

RECOGNITION OF 3-D OBJECTS WITH A CLOSEST POINT MATCHING ALGORITHM

Ch. Schütz and H. Hügli, University of Neuchâtel

Tivoli 28, 2000 Neuchâtel, Switzerland
Phone: ++41 38 30 16 53
Fax: ++41 38 30 18 45
Email: schutz@imt.unine.ch

ABSTRACT:

This paper is a contribution to the recognition of arbitrary 3-D shapes from range images. Recently, a closest point matching algorithm was proposed that is in a position to match 3-D shapes without the need of an explicit object description in terms of primitives. The advantage is that the comparison works directly on the 3-D coordinates of the surface points and needs neither segmentation nor feature extraction. Arbitrary shapes can potentially be matched with it.

In this paper we present a number of 3-D object recognition experiments conducted for assessing the capability of this approach to recognise objects measured by range images. The experiments performed with the designed methods show the fast convergence and reliable recognition of objects. We also address some of the key issues of this matching algorithm, which are the strong dependence on the initial conditions, like position of the object to be recognised, and the quadratic computing complexity.

KEY WORDS: 3-D free form object recognition, pose estimation, range imaging.

1. INTRODUCTION

The work described in this paper was carried out in the context of knowledge based 3-D object recognition. We developed a hybrid recognition system, which combines range and intensity images to generate and verify object hypotheses (Hügli, 1994b). The range images are acquired with a range finder working on the principle of space coding with projected stripe pattern and triangulation. One application of our vision system is to recognise known objects in a robot environment, which permits to update a virtual representation of the robots workspace (Natonek, 1995).

In our system the model and test objects are represented by a set of 3-D points, called a dense 3-D map. There exist mainly two approaches for geometrically matching dense 3-D maps (Zhang, 1994). The first one is based on the description of the dense 3-D map in terms of geometric primitives. Primitives of the model and the test are organised in a graph structure and mapped one onto the other with a graph matching method. The second approach considers the dense 3-D map as a surface and tries to find a transformation between the model and the test. The primitive-based approach exploits symbolic matching criteria and allows important changes in orientation and positioning between model and test, usually essential in object recognition. Unfortunately this approach depends largely on a robust and precise detection of primitives which remains a problem with current known methods (Zhang, 1994) (Maître, 1994). In contrast the surface-based approach uses all available information and does not need data segmentation. It can be easily applied to free form objects. The techniques for this approach usually need a priori knowledge of the transformation between test and model to perform a successful recognition. This approach has been used in visual navigation (Zhang, 1994) and object modelling (Chen, 1992) where motions between two surfaces to match are small while the primitive-based approach has been widely applied in object recognition.

Recently, an interesting technique for the registration of dense 3-D maps has been proposed (Besl, 1992). It uses a surface-based approach and needs, as one would expect, a priori knowledge to give a rough estimate of the transformation between the test and model object. Closest points between the two surfaces to match are coupled and used to calculate the transformation, which minimises the mean square error of the distances. The test object is then moved and rotated by the resulting transformation. This procedure is applied several times until the error falls below a threshold or the number of iteration exceeds a predefined constant. Other researchers have used similar algorithms (Zhang, 1994) (Feldmar, 1994) (Chen, 1992) to estimate motions, register surfaces and build models. We show some experiments to investigate the usefulness of the closest point matching algorithm for the recognition of objects measured by range images.

2. CLOSEST POINT MATCHING ALGORITHM

The algorithm, we use, was proposed in (Besl, 1992). It aims to find the transformation consisting of rotation \mathbf{R} and translation \mathbf{t} between two dense 3-D maps, referred as a model X and a test P . X and P are represented by their coordinates in space. Here follows a short description of the algorithm.

• input: Two dense 3-D maps X and P containing respectively N_X and N_P 3-D points.

• output: Transformation (\mathbf{R}, \mathbf{t}) between X and P found by the algorithm.

• iteration:

(1) Build the set Q_j of closest point pairs with $\mathbf{q} = (\mathbf{p}, \mathbf{x})$: $\forall \mathbf{p} \in P_j$ find $\mathbf{x} \in X$ with $\|\mathbf{p} - \mathbf{x}\| = \min(\|\mathbf{p} - \mathbf{x}_j\|)$, $j \in [1, \dots, N_X]$.

