Using the Correspondence Framework to Select Surface Matching Algorithms

B. M. Planitz and A. J. Maeder
e-Health Research Centre, ICT CSIRO
Brisbane, QLD 4000
Birgit.Planitz@csiro.au

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

Abstract

A correspondence framework has recently been proposed to unify a wide variety of surface matching algorithms, and provide a consistent structure for establishing new ones. When an algorithm is implemented using the framework, it is divided into five stages. A module is created for each stage of the framework, and that module is placed in a library (for that stage of the framework). Algorithms are created by connecting five appropriate modules from the library. It is envisaged that in the future, algorithms will be created by automatically connecting five suitable modules for their specific surface matching tasks. This paper takes a step towards this goal, by presenting a metric for assessing the outcomes of the final stage of the framework. The metric provides a quantitative value that determines the suitability of an algorithm for a specific task. Six algorithms are presented and their suitability over a range of surfaces is tested. Results show that the outcome of each experiment reflects the expected outcome. Thus, the metric is an appropriate tool for algorithm selection. Future directions at the end of the paper discuss the concept of using metrics at the other stages of the framework, so that the automatic algorithms selection process can be realised.

1 Introduction

A significant body of research is available in the field of three dimensional (3D) surface correspondence establishment. Correspondence computation is the process of establishing mappings between two rigid surfaces. It is used to determine which portions of the two surfaces overlap.

An abundance of algorithms has been developed for computing the coarse initial mappings between two surfaces. However, no single algorithm has prevailed, which can match any two arbitrary surfaces. This is due to the fact that algorithms are application specific, as they place restrictions on the types of the input surfaces they can match [10]. When given a particular matching task, a suitable algorithm must be selected (or created) for that task. Until recently however, research into the 3D surface correspondence problem was hindered by a lack of uniform technology, and the absence of a consistent model for comparing existing approaches and developing new ones.

A correspondence framework for surface matching algorithms has been presented to address these issues [10]. The framework has been derived from the perspective of rigid surface correspondence, which constitutes a major subcategory of surface correspondence. It is both a conceptual model and a software design tool, which facilitates the analysis, comparison, development and implementation of rigid surface matching algorithms. It is general, unifying a wide variety of existing algorithms using consistent terminology. It is also flexible, enabling the the synthesis of powerful new algorithms.

The framework divides the process of correspondence into five stages. Algorithms are implemented as a series of five modules, one for each stage of the framework. A future objective of the framework is to use it for automatic algorithm creation. That is, a method would be used to select the five best modules (from modules that are available in a framework library) for a given surface matching task. This paper takes a step in the direction of automatic algorithm selection, by presenting a quantitative metric for assessing the outcomes of the final stage of the framework.

The paper begins by outlining the framework in Section 2. Section 2 also presents six algorithms whose components already exist within the framework library. The metric and method for assessing the outcomes of the final stage of an algorithm are presented in Section 3. The metric is then used to assess the suitability of each of the six aforementioned algorithms over a variety of surfaces, in Section 4. The expected and actual results are compared. Section 5 then discusses future work with regards to completely automatic algorithm selection using the framework library. Finally, Section 6 summarises the paper with concluding remarks.

2 The Correspondence Framework

The correspondence framework is both a conceptual model and a software design tool for surface matching algorithms. The framework consists of five stages: region definition, feature extraction, feature representation, local matching, and global matching [10]. When matching pairwise surfaces, the framework is employed as demonstrated in Figure 1.

As a conceptual model, the framework enables the researcher to analyse each of the five stages of a surface matching algorithm on its own accord [10]. The stages of one algorithm are directly comparable to the stages of another. Algorithms are developed by connecting five appropriate stages of existing algorithms.

