The Correspondence Framework for Automatic Surface Matching

Birgit M. Planitz and Anthony J. Maeder
Cooperative Research Centre 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

Determining correspondence between sensed views of objects or scenes is a widely studied topic in computer vision. This paper examines correspondence, with particular focus on intrinsic rigid surface matching algorithms, where correspondences are generated using intrinsic surface properties. Many existing intrinsic correspondence algorithms use the same general approach to determine correspondences between surfaces. This paper details this general approach and shows how it forms the basis of the general correspondence framework. The framework is introduced as both a conceptual and an actual model that can be applied to a variety of correspondence tasks. Existing techniques are restructured to fit within the framework, and illustrate its generality. A novel algorithm is also constructed, to show how the framework facilitates innovation, that is, the synthesis of new techniques. Algorithms in the framework are tested to verify both generality and innovation. Future directions of the framework extend the potential testbed of the framework to incorporate a number of open computer vision correspondence issues.

1 Introduction

Correspondence is a fundamental task in computer vision, which typically precedes registration or a related task. In registration, correspondences between two or more views are used to compute the transformation parameters required to bring the views into alignment. A related task is stereo matching, where local correspondences between views are computed to determine the disparity between the left and right images of the same scene.

Formally, correspondence is the process of determining matching regions between views. In this paper, views are defined as renditions of objects or scenes, and regions are spatially consistent subsets of a view. An example of a view is a three dimensional (3D) surface, and region may be a small connected neighbourhood located on the surface.

This paper focuses on the problem of rigid 3D surface correspondence. This type of correspondence precedes the process of 3D surface registration. The problem where surfaces of the same object or view are acquired from different, unspecified viewpoints, is considered. The correspondence task is determining the matching surface regions to compute the registration parameters between them.

This paper specifically addresses *intrinsic* surface correspondence. Intrinsic correspondence refers to matching the intrinsic properties of regions between surfaces, rather than considering their spatial (*extrinsic*) differences. Because it relies only on surface properties, intrinsic correspondence is a powerful way of determining coarse matches between surfaces. Extrinsic methods are generally very good at refining the rough initial estimates. Therefore, surface correspondence methods are often intrinsic, followed by an extrinsic component (for examples, see [1, 8, 9]). The primary objective of this paper is to provide a framework for intrinsic correspondence methods, focusing on the rough initial correspondence estimates between views. Many suitable extrinsic methods exist, which can be used to refine the initial estimates (a popular example is [2]).

Determining initial correspondence estimates between surfaces is a challenging problem. This is primarily due to factors such as the unknown spatial relationships and unknown degrees of overlap between the surfaces. Existing correspondence algorithms combat these problems, but often suffer from weaknesses such as computational inefficiency [4], memory usage [8], and application specificity [9]. However, these problems are usually attributed to only one or two sections of the correspondence algorithms.

In this paper, the generalised correspondence framework is presented, which separates the task of intrinsic surface correspondence into five distinct components. By restructuring existing techniques to fit within the framework, it is easier to identify where the strengths and weaknesses of each method lie. The modular structure of the framework also allows for modules of different algorithms to be slotted into an existing algorithm, to test or improve the algorithm. Finally, new components of algorithms can easily be synthesised and added to the library of components that already exists in the framework structure.

Formally, the primary objectives of the framework are:
Generality: to show that a wide variety of existing algorithms follow the same generic approach; and
Innovation: to synthesise new correspondence methods, using the interchangeable modules of the framework.

The objectives outlined above are addressed in this paper by first examining a general approach taken by existing intrinsic correspondence methods, in Section 2. This approach is then formalised in Section 3, where the correspondence framework is outlined and each stage of the model is explained in detail. This section also addresses the framework's objectives by first restructuring existing algorithms to fit within the framework, and then presenting a new method that combines components of existing algorithms. The coarse registration resulting from testing the restructured and new algorithms are presented in Section 4. Finally, future directions of the framework are discussed in Section 5, and Section 6 concludes the paper with some summary remarks.

2 Background

This section examines a general approach taken by many existing intrinsic rigid surface correspondence methods. First, some terminology, used to understand each algorithm in a common frame of reference, is presented. This is followed by outlining the general approach taken, and finally discussing four distinct algorithms with respect to the approach.

This section and the remainder of the paper focus on pairwise correspondence between surfaces. It is important to remember that the framework provides a guide for computing initial correspondence estimates between views. Once initial correspondence estimates have been determined, multiview methods, such as [11] can be employed to register all the surfaces simultaneously, refining the initial estimates.

