Intrinsic correspondence using statistical signature-based matching for 3D surfaces

Birgit M. Planitz and Anthony J. Maeder
CRC for Satellite Systems
Queensland University of Technology
GPO Box 2434 Brisbane, QLD 4001
{b.planitz, a.maeder}@qut.edu.au

John A. Williams
School of ITEE
University of Queensland
St Lucia, QLD 4072
jwilliams@itee.uq.edu.au

Abstract

A wide variety of applications including object recognition and terrain mapping, rely upon automatic three dimensional surface modelling. The automatic correspondence stage of the modelling process has proven challenging. Intrinsic correspondence methods determine matching segments of partially overlapping 3D surfaces, by using properties intrinsic to the surfaces. These methods do not require initial relative orientations to begin the matching procedures. Hence, intrinsic methods are well-suited for automatic matching.

This paper introduces a novel intrinsic automatic correspondence algorithm. Local feature support regions are described using distance and angular metrics, which are used to construct cumulative distribution function signatures. Local correspondences are hypothesised by comparing the signatures of two surfaces. A geometric consistency test is then applied to select the best local correspondences. Finally, registrations are computed from the remaining correspondences and the best alignment is selected. Results demonstrating the algorithm's accuracy in selecting correspondences for mutual partially overlapping surfaces, are presented. The algorithm's parameters prove robust, with only the local region size being surface dependent.

1 Introduction

Automatic correspondence is an important step in three dimensional (3D) modelling. Automatic correspondence is essential in applications where the position of the sensor, with respect to the scene, is unknown. A typical example is terrain mapping [8]. It is desired that the images of all views (i.e. 3D surfaces) are input into the modelling system, where they are automatically manipulated to form a 3D model of the object/scene. This section outlines the 3D modelling process, discusses intrinsic correspondence, and briefly highlights existing intrinsic correspondence algorithms.

The 3D modelling process consists of four main stages. First, a sensing device is used to obtain 3D surfaces of different views of an object/scene in the **data acquisition** stage. Secondly, the matching segments of different surfaces are found using a **correspondence** algorithm. Thirdly, the corresponding surfaces are aligned by applying a **registration** scheme. Finally, the aligned surfaces are merged to form a complete 3D model of the object/scene, in the **integration and reconstruction** phase. The automatic data acquisition, registration, integration and reconstruction stages have been more or less solved. Automatic correspondence however, has proven challenging.

Correspondence methods can be categorised as either *intrinsic* or *extrinsic* [12]. Intrinsic methods form correspondences by comparing the intrinsic properties of surfaces, whereas extrinsic methods form correspondences using the relative orientations between the surfaces being matched. Extrinsic methods require a rough initial alignment between the surfaces to converge to the correct solution [3, 5]. Intrinsic correspondence-registration methods automatically form these initial alignments within the algorithm. Therefore, intrinsic techniques are the key to developing automatic correspondence algorithms, because theoretically no user interaction is required. However, no *fully* automatic technique exists.

Some key intrinsic methods are highlighted as follows. The Random Sample Consensus based Data Aligned Rigidity Constrained Exhaustive Search method defines regions on one surface X, and searches for regions of similar size on the other surface Y [4]. Triangles comprised of selected control points make up the regions. A similar method is graph matching, whereby a graph using distances between points is constructed on X [6]. The algorithm then attempts to build the same (or part of the same) graph on Y. Methods such as spin-image and geometric histogram matching treat the correspondence problem slightly differently. The former creates signatures on each surface, that are based on the horizontal and vertical distances from selected points to every other point on the surface [9]. The latter examines the angles and vertical distances from a given mesh facet, to other facets within a predefined distance [2].

Aspects of the outlined correspondence methods are referred to in other sections of this paper. Intrinsic correspondence is detailed in the following section. Section 3 then introduces a novel intrinsic correspondence algorithm. Section 4 provides results of matches between mutual partially overlapping surfaces. Finally, Section 5 summarises the work outlined in this paper.

2 Background

As discussed developing an intrinsic method is the best solution to constructing an automatic correspondence method. In this section a typical intrinsic algorithm is outlined, and its main components are reviewed in detail.

