Geometric Matching of 3D Objects:
Assessing the Range of Successful Initial Configurations

Heinz Hügli, Christian Schütz
University of Neuchâtel, Institute for Microtechnology
Rue Breguet 2, CH-2000 Neuchâtel, Switzerland
hugli@imt.unine.ch, schutz@imt.unine.ch

Abstract

This paper considers the matching of 3D objects by a geometric approach based on the iterative closest point algorithm (ICP), which, starting from an initial configuration of two rigid objects, iteratively finds their best correspondence. The algorithm does not converge always to the best solution. It can be trapped in a local minimum and miss the optimum matching. While the convergence of this algorithm towards the global minimum is known to depend largely on the initial configuration of test and model objects, this paper investigates the quantitative nature of this dependence. Considering the space C of relative configurations of the two objects to be compared, we call range of successful initial configurations, or SIC-range, the subspace of C which configurations bring the algorithm to converge to the global minimum. In this paper, we present a frame for analyzing the SIC-range of 3D objects and present a number of original experimental results assessing the SIC-range of a number of real 3D objects.

1. Introduction

Geometric matching represents a powerful method for the recognition of 3D objects. It applies to objects represented by their geometry and is typically used to compare the shape similarity of two objects subject to a rigid transformation.

When correspondences between the elements of two objects are known, the matching is solved classically by finding the optimal rigid transformation [12]. When correspondences between the two objects are unknown, the problem cannot be solved in closed form but a solution is provided by the iterative closest point (ICP) algorithm [2].

The ICP algorithm starts from an initial configuration of two rigid objects and iteratively finds their best correspondence. This algorithm is proven to converge but does not converge always to the best solution. It can be trapped in a local minimum and miss the global minimum. While the convergence of this algorithm towards the global minimum is known to depend largely on the initial configuration of test and model objects, few is known about the quantitative nature of this dependence.

Every well designed recognition system uses a priori knowledge to constraint the recognition search. The more the knowledge, the less complex the search. In previous works using geometric matching methods, initial configurations are derived from contextual knowledge about the range finder setup [8][3][4] or found automatically by extracting some features as curvature or edges [1][9][10][5]. In absence of a priori knowledge or robust features, the ICP algorithm is started with multiple different initial configurations. In a brute force approach, the ICP algorithm is even started with initial configurations covering the whole configuration space [6].

In any case, the recognition strategy and performance depend on the nature of the convergence towards the global minimum. The lower the constraints on the exact choice of an initial configuration, the less the number of initial configurations required to solve the problem. In other words, the larger the convergence range, the less complex the search. Knowing quantitatively the convergence range of objects is therefore a key towards the design of efficient recognizers.

This paper presents work carried out for assessing the quantitative convergence range for the geometric matching of 3D objects. After a definition of the convergence properties, we present a frame for analyzing the matching convergence range of 3D objects. We present a number of original experimental results assessing the matching convergence properties of several real 3D objects acquired by range imaging.

2. Convergence of ICP

This section discusses the convergence properties of ICP and introduces the notion of SIC-range. An example illustrates SIC-range for 2D objects

Geometric matching by ICP

The ICP algorithm used for geometric matching proceeds iteratively. In the case of objects represented by sets of points, it first pairs every point of the test set with the closest point of the model set. All pairs of closest points between two objects to be matched are then used to calculate the rigid transformation that minimizes some distance measure ε_k. The test object is then translated and rotated by the resulting transformation. This procedure is applied repeatedly until the distance ε_k falls below a threshold ε_0 or the number of iterations exceeds a chosen value k_{max}.

Fig. 1: Matching puzzle parts by ICP

An example of the successive steps of matching two puzzles is shown in figure 1. We see how the black test object is successively moved towards the gray model object as the algorithm proceeds.

Successful initial configuration

Not all initial configurations of test and model objects lead to optimal matching. As an example, figure 2 illustrates the configurations of a test and a model puzzle before and after matching. Note that in this example the test is only a subpart of the model. Optimal or successful matching is shown in a) while unsuccessful matching is illustrated in b). In this later case, the algorithm keeps locked in a local minimum and misses the global minimum that stands for successful matching.


Fig. 2: Successful (a) and unsuccessful (b) matching

SIC-range

Let us call configuration c, the relative pose of test and model and let C denote the space of possible initial configurations. Successful matching is obtained only from a limited range of C. We name it range of successful initial configurations or SIC-range.

The SIC-range being a subset of the space C of initial configurations, we can represent it by a binary function $f(c)$, $c \in C$, which is true iff c is leads to successful matching.