2.1 Terminology

Terminology that is used throughout the remainder of the paper is presented in Table 1. These terms have been chosen to consist with common usage elsewhere in image processing, for example [7].

Term	Definition
view	rendition of an object/scene
region	spatially consistent subset of a view
feature	intrinsic attribute of a region
feature representation	functions of accumulated features, that are comparable using distance metrics
anchor	location of a region, feature or feature representation on a view
local match	distance measure used to compare feature representations
group	set of local correspondences
group match	intrinsic metric used to compare group properties between views
global match	extrinsic metric used to evaluate registrations resulting from group mappings

Table 1. Terminology.

2.2 A General Approach for Correspondence

Many existing intrinsic methods compute correspondence using a similar generic approach. In this section, two surfaces X and Y are considered. The aim is to determine the correspondences between X and Y, and compute a rough initial registration between the two surfaces. A general approach followed by a number of intrinsic correspondence methods begins by specifying regions on both surfaces. These regions are computed using either a segmentation algorithm, or by selecting points of interest on the surfaces, and constructing regions around the points. Each region contains at least one anchor, describing its position on its parent view.

The next stage of the general approach is to extract features from each region, and represent the features using functions that are comparable using distance metrics. The local matching from one surface to the next then occurs by matching each representation on X with each representation on Y. The best local matches are then selected and local correspondences are formed by pairing anchors from X and Y that form the best local matches.

The final stage of the general approach is a global match procedure. This procedure generally incorporates a grouping process where sets of local correspondences are grouped and tested. If a group contains anchors on X that contain sufficiently similar properties as the anchors on Y, the group is used to transform the two surfaces into a common coordinate system. The surfaces are final tested using an extrinsic global match. If a number of groups are tested, the one that results in the best registration is selected as the set of correspondences between X and Y. The following section reviews existing methods that follow the general procedure outlined here.

2.3 Existing Intrinsic Methods

Four significantly different algorithms, selected from the intrinsic correspondence literature, are presented in this section. These algorithms represent a wide variety of intrinsic surface correspondence methods. This section shows how the algorithms relate to the approach discussed in the previous section. Examples of other algorithms that can be seen as following the same general approach are [6] and [12].

The four methods considered here are: spin-image [4], geometric histogram [1], Random Sample Consensus (RANSAC) based Data Aligned Rigidity Constrained Exhaustive Search (DARCES) [4] and intrinsic curve [9] matching. Spin-image [8] and geometric histogram [1] matching are similar techniques. Both methods build regions around interest points and use special 2D histogram feature representations to represent the regions. A set of pruned local correspondences is passed to each global match procedure. The match procedures vary, but both use a grouping and testing process.

The RANSAC-based DARCES method [4] does not extract feature information from regions. It initiates by computing point regions on one surface. The main component of the matching however occurs at the global match level, where collections of point regions on one surface are formed, and the second surface is searched for similar collections.

Intrinsic curve matching [9] extracts curvature information from regions, and then combines regions with similar curvature properties. Feature representations stem from these combined regions, and are matched from surface to surface. The global match procedure again relies on a grouping and testing scheme. This algorithm and the others outlined in this section are restructured to fit within the correspondence framework introduced next.

3 The Correspondence Framework

The correspondence framework is both a conceptual and an actual model that generalises the approach used by intrinsic surface correspondence outlined in the previous section. This section presents and details the five stages of the correspondence framework. It then re-casts four algorithms discussed in the previous section, to fit within the framework. A new algorithm is also presented, demonstrating how the framework facilitates the development of new correspondence methods.

3.1 A Model for Intrinsic Correspondence Computation

The correspondence framework consists of five stages: Region Definition, Feature Extraction, Feature Representation, Local Matching and Global Matching. Figure 1 illustrates the progression of a correspondence computation through these stages.

The function of each respective stage is:

Region Definition: The process of defining geometrically constrained regions on a surface. An example is constructing regions using neighbourhood information around anchors on surfaces [9].

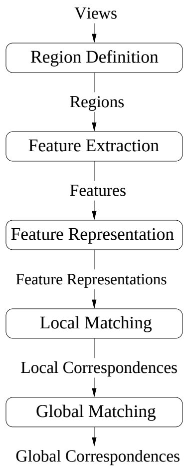

Figure 1. The correspondence framework.

Feature Extraction: The process of extracting geometric features inherent to each region. An example is extracting vertical and horizontal distances from the anchor to all surface points in the region [8].