2.1 Intrinsic Correspondence Dissected

Figure 1 illustrates a typical approach to intrinsic correspondence. This method assumes pairwise correspondence and registration. When registering a set of surfaces, it is sufficient to do the initial alignment in a pairwise manner, matching each surface with every other surface individually. The final accurate alignment between all surfaces can be obtained using a multiview registration scheme [14].

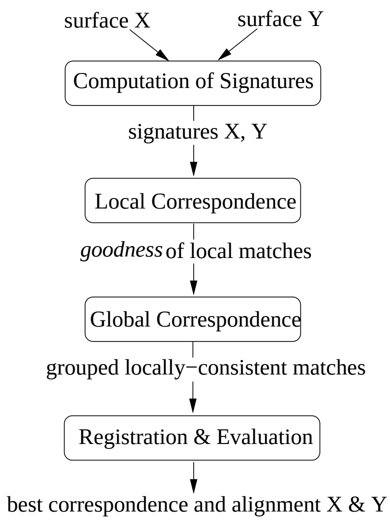

Figure 1. The steps in a typical intrinsic correspondence algorithm.

Figure 1 illustrates a procedure, where the number of potential correspondences between two surfaces are continuously reduced, such that only a few remaining hypothesised matches are passed to the registration-evaluation process. Essentially, this algorithm is an exercise in *correspondence pruning*. The first step of the pruning process is to select small local feature support regions on both surfaces X and Y. Then, these regions are given signatures based on their local features. The signatures of X are matched with the signatures of Y in the local correspondence phase, and the *goodness* of each match is passed to global evaluation procedure. Here, bad matches are discarded, reducing the number of possible local correspondences. The remaining matches are accumulated to provide evidence for consistent local matches. The sets of consistent matches are then used to compute registrations to bring X and Y into common coordinate systems, where their alignments are evaluated. The best alignment gives rise to the optimal correspondence between the two surfaces.

A more detailed review of the steps shown in Figure 1 is discussed in the following sections.

2.2 Signatures

Signatures of surface regions are constructed from intrinsic surface properties. The signature construction process encompasses two very important steps: choosing a viable descriptor of a region, and storing this descriptor in a signature that can easily be compared to other signatures. Some considerations when selecting descriptors, region sizes and signatures are highlighted below.

It is desired that a surface descriptor is robust, and unique to its local neighbouring region. Regions can be uniquely described using angular [2], distance [2, 4, 6], and differential [10] features. The following examples concern robustness. Distances between points are robust, as they use the original surface properties [11]. However, features such as normals, are less robust because smoothing is usually required so that the descriptors are accurate [5].

Selecting the size and shape of the local support regions, is a process that concurs with choosing a region descriptor. In some cases the regions' size and shape may be constant [2], and in others variable [4]. It is essential that each local support region is large enough to store a unique description of the area, but not too large, as the entire region may not be in the overlapping portion of the corresponding surfaces.

After selecting one or more surface descriptor(s) and defining the region size and shape, the descriptors of each support region must be stored as signatures. Typical signatures include graphs [6] and histograms [2]. The signatures must be of reasonable size to ensure that local matching is efficient. Signatures such as spin-images are less favourable because they require large data storage space [9].

Once the signature computation process is complete, the signatures of surfaces X and Y are passed to the local match algorithm.

2.3 Local Correspondence

The function of the local correspondence algorithm is to determine how well the signatures of surface X match with the signatures of surface Y. A suitable match metric is required for this operation. Examples include the Minkowski norm for shape distribution matching [11], and Bhattacharyya distance for histogram matching [2].

Local correspondences are generally computed in batches, as shown in Figure 1, where every signature of one surface is compared with every signature on the other. The *goodness* of all local matches is then passed to a global evaluation algorithm.

2.4 Global Correspondence

Global correspondence algorithms prune all local matches, so that only good locally-consistent matches are selected. An example is pruning a probability matrix, where the entries of the matrix are *probabilities* of matches between signatures on X and Y. The global correspondence algorithm is used to *thin* the p-matrix by only selecting matches with a probability greater or equal to an acceptance level $p0$.