Recalling that the SIC-range refers to two objects, we can write $f_{tm}(c)$ to refer to test t and model m. Note that $f_{tm}(c)$ is in general different from $f_{mt}(c)$, because the ICP algorithm, and more specifically the nearest neighbor computation used in it, is not commutative.

In the special case where an object is matched with itself, we write the related function $f_{mm}(c)$. Only in this special case the SIC-range refers to a single object.

Assessing the SIC-range of two objects means exploring C and testing each initial configuration c with ICP for successful matching, setting $f(c)$ accordingly. With limited computing resources, only a limited number of configurations can be tested: Therefore, the space of initial

configurations C is explored at discrete locations c_i leading to the discrete function $f(c_i)$.

The remainder of this paper is devoted to the assessment of the SIC-range of various real objects.

All objects considered from now on, both tests and models, are represented by sets of points. Therefore also, the ICP algorithm being used is the one relative to points, as described previously.

SIC-range of 2D objects

For 2D objects, the configuration space is three dimensional. A configuration c_i is defined for example by the coordinates (x_i, y_i, ω_i) referring to the translation $(x_i, y_i,)$ and relative orientation ω_i of the two objects.


Fig. 3: Configuration (x,y,ω) and associated SIC-range (black sectors)

To assess the SIC-range of the puzzle of figure 3, we consider sixteen translations of the test in a regular grid around the model and, for each translation, 36 evenly spaced relative rotations ω. Test and model represent the same object, i.e. $f(c)=f_{mm}(x,y,\omega)$. Figure 3 plots the obtained SIC-range. The center of each circle indicates the translation. The black sectors at every grid position indicate the angles ω for which the matching is successful.

We observe that the SIC-range is limited in the rotation range ω and that translation between two puzzles to be matched is of minor influence. Very roughly, and for small initial translations, the SIC-range of this puzzle can be described by $\omega \in [+30..-30]$ degrees.

3. SIC-range of 3D objects

This section extends the SIC-range to 3D objects and presents a useful representation of it as a SIC-map.

6-dimensional SIC-range

With 3D objects, the configuration space C is 6-dimensional and so is the SIC-range. The analysis of this general case is not simple. On one hand, limited computational resources impose a limited number of configurations to be tested, and do not allow a dense cover of C. On the other hand, a representation and a fortiori an interpretation of a 6-dimensional SIC-range is not simple, anyhow. Therefore, we search for a meaningful SIC-range representation of lower order.

Preferred configurations of 3D-objects

Let us consider object matching in a frame where following general hypotheses hold:
* Object are represented by sets of points
* The model is compact
* The test includes only points from a visible part of the object


φ zenith angle

θ azimuth angle

ω view axis rotation angle

— — view axis

◎ view point

Fig. 4. Preferred configuration setup of test and model

In this frame, we define a class of preferred initial configurations which we believe are well suited for successful convergence.

The related setup [7] is shown in figure 4 and defined as follows. The coordinate system is bound to the model and placed at its center of mass. The test object is placed on the view axis defined by the camera pointing towards the model center at a fixed distance from it. 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.

Motivations for the selected setup are as follows:
* The test is initially optimally placed near the center of the model.
* Placing the test between model and camera ensures the test is moved from outside towards the model and discourages any match with its invisible parts.

SIC-map assessment

Starting from 3D model and test objects, we assess their SIC-range by exploring the (ϕ,θ,ω) space at discrete locations, testing ICP for successful matching and setting $f(\phi,\theta,\omega)$ accordingly.

Being 3-dimensional, the SIC-range can now be visualized by a map similar to the one used for 2D objects. This SIC-map [11] is defined in figure 5 and examples are shown in figure 8. The small circles span the (ϕ,θ) space of the spherical coordinates. Each circle holds for the (ϕ,θ) position of the view axis and represents by its black sector the associated SIC-subrange for ω. The spherical coordinate system is projected on a plane tangent to the pole. View points with same zenith angle lie on the same circle around the pole and those with same azimuth lie on

the same radius. View points from the lower hemisphere are omitted.

The SIC-map assessment includes several steps. Start is from the pole configuration, which is the pose obtained after a first successful match of the model and test.

From this pole configuration, the complete (ϕ,θ,ω) parameter space is explored in discrete steps varying azimuth angle θ from 0 to 360, zenith angle ϕ from 0 to 180 and view axis rotation angle ω from 0 to 360 degrees. This procedure orientates the test at different view points and then rotates it around the view axis.


Fig. 5. The SIC-map

The exact sequence brings the pole axis to the spherical coordinates (ϕ,θ) by a rotation of the test object by zenith angle ϕ around the rotation axis which is perpendicular to the great circle of azimuth angle θ as illustrated in Fig. 6. Then, follows the rotation ω around the view axis.