Feature Representation: The process of representing the extracted features as functions that are comparable using distance metrics. An example is constructing special 2D histograms [1, 8].

Local Matching: The process of matching feature representations and culling bad matches. An example is comparing the Euclidean distance between intrinsic curvature signatures, and only accepting those signatures whose distance in difference fall below a preset threshold [9].

Global Matching: The process of grouping local correspondences, matching groups, and evaluating registrations derived from group anchors. An example is the RANSAC-based DARCES algorithm. This algorithm groups triplets of local match anchors based on the intrinsic distances between them. Registrations are applied to align the triplet pairs of the two surfaces. Finally, the registrations are evaluated by determining the number of closest points between the first surface and the registered second surface [4].

The following section examines the framework with respect to existing intrinsic methods.

3.2 Restructuring Existing Algorithms to Fit Within the Framework

The four algorithms outlined in Section 2.3 are revisited in this section. They are each modularised to fit within the correspondence framework. This modularisation demonstrates generality, implying that the framework is suitable for many intrinsic surface correspondence methods.

Figure 2 describes each stage of each algorithm. When modularised, the advantages and disadvantages of each algorithm outlined in Figure 2 are easier to identify. These are discussed below.

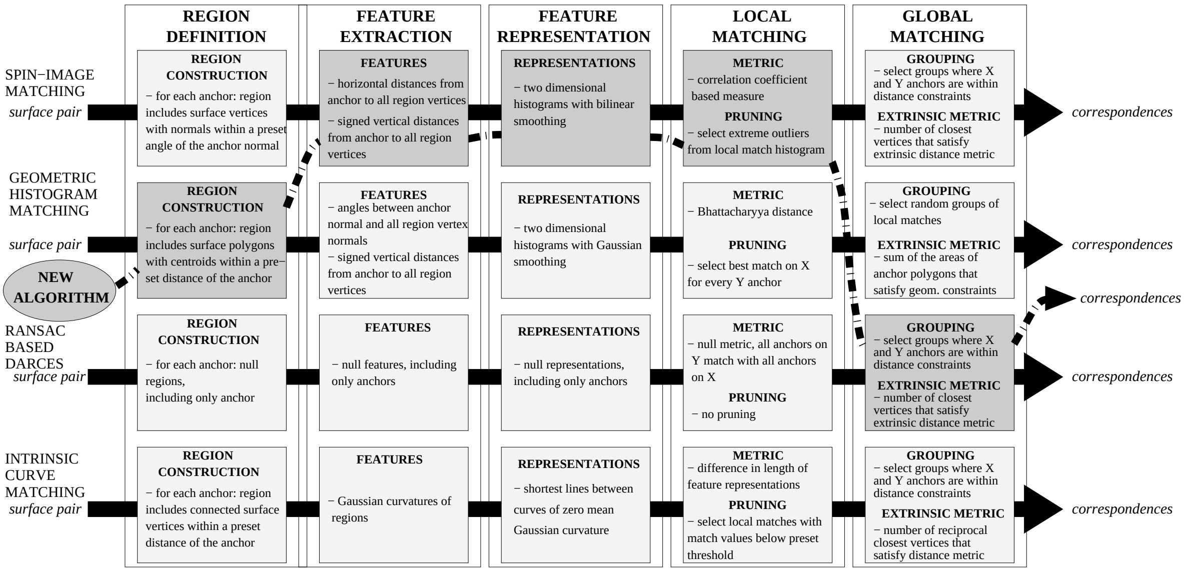

Figure 2. Existing intrinsic correspondence algorithms restructured to fit within the framework. The steps combined to construct the new correspondence algorithm are traced with the dotted line.

Spin-image matching [8] is a powerful algorithm using informative surface representations. It does however use large regions, which increase the storage space required to run the algorithm. The registration evaluation segment of the global matching process is also relatively slow, because it applies an iterative update scheme to each successful alignment.

Geometric histogram matching [1] uses smaller regions than spin-image matching, reducing the storage space required for the algorithm. The features extracted however, require that input surfaces are smooth, in order to retrieve accurate angle-between-normals data. The feature representations are 2D histograms, as in spin-image matching. The global match process selects random groups and tests their geometric consistency. This often requires a number of repetitions of the global match process, to achieve correct correspondences between two surfaces.

The computational effort in the RANSAC-based DARCES method [4] is reserved for the global match selection. The process relies on the random selection of the first point in each group, to build groups. Many iterations are required to achieve correct correspondence between the two surfaces. This process would operate much faster if fewer local correspondences were supplied to the algorithm.