There are a number of ways of selecting global correspondences. This process is often integrated into the local matching procedure, such that bad matches are discarded immediately [4], and not all matches are passed in a batch for global evaluation. However, analysing a batch of matches can be extremely valuable when using pruning methods, because the correspondences can then be evaluated at different acceptance levels.

The correspondences that pass global evaluation are a much smaller set of possible matches between two surfaces, and these are passed to a registration-evaluation algorithm.

2.5 Registration and Evaluation

The final analysis of potential matches between two surfaces is completed by evaluating how well two surfaces align when using selections of the remaining correspondences. Three corresponding pairs of regions must be selected to uniquely align two surfaces in a common coordinate system.

A number of registration-evaluation methods are used to select the best possible transformations to align two surfaces. In most cases all possible combinations of three corresponding region pairs are used to compute registrations between the surfaces. This is generally followed by another evaluation where extrinsic-type metrics [3, 5] are computed to further prune the correspondence space. The final registrations are examined by using methods such as an evidence accumulation scheme to test for consistent transformations [2], or a model based scheme to test for good alignments [4].

The best alignment between the two surfaces is selected as the final outcome of the intrinsic correspondence-registration algorithm.

3 Statistical Signature-based Matching

In this section, a novel intrinsic correspondence algorithm is introduced. The algorithm conforms to the structure outlined in Section 2, and the following subsections discuss signature selection, local and global correspondence, registration and alignment evaluation.

3.1 Signatures

In this correspondence algorithm, distance and angular descriptors, and statistical signatures are used to characterise regions on the surfaces being matched. It is assumed that the surfaces are stored as polyhedral (generally triangular) meshes. The following paragraphs discuss the signature derivation process, with respect to the mesh surfaces.

Two descriptors are utilised to describe a local region on the mesh. The $D1$ distance is a robust descriptor, and was derived for surface model matching [11]. The metric is the Euclidean distance between a centre vertex and points on the surface, in the local feature support region. The second metric is the $A1$ descriptor, which is the angle between the normal of the centre vertex and the normals of points, in the local feature support region. Angles between facet normals are used as feature metrics in geometric histogram matching [2]. Both metrics are used in a combined fashion to more uniquely portray each local support region on the surfaces being matched.

Local support regions are selected around each centre vertex as follows. First, the *border layers* of each mesh are determined. This is done by selecting the vertices on the border of the mesh (layer 1), then the vertices that connect to border vertices (layer 2), then the vertices that connect to layer 2 vertices (layer 3), and so on as shown in Figure 2. The mesh can now be evaluated at a certain level. For example, if there are many points on the mesh, it would be very inefficient to evaluate the potential match between every point on mesh X and every one on Y. Therefore, only the vertices of the innermost layers (say layer 3 and above are selected).

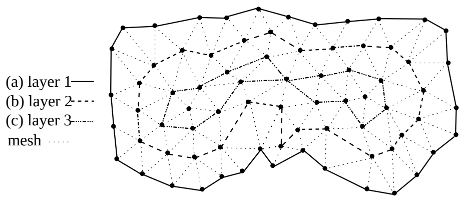

Figure 2. The border layers of a mesh, with (a) being the outermost layer, (b) the second, (c) the third and so forth.

Once the vertices to be evaluated have been selected, a neighbourhood radius is chosen (this is currently a user-defined value). This distance determines the surrounding neighbourhood of each vertex, by including all the points in the neighbourhood that fall into the sphere generated by the radius. The $D1$ and $A1$ metrics are then calculated for each local support region. The combination of distance and angular descriptors has proven powerful in correspondence algorithms [2, 9].

A separate signature is then built for both the $D1$ and the $A1$ metrics. The signatures are cumulative distribution functions (cdfs), and were selected in concurrence with the local match metric discussed in the following section.