(2) Find the transformation $(\mathbf{R}_j, \mathbf{t}_j)$ that minimises the mean square error $e_j(\mathbf{R}_j, \mathbf{t}_j) = 1/N_p \cdot \sum \|\mathbf{R}_j \mathbf{p}_k + \mathbf{t}_j - \mathbf{x}_k\|^2$ with $(\mathbf{p}_k, \mathbf{x}_k) = \mathbf{q}_k$, using the quaternion method (Faugeras, 1986).

(3) Apply the transformation, $P_{j+1} = (\mathbf{R}_j, \mathbf{t}_j)(P_j)$.

until e_j is below a threshold τ or j is greater than a limit μ .

The simplicity of this algorithm makes it easy to implement. Besides, it does not require any data pre-processing or local feature extraction. This facts allow to handle reasonably noise effects that occur in our range finder system. The algorithm converges after a few iterations as showed in (Besl, 1992). Most of the computing time of the algorithm is spent in the closest point search part (1) whose cost is $O(N_p \cdot N_x)$, when not using optimised searching techniques. The calculation (2) and application (3) of the transformation grow both linearly with N_p .

We will now show some experiments to demonstrate the usefulness of the closest point matching algorithm for free form object recognition. First, there will be a summary of experiments in the plane with puzzle pieces. Then, we will do a systematic investigation of the algorithm in 3-D space using a synthetic object.

3. EXPERIMENTAL RESULTS IN 2-D SPACE

We first present experiments in the 2-D space with puzzle pieces in order to better visualise the behaviour of the algorithm in two dimensions. The puzzle pieces are acquired with a b/w camera. After a thresholding of the image and the extraction of the contour 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 1 shows some examples of the matching of puzzles and parts of puzzles placed at different positions and different orientations. The example illustrates the correct matching of a complete piece and of a subpart as well as an incorrect match of a subpart. We observe that the test pieces converge in few iterations towards the model. However, the selection of the starting position becomes crucial to a successful registration.




Fig. 1: Matching iterations of puzzle pieces

To get a global view of the influence of translation and rotation we place successively the test piece in several different positions around the original puzzle. Moreover the so placed test piece will be rotated around its center of mass. Figure 2 plots the results for all rotations and sixteen positions arranged in a grid. The black sectors at every grid position indicate the angles for which the matching has been successful.




Fig. 2: Convergence zone for a puzzle piece placed and rotated differently

We observe that the translation between two pieces to be superposed 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 convergence zone for the rotation angle is about between ± 30 degrees.

Since the computing time of the closest point matching algorithm grows with $O(N_x \cdot N_p)$, we are interested in using as few points as possible. Experiments with subsampled data confirmed that good convergence is available for reduction factors up to eight, which allows to represent a puzzle border by only 60 points (Hügli, 1994a).

In object recognition it is important to know if the final distance between two different objects is significantly larger than between two representations of the same object. This fact must be satisfied to allow reliable decision. Experiments with puzzle pieces of different shapes showed that a threshold can be set for the final distance to separate reliably the object classes (Hügli, 1994a).

4. EXPERIMENTAL RESULTS IN 3-D SPACE

Encouraged by the promising first results, which are fast convergence, easy application to free form objects and relatively large zone of convergence, we now consider objects in 3-D space. We use a synthetic object to easily control all parameters like object size, point density and symmetry. A corner represented by a set of points forming its three sides is our choice. We now test how a copy of the corner object placed and oriented differently in space can be matched with its original. The number of iterations is fixed at ten, if not indicated otherwise. Figure 3 shows the model together with the test to be positioned at all 27 locations of the grid. Each position has a number which is referenced in the result plots. The error that corresponds to the mean square distance between test and model points will be indicated in square units.

