcsy is an associate editor of IEEE Transactions on Pattern Analysis, Machine Intelligence, Pattern Recognition Letters, Robotic Systems Journal and Computer Vision, Graphics and Image Processing Journal.

Explanations of Figures

Table 1. Classification rules for dividing mail pieces according to their shape into the following four classes: BOXES, ROLLS, FLATS and IRREGULAR MAIL PIECES. The input parameters are the recovered superquadric parameters a_1, a_2, a_3, ϵ_1, ϵ_2 and the residuum of the least square fit G. A mail piece is irregular either, if the goodness of fit G of the recovered model is not sufficient, as set by constant RES, or if the dimensions are not met by the preset size limits. Constant T_{FLAT} sets the maximal thickness of a flat. Constant W_{FLAT} sets the minimal width and length of a flat. Constant S_{BOX} sets the maximal size of a box along any dimension. Constant D_{ROLL} sets the minimal diameter of a roll and constant L_{ROLL} sets the minimal length for a roll. Only lower bounds on dimensions of particular classes of mail were considered in this classification table although upper bounds could be introduced also, depending either on regulations or capabilities of mail handling equipment.

Figure 1. Superquadrics (1: $\epsilon_1 = \epsilon_2 = 1$; 2: $\epsilon_1 = 0.1, \epsilon_2 = 1$; 3: $\epsilon_1 = \epsilon_2 = 0.1$) and deformed superquadrics (models 4 and 5 are tapered, model 6 is bent).

Figure 2. Shape recovery of a parallelepiped-like object (a box). On top is the original range image followed by the recovery sequence showing the initial estimate (E) and models after 1st and 9th iteration during which all 11 model parameters were adjusted. The above model recovery sequence took only about 20 seconds on a VAX 785 computer. The recovered parameters were: $a_1 = 54$ mm, $a_2 = 55$ mm, $a_3 = 63$ mm, $\epsilon_1 = 0.1$, $\epsilon_2 = 0.1$. According to the decision rules in Table 1, this object was classified as a BOX of *width* = 126 mm, *depth* = 108 mm and *height* = 110 mm.

Figure 3. Shape recovery of a tube formed out of a rolled up newspaper. On top is the original range image followed by the recovery sequence showing the initial estimate (E) and models after 3rd, and 15th iteration when all 11 model parameters were adjusted simultaneously. The recovered parameters were: $a_1 = 24$ mm, $a_2 = 28$ mm, $a_3 = 137$ mm, $\epsilon_1 = 0.1$, $\epsilon_2 = 1.1$. According to the decision rules in Table 1, this object was classified as ROLL of *length* = 274 mm and *diameter* = 52 mm.

Figure 4. Shape recovery of a mail flat. On top is the original range image followed by the model recovery sequence showing the initial model estimate (E) and models after the 1st and 15th iteration when all 11 model parameters were adjusted simultaneously. The recovered parameters were $a_1 = 5.5$ mm, $a_2 = 63$ mm, $a_3 = 98$ mm, $\epsilon_1 = 0.2$, $\epsilon_2 = 0.3$. According to the decision rules in Table 1, this object was classified as a FLAT of *length* = 196 mm, *width* = 126 mm and *thickness* = 11 mm.

Figure 5. Shape recovery of an irregular mail piece—a banana. A banana is certainly a highly unusual mail piece, however, an article published in the *Hartford Courant* on 10 October 1987 reported that a ripe yellow banana with stamps and address on it arrived by regular mail for a patient in a hospital in New Haven. On top of the figure is the original range image. Below is the initial model estimate (E) and the recovered model (N-30) after 30 iterations, without using any deformations. The fit of this model N-30 is quite poor—an indication that the object is an *irregular postal piece.* The model recovered using the built-in bending deformation achieves a better fit—shown in the 30th iteration (30) of the model recovery sequence when a total of 13 model parameters, including two bending parameters, were adjusted simultaneously. The recovered parameters were $a_1 = 13$ mm, $a_2 = 19$ mm, $a_3 = 87$ mm, $\epsilon_1 = 0.6$, $\epsilon_2 = 0.7$, *radius of the bend* = 85 mm.

Figure 6. Influence of coarser range maps on the recovered models. On the left, from top down, are the original range map (A) and coarser range maps, obtained by picking every 2nd (B), 4th (C) and 8th range points (D) in x and y axis of the original range map. On the right are the models recovered on the corresponding subsampled range maps but shown for comparison against the original range map A.

Figure 7. Residuum during recovery of models A, B, C, and D in Figure 6 as a function of CPU time on a VAX 785 computer. The jaggedness of the functions is due to the addition of Poisson noise after each iteration step which enabled escaping from shallow local minima. When the fitting function reached a plateau, the corresponding model did not improve any more. The dotted line (MR) shows the residuum for a multi-resolution fitting technique when model recovery started on the coarsest map and switched to a finer map when the fitting function did not improve any more.