The individual functions of the stages of the framework are described briefly below. For further information on the framework and algorithm selection/creation, the reader is referred to [10]. The first stage of the framework, region definition is the stage where localised regions are selected on both input surfaces. Feature extraction is the stage where intrinsic surface properties are extracted from regions. Feature Representation is the stage where features extracted from regions are represented in a way so that they are comparable to other feature representations. Local Matching is the stage where local correspondences are hypothesised between two surfaces, and grossly erroneous matches are rejected. Global Matching is the stage where global correspondence and the subsequent coarse initial alignment between two surfaces are computed.

Four existing algorithms and two new algorithms have been developed to fit within the framework: Spin-image Matching (SIM) [6], Geometric Histogram Matching (GHM) [1], Intrinsic Curve Matching (ICM) [7], Random Sample Consensus based Data Aligned Rigidity Constrained Exhaustive Search (RBD) [4], SIM with RBD (SIM-RBD) [9], and $D2$ Signature Matching with RBD (DSM-RBD) [11]. These algorithms and the types of surfaces they are designed to match are highlighted in Table 1. The following section introduces a quantitative method for assessing correspondence algorithms, which will be used to determine whether the expected suitability of each algorithm listed in Table 1 is correct.

Algorithm	Expected Suitability
SIM	a wide variety of surfaces, except those that exhibit symmetry about an axis of rotation
GHM	surfaces with a smooth topological variations and a significant amount of mutual overlap
ICM	smooth surfaces with relatively high resolution and significant topology
RBD	a wide variety of surfaces, particularly featureless pairs with significant overlapping segments
SIM-RBD	a wide variety of surfaces, more robust against symmetry than SIM
DSM-RBD	a wide variety of surfaces, particularly featureless surface pairs with fewer overlapping segments than RBD can handle

Table 1. Six correspondence algorithms, and the surface types they are designed to match.

3 Assessing the Quality of Global Correspondences

The general method for assessing the accuracy of a global correspondence (mapping) between two surfaces is performed as follows. First, the global correspondence is established. The mapping is then used to compute the registration parameters, which align both surfaces in a common coordinate frame. For rigid surfaces, the registration parameters are a rotation **R** and a translation **T**. The accuracy of the alignment is then assessed by determining the proximity between the overlapping segments of the surfaces.

There are two important factors in registration assessment. The first is the establishment of Extrinsic Point Correspondences (EPCs) between surfaces, and the second is the selection of the metric that is used to measure the proximity of the overlapping segments of two surfaces. Both these factors are discussed in the following subsections, where the most generic metric is selected to test the six algorithms that were presented in Section 2.

3.1 Extrinsic Point Correspondence Establishment

Given two surfaces X and Y, EPC establishment implies specifying a mapping between a point on X and one on Y, where the points are close to one another. Some common restrictions that determine whether or not an EPC is valid are [12]:

- the distance between the points must be below a preset threshold; and
- the angle between the surface normals of the two points must be below a preset threshold.

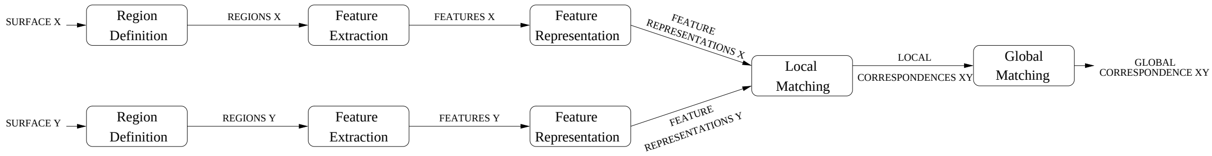

Figure 1. The correspondence framework

In addition to this, only the p% of closest correspondences may be used. The remaining $(100-p)$% are discarded to remove the possibilities of matching non-overlapping points. Also, only non-boundary points (on surfaces meshes) can be used as EPCs, to reduce boundary errors.

In some algorithms, one point X may only match with a single point on Y (for example [7]). However, generally more correspondences are used (for example [2, 13]). The method presented in this paper is the latter, as it is a more generic approach to EPC establishment.