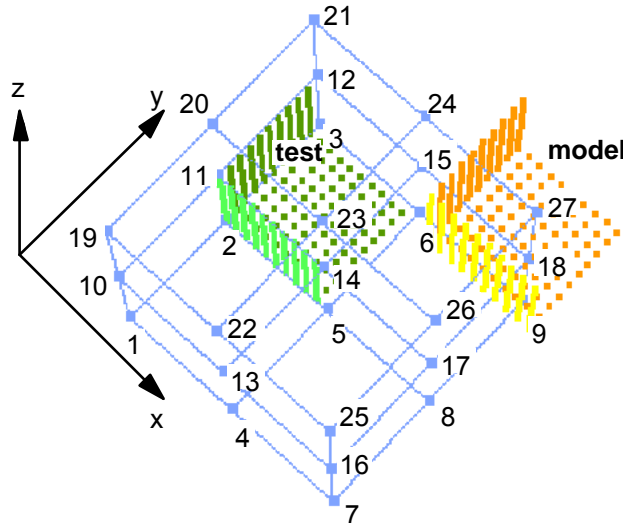


Fig. 3: Distribution of the starting positions

Translation. The influence of translation and rotation on the matching shall be observed separately. In a first experiment we keep the orientation unchanged and investigate how the convergence depends on the starting position. Figure 4 shows the final error of the matching for every starting point. A threshold, which is set at an error of one square unit, separates good and bad recognition cases (dashed line in figure 4). The black squares in figure 4 indicate the positions for which the matching failed. We see that the closest point matching algorithm can match the duplicate and the original for most of the starting positions. If we further investigate the positions for which the matching failed, we notice that the algorithm introduced a bad rotation at the first iteration step. This fact relies on the instable situation when many points in the test are coupled with only a few points in the model.




Fig. 4: Results for translated corners

Convergence. An iterative algorithm is only useful if it converges quickly. This criterion will be checked by observing the development of the error during the iteration. Figure 5 shows the measured matching error at successive iteration steps. Again the test objects are only translated. The errors are thresholded by appropriate limits at the iteration steps two, five and ten. At each step the result of the decision is the same and showed in the right half of figure 5. We observe that already after few iterations the successful cases with a minimum error can be extracted. This helps one to decide at an early stage of the iteration if a matching will be successful or not. Computing time can be saved because one does not really need an exact matching to find the interesting cases at the beginning of the object recognition procedure.

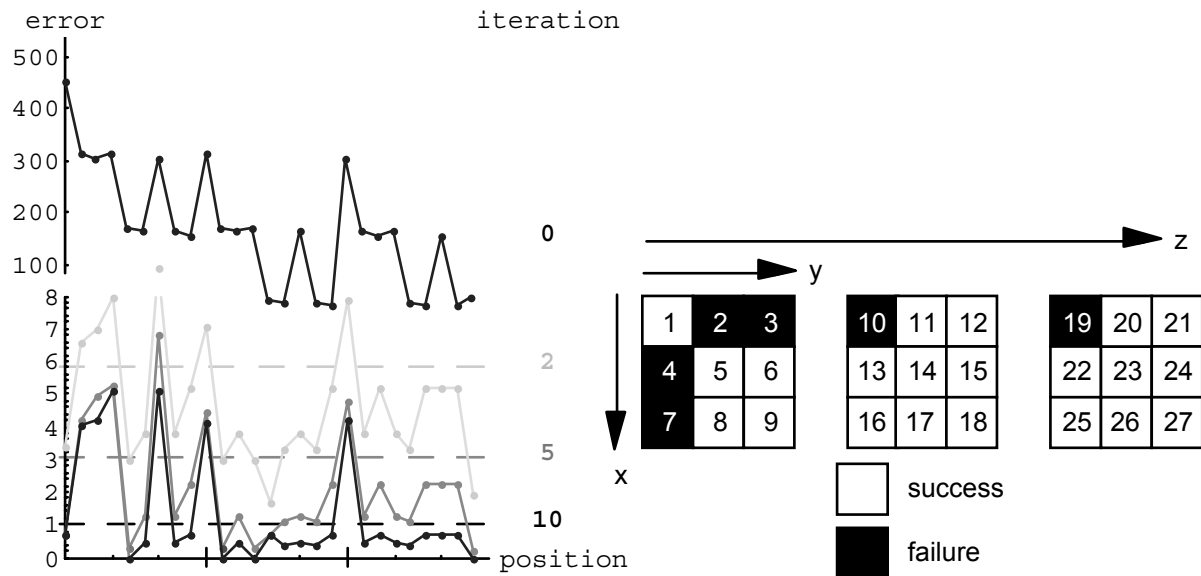


Fig. 5: Convergence of the closest point matching algorithm

Subsampling. Searching the closest point is the most time consuming part of the algorithm and grows with $O(N_p \cdot N_x)$. For a better performance we are thus interested in reducing the number of points representing the objects. In this experiment the number of points building the corners faces are subsampled uniformly to see if the matching still succeeds. As before we allow only translations. Figure 6 shows the results where, first, the number of points of the test has been reduced by two and, secondly, also the number of points of the model.