Fig. 6. Test transformation from pole to view point

For every initial configuration the ICP matching algorithm is launched, running enough (40) iterations to ensure convergence. Note that the criteria for successful convergence can be chosen differently. It could be a comparison of pose, it can also be a decision on the final error. We selected this last criteria for the decision.

4. SIC-maps of real objects

The experimental range data are from a light stripe range scanner. Light stripe sequences generated by a LCD-modulated light source are projected onto the 3D objects and recorded with an image camera. A 3D reconstruction process applying the triangulation principle delivers a single range image as well as the associated intensity

103

image. A separation of object and background is performed directly at this level by ignoring the black background of the scene. The acquisition step provides therefore directly the range data of the object seen under the corresponding viewing angle.

The test object, or simply test, is derived from a single range image obtained by above described scanner. The range image is transformed to a set of points in R^3 whose size is reduced by a decimation procedure [7] to a few hundreds points. A test is finally a set of few hundred points in R^3 which stem from the visible surface of an object seen by a camera.

The model object is reconstructed from several range images of the same object. First, range images are taken of the same real object observed under various viewing directions. Practically however, the range scanner is fixed and the object is rotated. Then, the set of range images are merged with a fusion method, giving rise to the 3D-model, represented by a set of few hundred points in R^3.

Example sizes of point sets used for SIC-ICP experiment are (see fig. 8):

	model	test
Base	586	189
Cover	365	178
Tape	377	123

In summary, a SIC-map stands for a model and a test. Its experimental assessment involves the following steps:

- Take multiple range images of a real object and build its 3D model
- Take a new range image of the same object and use it as test
- Perform a first successful match of test and model and use it as pole configuration
- Compute the SIC-map by exploring a number of discrete configurations in the (ϕ, θ, ω) space, applying ICP matching and deciding whether convergence is successful or not.

Series 1

Fig. 7 shows the SIC-maps of three toy objects obtained experimentally according to above exposed method. Zenith angles take the values $\phi=0°$, $10°$, $20°$, ...$60°$

As might have been expected, we observe that the SIC-subrange for ω (represented by the black sector) decreases for view points with increasing zenith angle. For the duck and the swan, the effective decrease starts at about $\phi=40°$ while for the fish, it starts earlier, at about $\phi=30°$. Globally, it is obvious that the duck and swan have larger SIC-ranges than the fish.

Noting ω_Δ the width of the ω rotation range, ϕ_Δ the size of the ϕ azimuth range, and $R_\Delta=(\phi_\Delta,\omega_\Delta)$ the subrange of C whose points belong to both ranges, we can have simplified subranges of a SIC-range, $R_\Delta \subset$ SIC-range, that provide a conservative description of it. i.e.

- $R_{\Delta fish}=(30°,120°)$
- $R_{\Delta swan}=R_{\Delta duck}=(40°,90°)$

In any case, it is interesting to note that the SIC-range for all objects is quite large, showing that the geometric matching of these objects can take advantage of a relative comfortable convergence behavior towards the optimal match.


a) Duck

b) Fish

c) Swan

Fig. 7. Toys and their SIC-maps

Series 2

A second series of experiments was conducted with 3 different parts of a scotch tape dispenser. For each part, a model object was built and three test objects were acquired, each with a different view point. 9 tests are thus matched, each with its own model. Images of the 9 test views are shown in figure 8 together with the associated SIC-maps.

The purpose of this new experiment is to analyze other objects and also the same object with different test views.
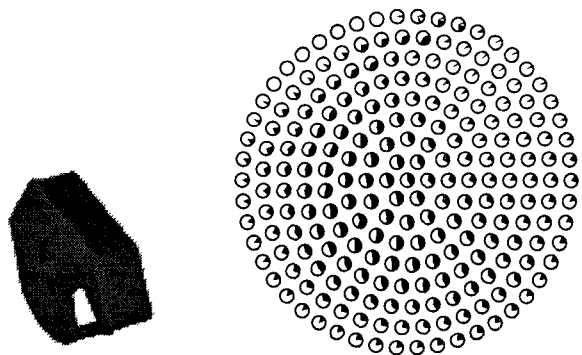
Compared to the first series, the objects are considered more complex, in particular, their overall shape exhibit more concavities.

a0) Base 0

b1) Cover 1

a1) Base 1

b2) Cover 2


a2) Base 2

c0) Tape 0

b0) Cover 0

c1) Tape 1

105

c2) Tape 2

Fig. 8. Scotch parts and associated SIC-maps

Let us discuss the 9 SIC-maps of figure 8 which cover here, compared to previous series, a somehow increased range of azimuth values: $\phi=0°, 10°, 20°, ...80°$.