3.2 Local Correspondence

A brute force local matching algorithm is employed, where the signature of every vertex under evaluation on mesh X is tested against the signature of every selected vertex on Y. The match metric employed is the Kolmogorov-Smirnov two sample test (KS-test). The KS-test tests whether two samples, that are drawn independently, belong to the same population [7]. The test statistic used in the two-sided KS-test is T, which is the greatest absolute distance between the two cdfs supplied for each match. The acceptance level, or probability of a match p, is calculated using T [7].

All $D1$ signatures of the two surfaces X and Y are compared, and all $A1$ signatures of X and Y are compared. The comparison results are stored in 2D probability matrices, P_D and P_A respectively. The respective elements of the matrices are then multiplied to form the p-matrix P, which is passed to the global correspondence algorithm for further evaluation.

3.3 Global Correspondence

The global correspondence method prunes local matches by examining geometric consistency. The first step of the global correspondence method however, is discarding local matches in the p-matrix that have probabilities below $p0$. The parameter $p0$ is the acceptance value that supports the hypothesis that two signatures belong to the same distribution.

Every possible combination of three local correspondences is then chosen from the remaining matches. The local correspondence pairs form triangles as shown in Figure 3. Geometric consistency must exist between the two surfaces if the formation of local correspondences are to be accepted. That is $d1_X \approx d1_Y$, $d2_X \approx d2_Y$, and $d3_X \approx d3_Y$. The measurement used to test similarity in the three distances is $|d_X - d_Y| \leq \tau$, where τ is a percentage (also user defined) multiplied by $max(d_X, d_Y)$. The varying nature of τ ensures that good matches based on larger triangles are more acceptable then those based on smaller ones. This increases the robustness of the algorithm, minimising the error introduced when considering small triangles as the best correspondences.

Every possible combination of three local correspondences is evaluated, and only the Q matches that pass the geometric consistency test are passed to the registration-evaluation algorithm, for further evaluation.

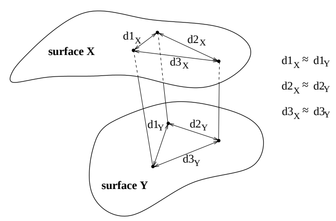

Figure 3. Three local correspondences between surface X and Y. The distance between the selected vertices on the surfaces must be similar for good geometric consistency.

3.4 Registration and Evaluation

Each possible combination of three centroids, of the Q remaining matches, is supplied to the alignment algorithm [1]. These registrations are then analysed visually. This process will soon be replaced by an extrinsic-type correspondence-alignment-evaluation scheme. A typical evaluation method is to compute the closest points from mesh X to mesh Y, and then determine the distances between them [13]. Finally, the number of point-pairs whose distances fall below a threshold τ are summed, and the match with the highest sum is selected as the best correspondence between the two meshes.

3.5 The Algorithm Summarised

The novel intrinsic correspondence method discussed in the previous sections is summarised in Figure 4.

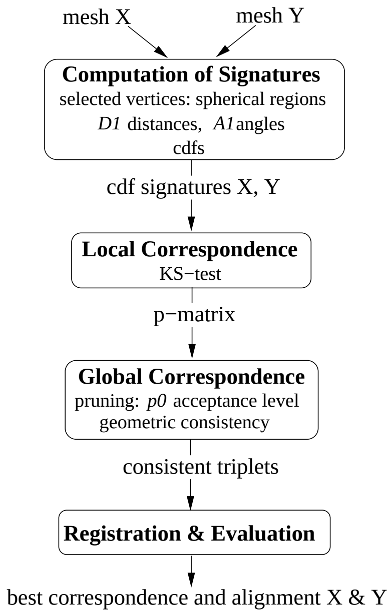

Figure 4. The steps of the statistical signature-based matching intrinsic correspondence algorithm.

4 Results

In this section, the results of the statistical-based signature matching algorithm are illustrated for surfaces with mutual partially overlapping segments. Both visual results and an analysis of the effects of the parameters on the algorithm are highlighted.

The surfaces matched are triangulated meshes and are displayed on the left hand side in Figure 5. The best triplets of corresponding points are shown in the right hand column. Note that at least three matches are required to register two surfaces.

