Aligning Arbitrary Surfaces using Pairwise Geometric Histograms

A. P. Ashbrook, R. B. Fisher, N. Werghi and C. Robertson
Department of Artificial Intelligence, University of Edinburgh
Edinburgh, Scotland

Abstract

In this paper we present a novel representation for arbitrary surfaces that enables local correspondences to be determined. We then describe how these local correspondences can be used to search for the transformation that best aligns all of surface data. If this transformation is found to align a significant proportion of the surface data then the surfaces are said to have a correspondence.

1 Introduction

Finding a correspondence between two or more surfaces is a fundamental problem in many 3-dimensional vision and modelling applications. In this paper a novel representation for describing arbitrary surfaces is presented, enabling local correspondences between two or more surfaces to be determined. These local correspondences provide evidence towards the hypothesis that a global correspondence exists and enable an efficient search for the transformation that best aligns the data.

The representation is based on a shape descriptor previously applied to 2-dimensional shape representation problems [1] and recently extended to 3-dimensional surfaces [2]. The main contribution made by this paper is the adoption of a new, more efficient algorithm for determining global surface correspondences based upon local surface matches.

2 A Novel Surface Shape Representation

2.1 Surface Reconstruction and Approximation

Initially a given surface S, acquired using a range sensor, is described by a set of points samples $P = \{p_1, \ldots, p_N\}$. The points may represent a single view of the surface or a number of different views, for example from different viewpoints around an object. If a number of views are used then the data must be registered so that surfaces common to more than one view are aligned. The point set is then used to construct a triangular mesh approximation $\hat{S}$ to the original surface, where $\hat{S} = \{t_1, \ldots, t_M\}$ and t_i is a triangular facet of the mesh.

It is important to clarify at this stage that the only requirement of the mesh is that it is a good approximation of the surface shape. No assumptions are made about the actual placement of facets over the surface as this is unlikely to be repeatable.

In the work presented here an initial, regular mesh was constructed using an algorithm proposed by Hoppe *et al* [4] and then simplified using an algorithm proposed by Garland & Heckbert [5].

2.2 Histogram Construction

A pairwise geometric histogram h_i is constructed for each triangular facet t_i in a given mesh which describes its pairwise relationship with each of the other surrounding facets within a predefined distance. This distance controls the degree to which the representation is a local description of shape. The histogram encodes the surrounding shape geometry in a manner which is invariant to rigid transformations of the surface data and which is stable in the presence of surface clutter and missing surface data.

Figure 1(a) shows the measurements used to characterise the relationship between facet t_i and one of its neighbouring facets t_j. These measurements are the relative angle, α, between the facet normals and the range of perpendicular distances, d, from the plane in which facet t_i lies to all points on facet t_j. These measurements are accumulated in a 2-dimensional frequency histogram, weighted by the product of the areas of the two facets as shown in Figure 1(b). The weight of the entry is spread along the perpendicular distance axis in proportion to the area of the facet t_j at each distance.

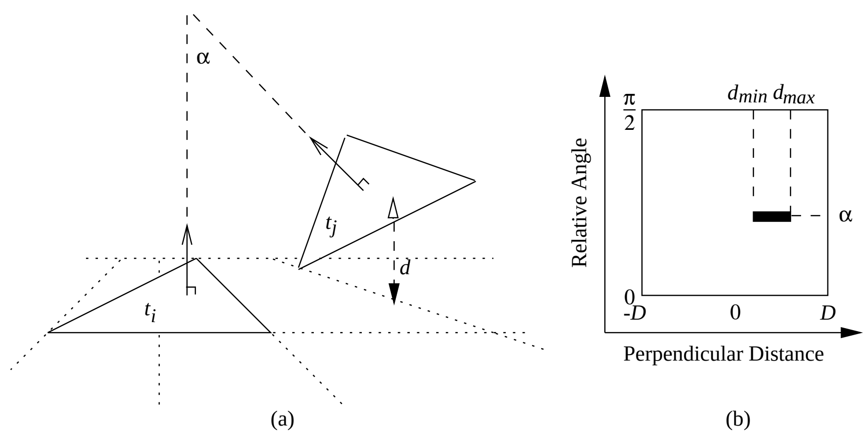

Figure 1: (a) The geometric measurements used to characterise the relationship between two facets t_i and t_j. (b) The entry made into the pairwise geometric histogram to represent this relationship.

To compensate for the difference between the measurements taken from the mesh and the true measurements for the original surface, each entry is convolved with an error function before being added to the histogram. Figure 2 presents the error in the relative angle and perpendicular distance measurements for a typical meshed surface. In practice, these error functions have been approximated with Gaussian distributions with a standard deviation of 10 degrees for the relative orientation and 0.5mm for the perpendicular distance.

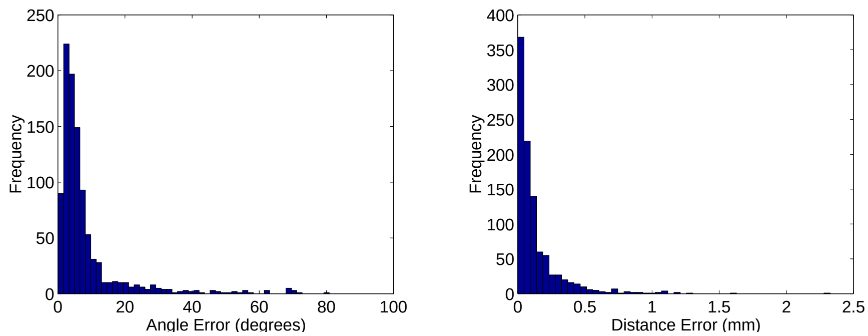

Figure 2: Relative orientation and perpendicular distance measurement errors for a typical meshed surface.