A number of observations can be made. Compared to the SIC-maps of the previous series which were characterized by an approximately pole centered SIC-range, the new SIC-ranges are characterized by a strong off-pole dominance. Comparing for example Base0, Base1 and Base2, the dominant part of the SIC-range is found respectively in the NW (Base0), SW (Base1) and SE (Base0) regions on the SIC-map. Other maps show similar behavior.

The second observation refers to the inter-object and intra-object variability. If intra-object variability is smaller than inter-object variability, we would expect the existence of some model specific features. All results obtained in this series reveal a comparable intra-object and inter-object variability. We thus do not see any object specific features that could be used to improve the convergence.

Notice here that the object tape displays a very dark SIC-map and seems therefore to behave differently. This is in fact not the case because the darkening of the SIC-map has a different cause, namely the existence, with this scotch tape, of two optimal matching poses due to the object symmetry.

So far, the reason for the strong off-pole dominance of the SIC-range as well as its variation as a function of the view axis are still unknown.

Nevertheless, all measured SIC-maps exhibit a rather large SIC-range. To give but a partial idea of its size, we introduce again the conservative notation of a simplified subrange, and in addition, translate it to the dominant part of the SIC-range. A range which is roughly valid for all scotch part is:

* $R_{\Delta scotch}=(30°,60°)$

5. Conclusions

Considering geometric matching of 3D objects by the ICP algorithm, we presented the SIC-map as a means to analyze the convergence behavior towards optimum matching.

The SIC-maps of the toy objects revealed a pole centered SIC-map providing relatively comfortable

convergence towards optimum matching.

The SIC-maps of the scotch parts revealed the existence of possibly off-centered SIC-ranges but with similarly comfortable convergence towards optimum matching.

The experimental SIC-ranges have sizes which exceed a common conservative subrange of $\Delta\phi$ <30° for zenith angles and $\Delta\omega$ <60° for rotation angles.

Far from being complete these results give nevertheless a first insight in the quantitative nature of SIC-ranges.

6. Acknowledgments

7. References

[1]   R. Bergevin, D. Laurendeau and D. Poussart, "Estimating the 3D Rigid Transformation Between Two Range Views of a Complex Object," IEEE Proceedings of Int. Conf. on Pattern Recognition, pp. 478-482, 1992.

[2]   P.J. Besl and N.D. McKay, "A Method for Registration of 3-D Shapes," Proceedings of IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), vol. 14, no. 2, pp. 239-256, 1992.

[3]   Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," International Journal of Image and Vision Computing (IVC), vol. 10, no. 3, pp. 145-155, 1992.

[4]   Z. Zhang, "Iterative point matching for registration of free-form curves and surfaces," International Journal of Computer Vision, vol. 13, no. 2, pp. 119-152, 1994.

[5]   J. Feldmar and N. Ayache, "Rigid and Affine Registration of Smooth Surfaces using Differential Properties," Lecture Notes in Computer Science, vol. 801, no. 2, pp. 397-406, 1994.

[6]   H. Hugli, Ch. Schutz, D. Semitekos, "Geometric matching for free-form 3D object recognition," ACCV, Singapore, vol. 3, pp. 819-823, 1995.

[7]   Ch. Schutz, H. Hugli, "Towards the recognition of 3D free-form objects," Intelligent Robots and Computer Vision XIV, Algorithms, Techniques, Active Vision and Materials Handling, Proc. SPIE, Vol. 2588, pp. 476-484, 1995.

[8]   G. Blais and M. Levine, "Registering multiview range data to create 3D computer objects," Proceedings of IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), vol. 17, no. 8, pp. 820-824, 1995.

[9]   J.-L. Chen and G.C. Stockman, "Determining pose of 3D objects with curved surfaces," Proceedings of IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), vol. 18, no. 1, pp. 52-57, 1996.

[10]   C.S. Chua, R. Jarvis, "3D Free-Form Surface Registration and Object Recognition," International Journal of Computer Vision, Kluwer Academic Publishers, vol. 17, pp. 77-99, 1996.

[11]   Heinz Hügli, Christian Schutz, "How well performs 3D object recognition from range images?", Proc. SPIE Vol 2904-9, 1996

[12]   O.R. Faugeras and M. Hebert, "The Representation, Recognition, and Locating of 3-D Objects", The International Journal of Robotics Research, Vol 5, No.3, 1986