Figure 5(a) shows two views of a Renault figurine, seen at 90 and 135 degree viewing angles. The correspondence triplets demonstrate accurate matches between the two surfaces. Figure 5(b) shows two very similar views of a toy dinosaur (viewing angles 0 and 360 degrees). The parameter $p0 = 0.99999999$ was chosen to reduce the possible number of remaining matches to four. When two surfaces are almost identical, the algorithm results in many accurately matched local support regions. Layers four and above were chosen for the match, although far fewer could have been selected because of the similarities of the two surfaces.

Figure 5(c) illustrates the best match when two surfaces are matched that contain a smaller percentage of mutual partially overlapping segments (views 0 and 36 degrees of the toy dinosaur) . The triplet of matches are not as accurate

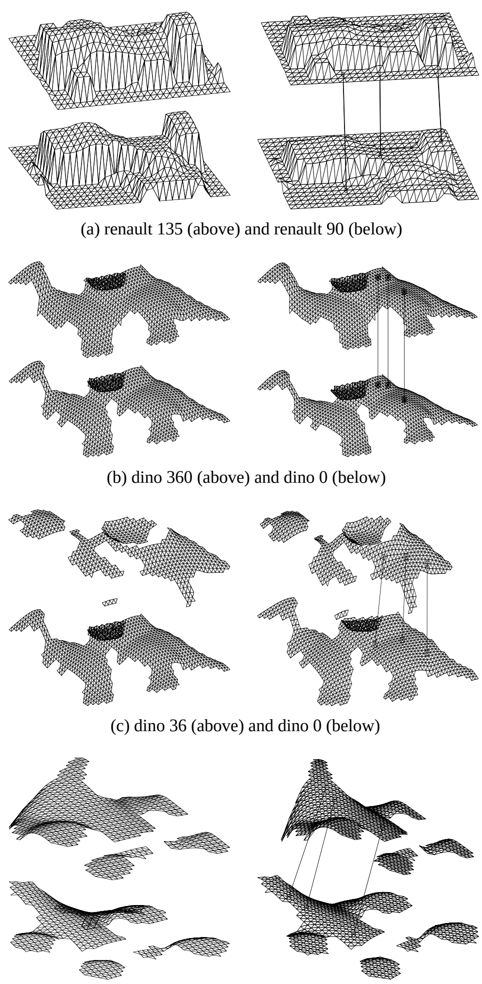

(a) renault 135 (above) and renault 90 (below)

(b) dino 360 (above) and dino 0 (below)

(c) dino 36 (above) and dino 0 (below)

(d) dino 252 (above) and dino 216 (below)

Figure 5. The surfaces tested (left), and the best resulting correspondence triplets (right).

as that in part (b), but correspond to regions of very similar surface variation. Figure 5(d) too shows the resulting correspondences when matching two surfaces contain a smaller percentage of mutual partially overlapping segments (views 216 and 252 degrees of a toy dinosaur). The triplet of best matches re-affirms the accuracy of the algorithm.

Table 1 summarises the parameter selection for the surfaces used to test the algorithm. The table also includes the radius size selected for each match procedure. The radius size was constant for each of the dinosaur match procedures, highlighting the robustness of the algorithm. The radius size selection will be the most important factor when fully automating the intrinsic algorithm. Other parameters such as τ are robust for a variety of surfaces and do not require adjustment. The neighbourhood layer region is mainly selected to increase the computational efficiency of the algorithm, and will only affect the correspondences if too few points are

Parameters	renault (90,135)	dino (0,360)	dino (0,36)	dino (216,252)
$p0$	0.90	0.9999 9999	0.95	0.945
radius	50	25	25	25
layers	≥ 3	≥ 4	≥ 3	≥ 3
tau	0.05	0.05	0.05	0.05

Table 1. The values selected for the parameters of the novel correspondence algorithm.

selected for matching.

The algorithm is currently implemented in MATLAB and takes no more than 30 minutes to run for larger meshes (e.g. *dino 0* and *dino 360*, with 964 and 1038 vertices respectively), depending on the parameter selection. Note that MATLAB is very slow in comparison with C or C++. The final algorithm will be implemented in C or C++ to dramatically reduce the computation time.