3.2 Measuring the Proximity of Two Surfaces

Given a set of EPCs, that adhere to the aforementioned restrictions, a metric is required that quantifies the proximity of two surfaces. This section lists a few metrics, and selects the most commonly used one to measure the performance of global correspondences.

Given a set of EPCS, some common metrics are:

- counting the Number of Point (NP) correspondences in the set [4];
- accumulating the Surface Area (SA) of the immediate neighbourhoods surrounding the EPCs [1]; and
- computing the Mutual Information (MI) between the surfaces using the EPCs [13].

The metric that is used in this paper is NP. NP is generally more robust than SA and MI for the following reasons. For MI, a greater number of EPCs need to be established than for NP. NP selects only the best EPCs, and is thus a more robust metric. SA is very sensitive to surface resolution, whereas NP can be applied to a greater variety of data. In the next section NP is used to test the performance of the six algorithms presented in Section 2.

4 Results

The objective of this paper is to provide a quantitative metric that can be used to assess the suitability of an algorithm for a particular surface type. This section presents six different surface pairs, which are matched using the algorithms presented in Section 2. The surfaces are compared in terms of acquisition, topology, and degree of overlap. The results of matching each surface pair using each algorithm are then presented, and the actual versus expected outcomes for each algorithm are discussed.

4.1 Test Data

The test pairs used in the experiment are presented in Figure 2. Note that the surfaces are highly subsampled versions of the original data, so that the robustness of the algorithms can be examined. The registered surfaces column of Figure 2 demonstrates that a perfect alignment between two low resolution surfaces is not possible. Thus, the relative heights of the two surfaces are shown. The surface segment (light for X and dark for Y) closest to the reader is highlighted. A summary of the mode of acquisition, degree of overlap, and topology of the surfaces is presented below.

The SCENE surface pair was captured using a mobile unit equipped with a structured light sensor [14]. The surfaces are displayed as triangular meshes, containing over 2500 vertices each. The data is typical of an indoor scene, containing sharp edges and planar facets.

The ANGEL surfaces were captured by placing an angel figurine on a turntable, and using a Minolta 700 range scanner to acquire views of the figurine at different rotations [3]. The triangular meshes shown are similar in size, both over 800 vertices each. The surfaces have distinct topologies and overlap significantly.

The DINO surface pair was acquired using the same scanner and process as the ANGEL pair [3]. The two DINO meshes vary greatly in size, with the first having 964 and the second having 667 vertices. Both surfaces have distinct topologies. However, there is much less mutual overlap between them than the ANGEL pair. The overlap is limited to the back leg and tail of the dinosaur, and only small patches on the head and front leg.

The HUB surfaces are mesh representations of synthetic range images, which were created to test an object recognition algorithm [5]. The two meshes are similar in size, with the first and second consisting of 1096 and 1132 vertices respectively. Although the surfaces have a large percentage of overlap, they are highly symmetrical about the z-axis,