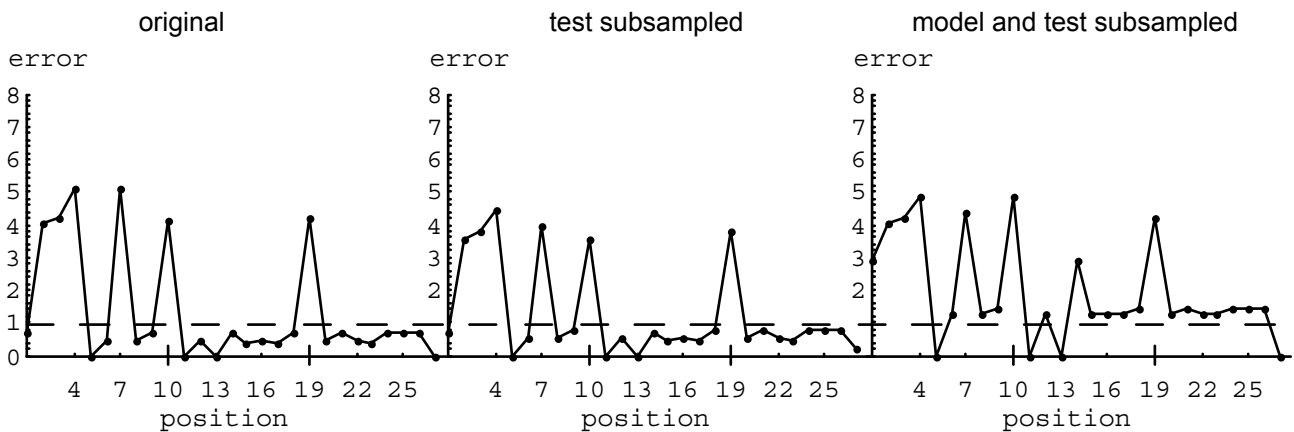


Fig. 6: Matching with objects subsampled in space

A reduction by a factor of two of the test has no influence on the result. If both the test and the model are subsampled in space the performance degrades. This fact will allow us to work with a test at a relatively poor resolution, which again will enhance speed in the recognition process.

Geometry. The next experiments investigate the influence of object geometry on the matching. We repeat the first experiment with an asymmetric corner. The corner faces have edge lengths with a ratio $x:y:z$ of 10:7:5. The left drawing in figure 7 shows the positions for which the matching is successful. These positions differ from the symmetric corner (figure 4) and there are slightly more failures.

The results of the matching of puzzle subparts showed that convergence depends largely on the starting position. To see if the same problem occurs in space we compare a small corner to a larger one. The subpart matching works only for a few initial conditions (see right drawing in figure 7). These cases have an interesting characteristic. All successful starting points are located near to the symmetric axis of the model corner.

