Computer vision of free-form 3D objects by geometric matching

Heinz Hügli, Christian Schütz
Institute of Microtechnology, University of Neuchâtel
Rue A.-L. Breguet 2, CH–2000 Neuchâtel, Switzerland
E-mail: hugli@imt.unine.ch

The automatic recognition of free-form 3D objects is still a major problem. Recent advances in computer vision have brought two basic contributions that help to solve the problem: range imaging and geometric matching. It is believed that a recognition approach which builds upon these two contributions will solve a number of free-form 3D object recognition problems In this paper, we will show how range imaging and geometric matching can be combined in a promising recognition approach. We will also present examples of recognition of free-form 3D objects and an application in assembly.

KEYWORDS: 3D vision, free-form object recognition, pose estimation, range imaging, geometric matching

INTRODUCTION

In recent years, much attention has been given to the recognition of 3D objects. Although successful recognition has been obtained with simple objects, it is difficult to extend the traditional segmentation and primitive extraction approach to deal with free-form objects. It is not clear, how object parts should be defined and how reliable segmentation may work, as pointed out for instance in [2] Therefore, we opted for a recognition principle based on geometric matching. This approach works directly on the 3D coordinates of the object surfaces as obtained from range data. Hence, the recognition is independent on assumptions related to object features.

According to this geometric approach for object recognition, the comparison of test and model surfaces is performed with an iterative closest point matching algorithm (ICP) [1] In our framework, the test surface is represented by a range image whereas the model surface can be represented by a set of 3D points, polygons or other geometric primitives. For this reason, data obtained from object samples or from CAD models as well can be easily used as models.

Convergence is an important aspect of geometric matching by ICP. Since the closest point algorithm converges to a local minimum, successful matching is not guaranteed [6] and it is often an open question how well geometric matching can contribute to object recognition in practical situations. In order to assess the recognition performance of geometric matching by ICP, we investigated successful configurations for a set of test objects. We define the relative initial pose of test and model objects for which the ICP algorithm converges to the global minimum as successful initial configuration (SIC). We experimentally measure the range of SIC and derive the number of required initial configurations (RIC) to allow the recognition of the test object in any pose. The larger the range of SICs, the smaller the number of RICs and the smaller the computation cost [3].

Matching experiments performed by varying initial configurations lead to results which show the existence of a zone that guarantees successful recognition: it is referred as the range of SICs and can be used as a measure for recognition performance. We will present recent experiments performed on various 3D objects that give an estimate of the range and variation of the SICs. The knowledge of the range of the SICs is important for the design of the recognition system, especially when dealing with different objects.

Final recognition results obtained with the example objects show the conditions under which the closest point algorithm can be successfully applied to free-form 3D object recognition.

As an example of the application of free-form 3D object recognition by range imaging and geometric

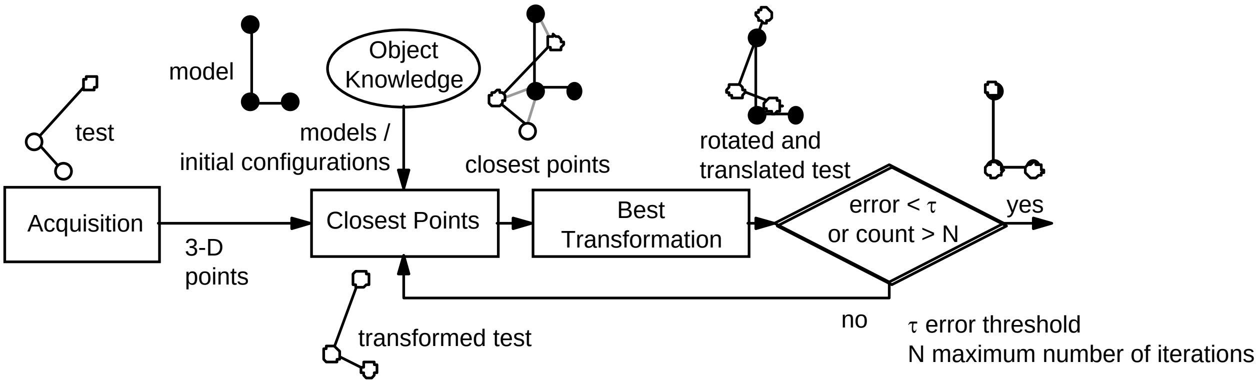

Fig. 1. Matching two objects with the iterative closest point algorithm (ICP)

Heinz Hügli, Christian Schütz, "Computer Vision of free-form 3D objects by geometric matching", Proc. Mecatronics, Europe-Asia congress on Mechatronics, Oct. 1-3 1996, Besançon

matching, we will show the automated assembly of scotch tape dispensers in a virtual reality robotics system [4] where 3D vision is used as a sensing system feeding back information from the real to the virtual world.