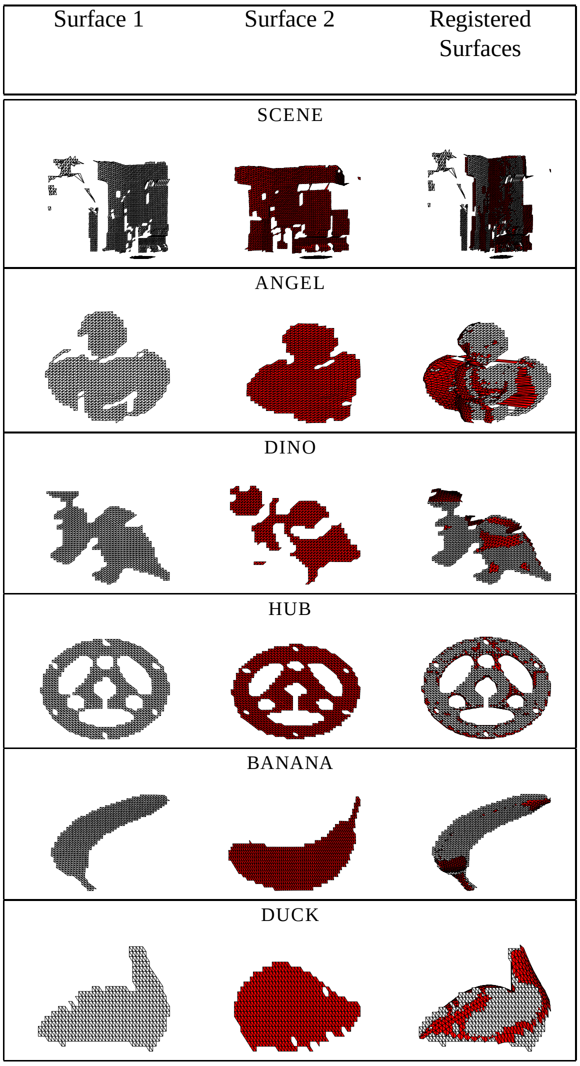

Figure 2. Test data: registered surfaces that have mutual partially overlapping segments.

which makes them difficult to match.

The BANANA surfaces originate from the same database as the HUB surfaces [5]. They are also mesh representations of synthetic range images, with the first and second containing 783 and 851 vertices respectively. The two surfaces also have a large percentage of overlap, however they lack distinct topology and varying curvature.

The DUCK surface pair was captured using a turntable, and a 3D-colour laser scanner [8]. The triangular meshes contain fewer than 550 vertices each. The only distinct feature in both surfaces is the sharp upward curve at the neck of the duck.

4.2 Outcomes

The NP scores achieved by matching each surface pair shown in Section 4.1 using each of six algorithms discussed in Section 2 are presented in Table 2. The NP scores are given as a percentage of the greatest number of possible correspondences that can be computed between two surfaces. These values are used to compare the actual with the expected outcome of each algorithm, which is discussed next.

Data	Algorithm					
	SIM	GHM	ICM	RBD	SIM-RBD	DSM-RBD
SCENE	84	80	19	74	75	77
ANGEL	85	66	70	67	23	47
DINO	83	20	71	55	50	70
HUB	0	75	34	96	95	70
BANANA	75	0	55	88	82	62
DUCK	86	0	89	85	80	85

Table 2. NP scores (%).

As expected, SIM produced highly accurate global correspondence results. Its only failure occurred on the HUB data set. This was expected, due to the symmetry of both HUB surfaces about the z axis. The NP values for SIM were generally very high ($> 75\%$) in all cases. This implies that a significant degree of overlap was found between surfaces. SIM performed better than all other algorithms for the SCENE, ANGEL, and DINO data sets. However, for the less topologically distinct BANANA data set, RBD and SIM-RBD produced higher NP values. This is due to the robustness of these algorithms for data with less distinct features. ICM produced a high, but only slightly better NP value than SIM for the DUCK surface pair, indicating that both algorithms match local feature representations accurately.

With the exception of the HUB surfaces, GHM produced poorer results than the SIM algorithm on all accounts. A high NP value ($> 80\%$) was achieved for the SCENE data set, and the ANGEL and HUB data sets achieved moderately high NP values ($65\% <$NP$< 75\%$). The NP scores indicate that GHM is not ideal for computing the correspondence between surfaces with fewer mutual overlapping segments, such as the DINO set. This is because only small segments overlapped in the coarse intial registration. GHM is also unsuitable for surfaces with few distinct topological variations, such as the BANANA and DUCK sets. The failure to achieve NP scores for these surfaces pairs was expected, as outlined in Table 1.