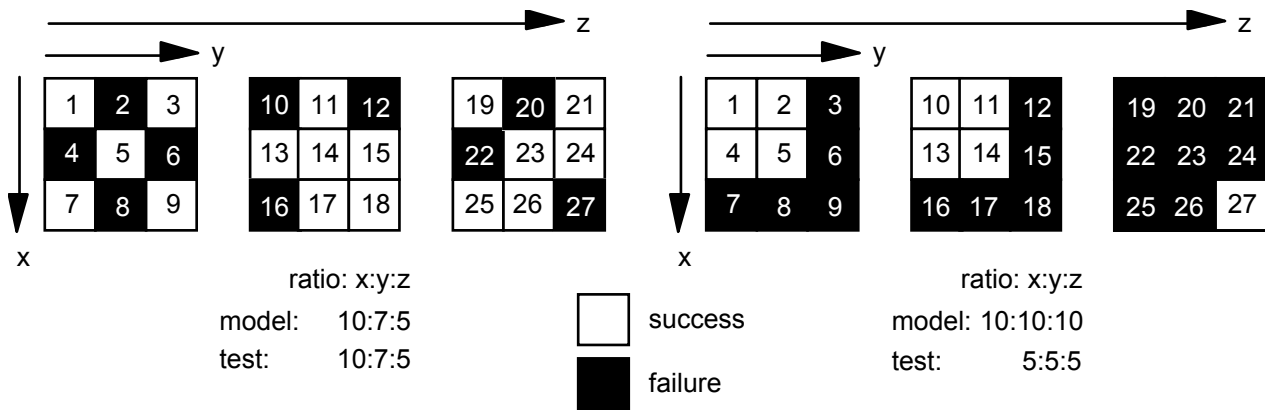


Fig. 7: Results for a symmetric corner and a subpart of a corner

Rotation. We analyse now the effect of rotation. The test will thus be rotated around the z-axis. Because we know the matching to be influenced by principal moments (Besl, 1992), the same experiment is also performed with an asymmetric corner. The final error for a symmetric and an asymmetric corner object is plotted in figure 8.




Fig. 8: Matching results for rotated corners

The plot reveals a zone of convergence (final error below one) that is about ± 40 degrees for the symmetric corner and slightly larger for the asymmetric one. Comparing these results with the ones obtained for the translated corners we see that the rotation angle is the major constraint for the closest point matching algorithm.

5. CONCLUSIONS

We presented a number of recognition experiments conducted in order to assess the capability of the closest point matching algorithm to recognise objects measured by range images.

The presented results show that the translation between test and model is of minor influence to the success. On the contrary, the rotation limits the convergence to a zone of about 80 degrees. Asymmetric objects show a slightly better performance, whereas subparts can only be matched for some starting points.

The limited performance of subpart matching will be important for designing recognition system, where the closest point matching algorithm is used to recognise an object measured from a single point of view. To overcome the large dependence of the convergence on the starting points, one has to place the test object at different positions and orientations. The distance between these heuristic positions should be smaller than the convergence zone of the algorithm.

At a first glance the necessity of multiple starting points seems to introduce much overhead. But since the convergence zone is relatively large, the number of starting points is quite low. Furthermore the algorithm converges quickly and allows subsampling of the point set representing the objects, which significantly reduces the computing time. These facts lead to a feasible method for the recognition of arbitrary shapes.

6. ACKNOWLEDGEMENT

This research has been funded by the Swiss national Priority Program in Informatics, Knowledge-based systems, under project number 5003-344336.

7. REFERENCES

- Besl, P. and McKay, N., 1992. A method for registration of 3-D shapes. In: IEEE Transactions On Pattern Analysis and Machine Intelligence, Vol. 14, No. 2, pp. 239-256.
- Chen, Y. and Medioni, G., 1992. Object modelling by registration of multiple range images. In: Image and Vision Computing, Vol. 10, No. 3, pp. 145-155.
- Faugeras, O. and Herbert, M., 1986. The representation, recognition, and location of 3-D objects. In: Int. J. Robotic Res., Vol. 5, No. 3, pp. 27-52.
- Feldmar, J. and Ayache, N., 1994. Rigid and affine registration of smooth surfaces using differential properties. In: Lecture Notes in Computer Science, Vol. 801, No. 2, pp. 397-406.
- Gingins, P. and Hügli, H., 1994. Model-based 3D object recognition by a hybrid hypothesis generation and verification approach. In: Intelligent Robots and Computer Vision XIII: Algorithms and Computer Vision, Boston, SPIE 2353, pp. 67-77.
- Hügli, H., Gingins, Ph., 1994a. Expériences de reconnaissance avec l'algorithme ICP. Rapport IMT 369 HU 09/94, Université de Neuchâtel.
- Hügli, H., Schütz, Ch., Baur, Ch., Natonek, E., 1994b. Knowledge-based 3-D vision system. In: Proceedings of Information Conference Module 2, Knowledge based systems, Swiss National Science Fundation, pp. 46-53.
- Maître, G., 1994. Segmentation et traitement préliminaire des images de profondeur. Thèse de doctorat à l'Université de Neuchâtel.
- Natonek, E., Zimmermann, Th., Flückiger, L., 1995. Model based vision as feedback for virtual reality robotics environments. In: IEEE Virtual Reality Annual Symposium, North Carolina.
- Zhang, Z., 1994. Iterative point matching for registration of free-form curves and surfaces. In: International Journal of Computer Vision, Vol. 13, No.2, pp. 119-152.