OBJECT RECOGNITION BY GEOMETRIC MATCHING

The aim of geometric matching is to compare two objects by measuring the similarity of the objects subject to rigid transformations. Geometric matching can be performed with the iterative closest point (ICP) algorithm [2]. Figure 1 gives an overview of the basic principle of ICP. Consider for the moment that objects are described as set of points in the space. One object is the test; the other is a model. The algorithm proceeds iteratively. First, it pairs every point of the test set with the closest point of the model set. These pairs of closest points between two objects to be matched are then used to calculate the rigid transformation (translation and rotation) which minimizes some distance measurement. The test object is then translated and rotated by the resulting transformation. This procedure is repeated until the predefined distance falls below a threshold τ or the number of iterations exceeds a chosen constant N.

The algorithm converges after some iterations to a solution which is characterized by its distance or error value. During the successive iterations, the test object undergoes successive rigid geometric transformations, bringing it progressively towards a better matching position. The algorithm is proven to converge. The error value decreases necessarily to a minimum and characterizes the solution found.

Successful matching occurs if the matching error, which is calculated as the sum of the mean and the deviation of the square coupling distances, is below a predefined threshold [11]. Otherwise, the algorithm is trapped in a local minimum.

Although convergence is guaranteed with the ICP algorithm, the found minimum is not always the global minimum: the algorithm can be trapped in a local minimum. A local solution can hide a best, optimal solution. Under these conditions, the problem is to find a way towards the best geometric match. Fortunately, this is not so much a problem practically. It was found that only global minima reach a very low value of the matching error. We therefore designate by successful matching one which error falls below a predefined threshold.

FROM RANGE IMAGING TO OBJECTS

Some sensors, such as laser scanners, yield range data. They measure the range or distance to the visible surface of the objects in the scene. Therefore, spatial location is determined for a great number of points on this surface. Formally, a range data set is a large collection of 3-D coordinate data, sampled at visible object surfaces in a scene. Rangel designates a range data element. Often sensors deliver range images which are range data sets with rangels ordered in a two-dimensional array. The array ordering is bound to the sensor configuration and has the same interpretation as for a conventional image.

The advantage of range images with respect to conventional intensity images (figure 2) derives directly from the fact that range measures the pure geometry of an object whereas the intensity image measures the amount of light reflected by it. While the geometry - the shape - of an object must be derived in a complicated way in the case of intensity images, this information in readily available in range images.

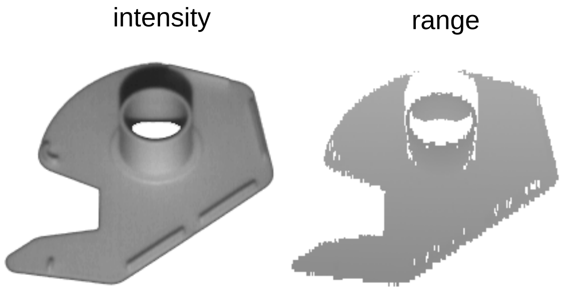

Fig. 2 Intensity and range images of a scotch tape dispenser part

For the purpose of recognition, a test object to be recognized is typically measured by a single range image and some kind of thresholding or segmentation can be applied to isolate it from its background, as for example described in [8]. The test data set is a subset of rangels from a single range image, it corresponds to a single view and represents only a part of the object surface.

A reference or model object must be represented by the entire object surface. Typically, it is obtained from CAD models or reconstructed from multiple range images of the real object. In any case, in this paper, we will consider them to be defined by an appropriate range data set of points distributed evenly on their complete surface.

The algorithm uses directly the sensed range data points and does not require any data pre-processing or local feature extraction, which makes it easily applicable to free-form objects and several types of object representations.

SIC-RANGE

Convergence is an important aspect of geometric matching by ICP. To illustrate it, we first present the results of a matching experiments in the 2-D space with puzzle pieces in order to better visualize the convergence behavior of the algorithm. The puzzle pieces are acquired with a video camera. After a thresholding of the image and the extraction of the contours we obtain a set of points describing the border of the puzzle piece. Such sets and subsets of puzzle pieces are fed into the closest point matching algorithm. Figure 3 shows some examples of the matching of puzzles and parts of puzzles. It shows initial and final configurations - relative positions and orientations - of test and model. It illustrates the successful matching of a complete piece; the successful matching of a subpart; a suboptimal match of a subpart. In all cases, the test pieces converge towards the model in few iterations. However, the selection of the initial configuration becomes crucial to a successful matching.