The results in this section highlighted that the novel intrinsic correspondence technique accurately selects correspondences between partially overlapping surfaces. The algorithm provides results that can be input into an extrinsic algorithm to achieve accurate final correspondences. Some further analysis required before the algorithm is complete includes examining the effects of the algorithm's parameters, and the algorithm's efficiency.

5 Conclusion

This paper highlighted the importance of intrinsic correspondence techniques, and introduced a novel intrinsic algorithm. The new method is a signature matching technique that uses the $D1$ distances and $A1$ angular measurements as descriptors of local support regions. The signatures are cdfs, and are compared by the local correspondence algorithm. The comparisons are made using a KS-test and are stored in two probability matrices, whose respective elements are multiplied. The resulting p-matrix is pruned by accepting only those local matches which are greater or equal to the acceptance value $p0$. Global correspondences are then tested for geometric consistency to further reduce the match search space. Finally registrations are applied to the best matches, and visual evaluations of the best alignments are made.

The initial results of the statistical-based signature matching algorithm are promising. The algorithm provided accurate alignments that can be input as initial relative orientations in extrinsic methods. The algorithm's only surface dependent parameter is the local support region radius size, making the technique robust. Future work includes fully automating the parameter selection process and coding a registration-evaluation algorithm to finalise the automatic correspondence process.

ACKNOWLEDGMENT: The Cooperative Research Centre for Satellite Systems is established and supported under the Australian Government's Cooperative Research Centres Program.

References

[1] K. S. Arun, T. S. Huang, and S. D. Blostein. Least-squares fitting of two 3-D point sets. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 9(5):698–700, Sept. 1987.

[2] A. Ashbrook, R. Fisher, C. Robertson, and N. Werghi. Finding surface correspondence for object recognition and registration using pairwise geometric histograms. In *Proc. European Conference on Computer Vision (ECCV98)*, pages 674–686, 1998.

[3] P. J. Besl and N. D. McKay. A method of registration of 3-D shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 14(2):239–256, Feb. 1992.

[4] C. Chen and Y. Hung. RANSAC-Based DARCES: A new approach to fast automatic registration of partially overlapping range images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(11):1229–1234, November 1999.

[5] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, Apr. 1992.

[6] J. Cheng and H. Don. A graph matching approach to 3D point correspondences. *Int. Journ. of Pattern Recognition and Artificial Intelligence*, 5(3):399–412, 1991.

[7] W. J. Conover. *Practical Nonparametric Statistics*. John Wiley & Sons, Inc., 1999.

[8] D. F. Huber and M. Hebert. A new approach to 3-d terrain mapping. In *Proceedings of the 1999 IEEE/RSJ Int. Conf. on Intelligent Robotics and Systems (IROS '99)*, pages 1121–1127. IEEE, October 1999.

[9] A. Johnson and M. Hebert. Surface registration by matching oriented points. *IEEE Proc. Int. Conf. on Recent Advances in 3-D Digital Imaging and Modeling*, pages 121–128, May 1997.

[10] P. Krsek, T. Pajdla, V. Hlavac, and R. Martin. Range image registration driven by hierarchy of surface differential features. In *Proceedings of the 22nd Workshop of the Austrian Association for Pattern Recognition*, pages 175–183, May 1998.

[11] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. Matching 3d models with shape distributions. In *Int. Conf. on Shape Modeling and Applications 2001*, pages 154–166, Genova, Italy, May 2001.

[12] B. M. Planitz, J. A. Williams, and M. Bennamoun. Automatic correspondence of range images. In *The Fifth Asian Conference on Computer Vision (ACCV2002)*, pages 562–567, January 2002.

[13] G. Roth. Registering two overlapping range images. In *Proc. IEEE Conf. on 3-D Digital Imaging and Modeling*, pages 191–200, 1999.

[14] J. A. Williams and M. Bennamoun. Simultaneous registration of multiple corresponding point sets. *Computer Vision and Image Understanding*, 81(1):117–142, Jan. 2001.