The only high NP value ($> 80\%$) achieved by ICM was for the DUCK data set. ICM produced accurate results for this data due to the data's smooth changes in curvature, which are required for feature extraction. The algorithm achieved moderately high results (65% <NP< 75%) for the ANGEL and DINO data, which also exhibited relatively smooth variations in curvature. A moderate NP value (55% <NP< 65%) was obtained for the BANANA surfaces. Because of their lack of smooth topology, NP scores of less than 50% were obtained for the SCENE and HUB surfaces. In summary, ICM performed as expected: better for surfaces with smoother curvature variation.

RBD is a recommendable algorithm for surfaces with few distinct topological features. This was evident in its very high NP scores ($> 85\%$) for the HUB, BANANA, and DUCK surface pairs. Moderately high NP values (65% <NP< 75%) were also obtained for the SCENE and ANGEL data sets, further demonstrating the robustness of the algorithm. RBD achieved a NP score of only 55% for the DINO surface pair. This was expected, as the algorithm is less likely to produce accurate matching results when the degree of mutual overlap between surfaces diminishes. In summary, it is recommended that this algorithm is very suitable for featureless surface pairs which have significant overlap.

SIM-RBD was expected to improve the robustness of the original SIM algorithm where surface symmetry is concerned. The NP score show that SIM-RBD did perform well on the HUB surface pair. The robustness of the RBD global matching module eliminated any false positive local matches produced by the SIM modules. SIM-RBD also it provided satisfactory results for surface pairs with fewer topological variations (BANANA, and DUCK), but was not as accurate as RBD. The SCENE result was almost equivalent to the RBD outcome. The ANGEL result was very poor, indicating that the algorithm is generally not as widely applicable as either the SIM or RBD.

DSM-RBD was expected to be a superior algorithm than the RBD for cases where surfaces contain a smaller degree of mutual overlap. DSM-RBD performed as expected. It produced a moderately high NP value of 70% for the DINO data set, almost 15% higher than the RBD result. Moderate to high results (NP$> 60\%$) were also achieved for the SCENE, HUB, BANANA, and DUCK surface pairs. The algorithm had difficulty with the ANGEL data, most likely due to the small regions, and non-optimised parameter values selected. Generally, this algorithm is recommendable for surfaces with few distinct topological features, and lower degrees of overlap. It is a solution to the problem that RBD is not suitable to handle, that is, the case where less mutual overlap exists between two surfaces.

In summary, it can be stated that each algorithm generally performed as expected. Therefore, using the NP metric metric with the specified EPC establishment scheme, is a suitable means of assessing global correspondences. This is an important step in the area of automatic correspondence algorithm selection (for given surface matching applications). The following section discusses using quality metrics at the other four stages of the framework, such that concept of complete automatic algorithm selection becomes conceivable.

5 Future Work

The correspondence framework provides a systematic approach for developing and implementing surface matching algorithms. This systematic approach gives rise to the possibility of using the framework to automatically select application specific algorithms. Given two surfaces, the five most appropriate modules (one for each stage of the framework) will be selected to compute the correspondences between the surfaces.

A step towards automatic algorithm selection was made in Section 4, where a quality metric was used to assess the final correspondences of each algorithm. Future work includes specifying evaluation metrics at each stage of the framework, such that the suitability of a module with respect to a particular surface type can be assessed. An example of an evaluation metric is as follows. For region definition, the metric may include information regarding storage requirements, size of regions, number of regions, and so on.

The five evaluation metrics would be included in an algorithm that sits outside the framework library. This algorithm would automatically select the five best modules for the particular task at hand. Examples of possible schemes are genetic algorithms and neural networks. It would be imperative to incorporate some learning capability into the scheme, such that particular modules are automatically selected for specific surface types. Note that the possibility of having a tool for automatic algorithm selection is only conceivable now that a systematic model for surface matching is available. Prior to the development of the correspondence framework, no such model existed.