Intrinsic curve matching [9] is relatively efficient. However, the feature representations are constructed from curves of zero mean Gaussian curvature. Surfaces without the curvature changes required to form the zero mean curvature curves, result in too few surface representations to achieve useful local and global matches.

3.3 Synthesising a New Algorithm

In the previous section, existing algorithms were decomposed into stages, to assist in the analysis of the algorithms. The best stages of each of these algorithms are now combined, to construct a new intrinsic correspondence method. The algorithm primarily combines the descriptive power of spin-image matching with the RANSAC-based DARCES global matching process. The steps used to construct the new algorithm are highlighted using the dotted and shaded path in Figure 2. Note that the local match path in Figure 2 does not show that both the spin-image matching and intrinsic curve matching pruning modules are combined to form the local match component of the new algorithm.

The new algorithm selects regions using the geometric histogram matching region selection process. This process ensures that regions are sufficiently large, but do not require as much storage space as spin-image regions. The feature extraction, feature representation and local match modules are all taken from the spin-image matching algorithm. As mentioned, spin-images generate powerful surface descriptions. The pruning section of the local correspondence algorithm uses the intrinsic curve pruning method, which reduces the number of best local matches passed to the global match module. The global match module then runs the RANSAC-based DARCES algorithm, which increases in speed dramatically, when using fewer local matches to achieve global correspondence.

The algorithm is presented below, with the sections incorporated from the four algorithms outlined labeled appropriately.

Region Definition: *geometric histogram matching*
regions defined by constricting the distance between the anchor and surrounding points
Feature Extraction: *spin-image matching*
features extracted are horizontal and signed vertical distances from anchor to surrounding region points
Feature Representation: *spin-image matching*
features are represented using 2D histograms
Local Matching: *spin-image matching*
representations are matched using a correlation-based coefficient
intrinsic curve matching
local matches with match values below a preset threshold are returned
Global Matching: *RANSAC-based DARCES*
groups are selected where the anchors on each surface fit the same distance constraints, and an extrinsic distance based measure is used to evaluate the registrations resulting from matched groups

The correspondence results achieved when using the modularised algorithms, including the new algorithm, are presented in the following section.

4 Results

Each of the algorithms outlined in the previous section was tested using three test surface pairs. The test data are Dinosaur and Angel surface pairs [3], and a Duck surface pair [10]. The test surfaces and their corresponding 2D images are shown in Figure 3. The test surfaces are low resolution surfaces, averaging approximately 1000 vertices per surface.

The algorithms outlined in the previous section were tested, and each returned the best correspondences achieved between each surface pair. These correspondences were subsequently used to align each pair in the same coordinate frame. These registrations are shown in Figure 4. Each registration is a rough initial alignment, which is all that is required of an intrinsic method.

Figure 4 shows that spin-image matching results in correctly aligned surfaces for all three pairs. It must be reiterated however, that this technique is inefficient in terms of storage, as large regions are used to create the spin-images. The geometric histogram also resulted in correct alignments for the Dinosaur and Angel surface pairs. A slight misalignment of the Duck surface pair occurred, possibly due to insufficient information extracted, to clearly distinguish between corresponding regions. The intrinsic curve matching algorithm correctly aligned the Angel pair, slight misaligned the Dinosaur pair, and failed to correctly register the Duck pair. These results were expected, because the algorithm retrieves a greater number of intrinsic curves for surfaces with more curvature variation. The Angel surfaces

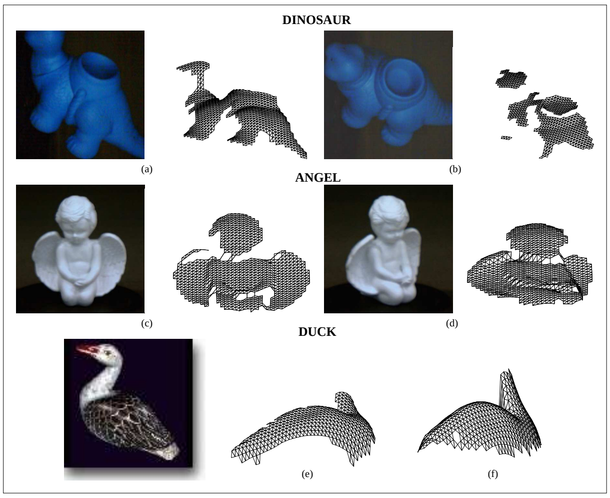

Figure 3. Surface pairs selected for testing: Dinosaur viewed at (a) 0 and (b) 36 degrees rotation, Angel viewed at (c) 0 and (d) 40 degrees rotation, and Duck viewed at two unspecified rotations (e) and (f).