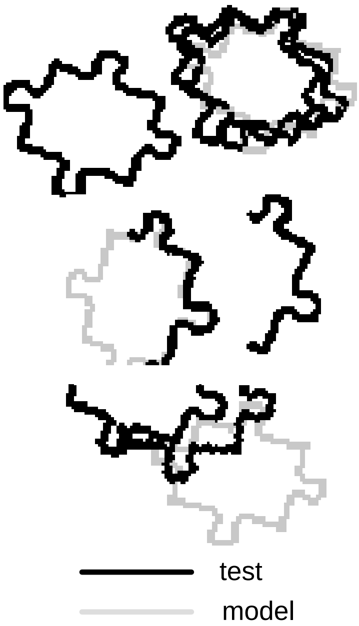

Fig. 3: Matching iterations of puzzle pieces

Knowing that the choice of this initial configuration has great effect on the success of the matching, we present and analyze here this influence quantitatively.

Considering the space of possible initial configurations, successful matching is obtained only for a limited range of it. We name it successful initial configuration range or SIC-range.

SIC-RANGE FOR 2D-SHAPES

To measure it, we place successively the test piece in several different positions around the model puzzle (x,y) and rotate it around its center of mass φ. Figure 34 plots the SIC-range for all rotations and sixteen positions arranged in a grid. The black sectors at every grid position indicate the angles φ for which the matching was successful.

We observe that the translation between two pieces to be matched is of minor influence to the convergence. However, if the rotation angle between the test and the model piece exceeds a certain value the registration fails. The SIC-range for the rotation angle φ is about [+30..- 30] degrees.

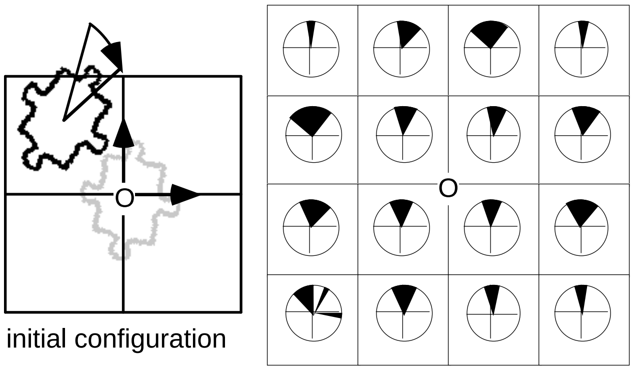

Fig. 4: SIC-range (in black) in (x,y,φ) space

SIC-RANGE FOR 3D-SHAPES

To practically measure the SIC-range in 3D space, we first reduce the dimension of the 6D space of initial configurations to the 3D space proposed in [12] and corresponding to the setup in figure 5. The test object is placed on the view axis defined by the camera pointing towards the model object at a fixed distance between model and camera. The space of initial configurations can then be defined by the triple (ϕ, θ, ω) where ϕ and θ are respectively zenith and azimuth angles of the view axis in the model spherical reference system and ω designates the camera rotation angle around this axis. From now on, we name the pair (ϕ, θ) view point.

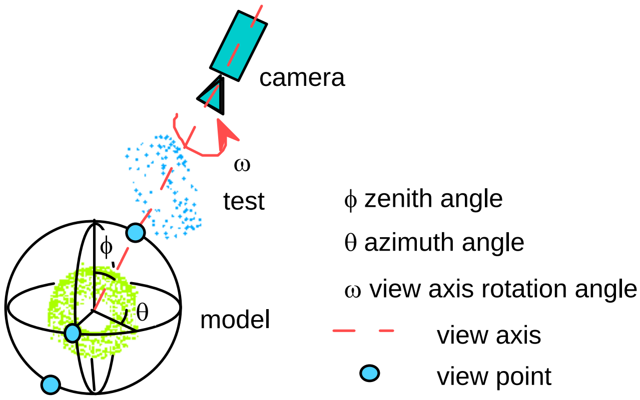

Fig. 5: SIC-range (in black) in (x,y,φ) space

We thus measure the SIC-range in the (ϕ, θ, ω) space. In order to do so, the (ϕ, θ, ω) parameter space of all possible initial configurations has to be inspected for successful matching. Successful configurations form the SIC-range. We estimate the SIC-range using the three free-form objects shown in figure 6. They represent injected plastic toys and are called duck, fish and swan.