6 Conclusion

This paper presented the results six surface matching algorithms that have been encoded within the correspondence framework. Four restructured and two new algorithms were tested. The objective of the paper was to demonstrate that the framework can be used to select algorithms for particular surface types. Each algorithm was used to match six surface pairs, and their correspondence results were evaluated by assessing the NP values of the registrations computed from the mappings. It was shown that each of the six algorithms does indeed favour particular surface types:

- SIM generally performs well across a wide variety of surfaces, but has difficulty in matching surfaces that exhibit symmetry about an axis of rotation;

- GHM is generally less accurate than the SIM, and would be more applicable to match surfaces of higher resolution, and with more topological variations;
- ICM only performs well on surfaces with smooth curvature variation;
- RBD is ideal for featureless surface pairs with significant degrees of overlap;
- SIM-RBD improves the robustness of SIM for surfaces that exhibit symmetry about an axis of rotation; and
- DSM-RBD is a good algorithm for surface pairs with fewer features and a smaller degree of mutual overlap than the RBD algorithm is accustomed to handling.

These results reflect the expect outcomes for each algorithm. Thus, the correspondence framework, in conjunction with the NP metric, is a suitable tool for selecting application specific algorithms.

Using the correspondence framework, future work will include developing a scheme for automatic algorithm selection. Section 5 discussed the concept of having a evaluation metrics at each stage of the framework, such that the best algorithm can be constructed for each particular application. It must be re-emphasised that automatic algorithm selection is only conceivable now that the framework, which is a systematic model for surface matching, has been developed.

References

[1] A. Ashbrook, R. Fisher, N. Werghi, and C. Robertson. Aligning arbitrary surfaces using pairwise geometric histograms. In *Proc. of Noblesse Workshop on Non-linear Model Based Image Analysis*, pages 103–108, Glasgow, Scotland, 1998.

[2] 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.

[3] R. Campbell and P. Flynn. A WWW-accessible 3D image and model database for computer vision research. In K. Bowyer and P. Phillips, editors, *Empirical Evaluation Methods in Computer Vision*, pages 148–154. IEEE Computer Society Press, 1998.

[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] G. Hetzel, B. Leibe, P. Levi, and B. Schiele. 3D object recognition from range images using local feature histograms. In *Proc. of the Int. Conf. on Computer Vision and Pattern Recognition*, volume 2, pages 394–399, Hawaii, 2001.

[6] A. E. Johnson. *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, Carnegie Mellon University, Pittsburgh, Pennsylvania, August 1997.

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

[8] Visual Information Technology, National Research Council of Canada, http://www.vit.iit.nrc.ca/VIT.html. Website, accessed: August 2003.

[9] B. M. Planitz, A. J. Maeder, and J. A. Williams. The correspondence framework for automatic surface matching. In *Proc. of Australian and New Zealand Conf. on Intelligent Information Systems*, pages 319–324, Sydney, December 2003.

[10] B. M. Planitz, A. J. Maeder, and J. A. Williams. The correspondence framework for surface matching algorithms. *Computer Vision and Image Understanding*, 97(3):347–383, Mar. 2005.

[11] B. M. Planitz, A. J. Maeder, and J. A. Williams. Synthesising surface matching algorithms using the correspondence framework. In *Int. Conf. on Pattern Recognition and Image Analysis: New Information Technologies*, St. Petersburg, October 2005. To appear: 2005.

[12] K. Pulli. Multiview registration for large data sets. In *Proc. 2nd Int. Conf. on 3D Digital Imaging and Modeling (3DIM'99)*, pages 160–168, Ottawa, Canada, Oct. 1999.

[13] A. Rangarajan, H. Chui, and J. S. Duncan. Rigid point feature registration using mutual information. *Medical Image Analysis*, 3(4):425–440, 1999.

[14] The Ohio State University, Department of Electrical Engineering. Website. http://sampl.eng.ohio-state.edu/s̃ampl/data/3DDB/RID/Struc_Light/index.html (accessed: November 2003).