produced more intrinsic curves then the Dinosaur or the Duck surfaces. The RANSAC-based DARCES correspondences correctly aligned all three surfaces. It did however, take longer to run then the other three algorithms.
The results of novel algorithm described in the previous section, are illustrated in the fifth row of Figure 4. All three surface pairs aligned correctly. The modules taken from the existing algorithms, which were combined to construct the new algorithm, performed as follows. The geometric histogram matching region selection module greatly reduced storage space required. The spin-image feature extraction, feature representation and local match modules applied the power of spin-image matching to the algorithm. The local match module also used the intrinsic curve matching culling process to reduce the number of local matches. Finally, the RANSAC-based DARCES global match module proved successful again, but ran much faster, as the number of local matches were supplied to it was greatly reduced.

The results in this section show that existing intrinsic algorithms perform as originally intended when re-cast in the proposed correspondence framework. They produce accurate results, and are easily combined to construct new algorithms. The new algorithm outlined shows how the framework facilitates the development of new and improved correspondence techniques.

5 Future Work

This paper focused on the correspondence framework for the purpose of rigid 3D surface correspondence. However, the correspondence framework is not limited to this domain.

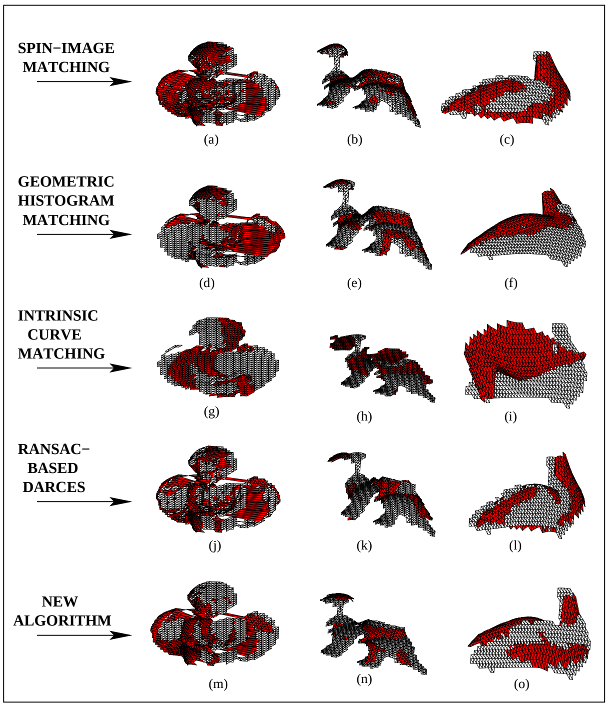

Figure 4. Registrations resulting from the best correspondences for the Angel (a,d,g,j,m), Dinosaur (b,e,h,k,n), and Duck surface pairs (c,f,i,l,o) using the algorithms indicated.

Correspondence algorithms for applications such as 2D image mosaicing can also be restructured to fit the framework.

A simple image mosaic algorithm has been implemented that first determines pairwise correspondence between images, before merging the set of images into a mosaic. Future work includes examining other domains, such as stereo matching, where the correspondence framework is applicable.

6 Conclusion

A correspondence framework was introduced in this paper, to provide a generic model for automatic surface correspondence. The framework was developed by examining the general approach to correspondence establishment, followed by many existing intrinsic rigid surface methods. The restructuring of algorithms, which were selected from a broad class of methods, to fit within the framework, demonstrated generality. Each algorithm produced correspondences that were the expected results for the method in question. A new algorithm was also developed and tested, showing how the framework facilitates the synthesis of new and improved correspondence methods. Finally, the expansion of the framework testbed was discussed, highlighting the framework's applicability to other open computer vision correspondence problems.

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

References

[1] 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.

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

[6] K. Higuchi, M. Hebert, and K. Ikeuchi. Building 3-D models from unregistered range images. *Graphic AI Models and Image Processing*, 57(4):315–333, July 1995.

[7] B. Jaehne. *Practical handbook on image processing for scientific applications.* CRC Press LLC, Florida, USA, 1997.

[8] 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.

[9] 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.

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

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

[12] J. V. Wyngaerd, L. V. Gool, R. Koch, and M. Proesmans. Invariant-based registration of surface patches. In *Proc. 7th IEEE International Conference on Computer Vision (ICCV '99)*, volume I, pages 301–306, Kerkyra, Greece, Sept. 1999.