The matching results are presented for every view point as a small circle where the black segment represents the SIC-range for ω as shown in figure 7. The view points on the sphere are projected on a plane tangential to the pole for visualization in 2D giving rise to the SIC-map as defined in figure 7. View points having the same zenith angle lie on a circle around the pole. Note that view points from the lower hemisphere are omitted since there is no successful matching at all.

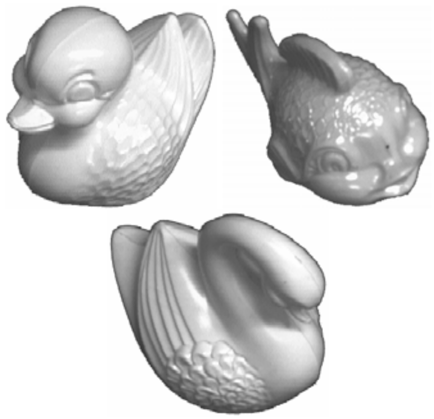

Fig. 6. Free-form objects: duck, fish and swan

Figure 8 shows the SIC-maps of the three toy objects obtained experimentally according to the above explained method. We observe that the SIC-range for ω decreases for view points with growing zenith angle. While this decrease is azimuth dependent for the fish it is nearly azimuth independent for the duck and swan.

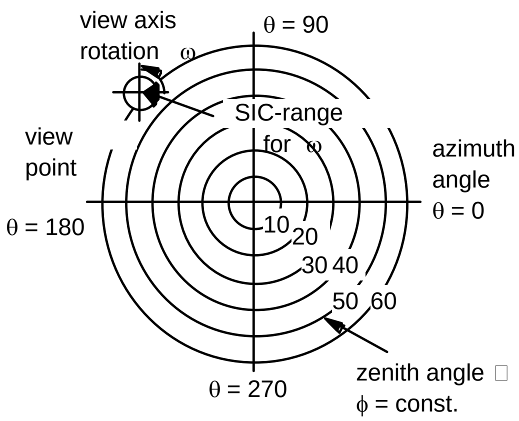

Fig. 7. The SIC-map

In order to give some general statement, let us define a SIC-range by limited zenith angle and rotation angle values as

$$R(\phi_0,\omega_0)=\{ \phi<\phi_0, |\omega|< \omega_0 \}$$

From the data of figure 7 follows the validity of following conservative and quantitative statement

$$R(30,50) \wp R_{duck}, R_{fish}, R_{swan}$$

giving a SIC-range which is completely included in all three measured SIC-ranges. The measured objects exhibit a SIC-range which is larger than 30 degrees in zenith angle and larger than [+50..-50] degrees in rotation.

The results obtained show a rather large SIC-range for the objects under evaluation. This fact is a very positive point for the practical application of geometric matching for 3D object recognition because it confers the method good stability. Moreover, it helps to keep computational costs low for every application where no a priori knowledge is available about the configuration of the object to be recognized and an exhaustive search comprising all possible initial configurations must be performed.

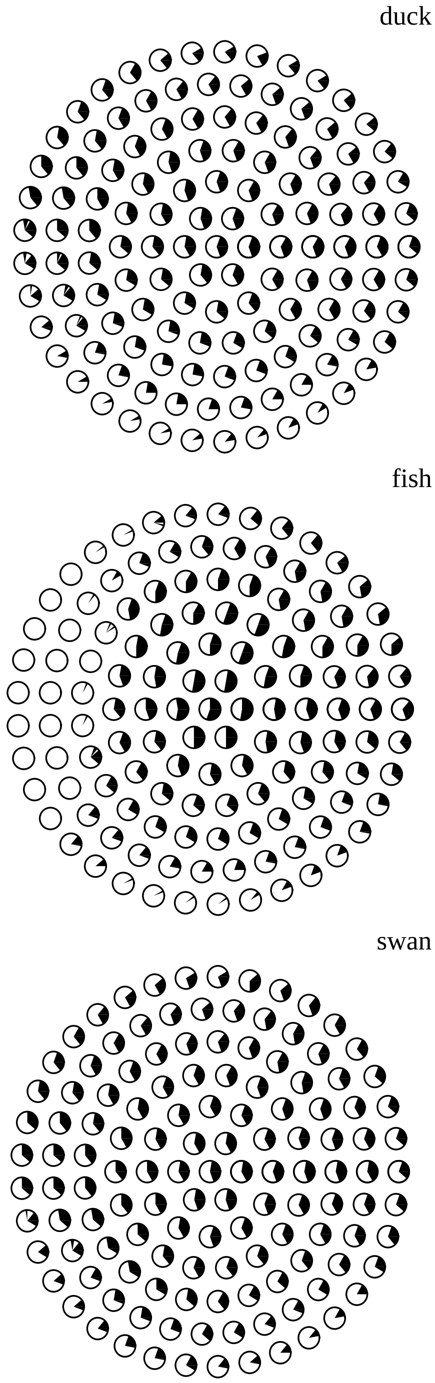

Fig. 8. SIC-map for three free-form objects (scale as in figure 7)

RECOGNITION APPLICATION

We consider the automatic assembly of scotch tape dispensers as a test for 3D vision by geometric matching. The task is to assemble the roll and the two side parts into a complete tape dispenser. As the three parts are available at the same location and have to be taken by the same robot, the task of vision is to identify the parts and to determine their pose.

Vision is by range imaging, with BIRIS (Vitana Corp.) or alternatively with ABW 320 projector. The robot is a 5 axis robot arm Mitsubishi Movemaster RV-M1 equipped with a special gripper [4].

A first vision approach is based on a classical recognition paradigm and proceed by range image segmentation into surface patches, patch grouping, object hypothesizing and verification [7]. The system is not in a position to recognize the 3D-shapes.

Using the geometrical recognition approach presented above, identification and pose estimation are performed with satisfaction. Furthermore, as no hypothesis is made about the object shape, objects of any shape can be considered. The model is built up very easily, from CAD or a simple example.

CONCLUSIONS

In this paper, we presented some aspects of a recognition approach which builds upon range imaging and geometric matching to solve a number of free-form 3D object recognition problems.

Regarding the convergence properties of the geometric matching algorithm, we introduced the notion of range of successful initial configurations (SIC) and an adequate graphical representation of it known as SIC-maps. These maps provide a simple view of the 3-dimensional SIC-range and represent a useful tool for the analysis of the quantitative convergence properties of physical objects. SIC-maps were measured for three real objects. The observed SIC-ranges appear to be rather large which indicates good convergence properties towards a successful matching, i.e. successful recognition. Also, it helps to keep computational cost low, providing good perspective for a universal application of this matching method.

Regarding the application of this recognition approach, we showed the automated assembly of scotch tape dispensers in a virtual reality robotics system where 3D vision is used as a sensing system feeding back information from the real to the virtual world. Whereas classical recognition approaches are sufficient for recognizing simple shaped objects, like polyhedrons, they were not in a position to recognize the more complex shaped scotch tape dispenser parts. The presented approach based on geometric matching was however in a position to recognize it in a simple way. In fact, no assumption was made on the objects geometry, the object was recognized as an arbitrary shaped object. This absence of hypothesis on the object shape, this capability to recognize objects of arbitrary shape makes this geometric matching approach a very general, universal and promising recognition tool.

ACKNOWLEDGMENTS

This research was partly funded by the Swiss national Priority Program in Informatics, Knowledge-based systems, under project number 5003-344336 and by Swiss national science foundation under project number 2100-43530.

REFERENCES

[1] Besl P. and McKay N., "A method for registration of 3-D shapes", In: IEEE Transactions On Pattern Analysis and Machine Intelligence, Vol. 14, No. 2, pp. 239-256, 1992.

[2] Hebert M., Ponce J., Boult T.E., Gross A. and Forsyth D., "Representation for computer vision", Report of the NSF/ARPA workshop on 3D object, CUNY Graduate School and University Center, 1994.

[3] Hügli H., Schütz Ch., Semitekos D., "Geometric matching for free-form 3D object recognition", In: ACCV, Singapore, Vol 3, pp.819-823, 1995

[4] Natonek E., Zimmermann Th., Flückiger L.," Model based vision as feedback for virtual reality robotics environments", Proc. IEEE Virtual Reality Annual Symposium, North Carolina, 1995

[5] Schütz Ch., Hügli H., "Free-form 3D object recognition", Optical 3-D Measurement Techniques III, Heidelberg, Wichmann, pp. 516-525, 1995.

[6] Schütz Ch., Hügli H., "Towards the recognition of 3D free-form objects", In: Intelligent Robots and Computer Vision XIV, Algorithms, Techniques, Active Vision and Materials Handling, SPIE, Philadelphia, Vol. 2588, pp. 476-484, 1995.

[7]
H. Hügli & P. Gingins, "Vision by range and intensity for model-based object recognition", Proc. 2nd Japan-France Congress on Mechatronics, Takamatsu, Japan, 1-3 Nov. 1994

[8] Christian Schütz and Heinz Hügli "Change detection in range imaging for 3D scene segmentation", International Symposium on Lasers, Optics and Vision for Productivity in Manufacturing I, Besançon, 10-14 June 1996, SPIE, Vol 2786, 1996