The complete pairwise geometric histogram for facet t_i is constructed by accumulating these entries for each of the neighbouring facets.

3 Classification of Scene Surface Features

Given two surface meshes, $\hat{S}^A$ and $\hat{S}^B$, the geometric histogram representation allows correspondences between all facets, t_i^A and t_j^B, from each of the meshes to be determined. A match for facet t_i^A is determined by finding the best match between its respective pairwise geometric histogram and all of the histograms representing the facets in surface $\hat{S}^B$. These *local* correspondences are treated as hypotheses for the correspondence between the two surfaces S^A and S^B.

The similarity, D_{ij}, between two pairwise geometric histograms h_i and h_j is defined using the Bhattacharyya metric [1, 2]. This is given by the expression:

$$D_{ij} = \sum_{\alpha,d} \sqrt{h_i(\alpha,d)h_j(\alpha,d)} \tag{1}$$

4 Estimating Surface Alignment

Good matches between surface facets provide evidence for the correspondence between the surfaces and provide constraints on the transformation that aligns them. To determine whether a pair of surfaces have a global correspondence these local correspondences are used to determine the transformation that best aligns all of the surface data. This is done here using a variant of the RANSAC (Random Sample Consensus) algorithm [3] which was developed for robust parameter estimation

To estimate the alignment transformation two passes of the RANSAC algorithm are used. In the first pass N_r pairs of surface patches are picked at random from the scene and these are used to generate N_r estimates of the rotation component of the alignment transformation. The amount of consistency associated with each estimate is determined by summing the area of matched surface facets which are consistent with the estimate. Matched surface facets are said to be consistent with the estimate if the direction of the surface normal of the aligned facets is within a specified degree of tolerance.

In the second pass of the RANSAC algorithm N_t triplets of surface facets are picked at random from the set of facets which were consistent with the best estimate in the first pass of the algorithm. An estimate of the translation that aligns the surfaces is then determined for each triplet and the amount of consistency is determined as before. In this case, matched surface facets are said to be consistent if the perpendicular distance between the aligned facets are within some tolerance.

The complete algorithm is repeated a fixed number of times. If the estimate with the maximum overall consistency is above a specified threshold then the surfaces are said to have a correspondence and the alignment transformation is improved by least squares fitting.

5 Demonstration

The results presented here demonstrate the effectiveness of using the proposed pairwise geometric histogram representation for finding the correspondence between scene surfaces and a set of model surfaces.

Figure 3(a) and 3(b) presents a pair of scenes containing a selection of objects. Each scene was generated by taking a single range image using a laser striper and then approximating the acquired surface points by a triangular faceted mesh. The first scene was approximated with 1000 triangular facets whilst 2000 facets were used to represent the second.

The set of model objects used as training data in this experiment are presented in Appendix A. To build each of the first three models enough range images were acquired to cover all of the surfaces. The range images for each object were then registered using the Iterated Closest Point algorithm [6] and a surface mesh of 1000 facets constructed. The remaining three models were each constructed from a pair of range images taken from different sides of the object and registered by hand. These surfaces were then approximated by 2000 facets each. Both scene and model surfaces were represented using geometric histograms with a resolution of 20x20 bins along the distance and relative orientation axes respectively. Pairwise measurements were constrained within a neighbourhood of 15mm.

Table 1 and Figure 3(c) and 3(d) present the object recognition and pose estimation results for each of the scenes. The table presents the percentage area of each scene which was found to be consistent with each of the six models, providing evidence for the presence of the models in each of the scenes. The

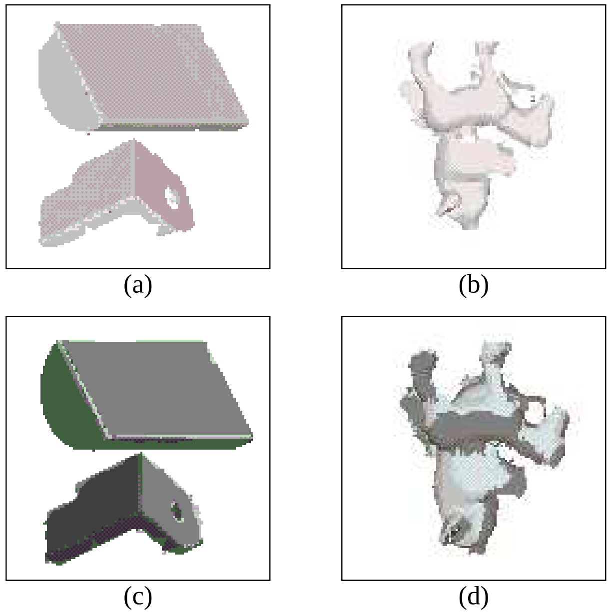

Figure 3: (a),(b) A pair of scenes. (c),(d) The estimated pose of the models.

	Cylinder	Block	Widget	Calf	Pig	Pony
Scene 1	48.5%	2.4%	36.6%	0%	1.1%	0%
Scene 2	0%	0%	0%	0%	19.4%	30.9

Table 1: The percentage area of each scene which was found to be consistent with each of the known models.

figure presents all of the detected models, in the lighter shade, superimposed over the scene data, in the darker shade, at the estimated poses. In all cases the models present have been detected successfully and the pose of each model determined. For each scene the RANSAC algorithm was run for 5000 trials to determine the best orientation of each model and then for 10000 trials to determine the best translation of each model.

6 Conclusions

In this paper a novel approach for representing 3-dimensional surface data using pairwise geometric histograms has been described. The representation allows local correspondences between pairs of arbitrary surfaces to be determined and these local matches may then be used to determine global surface correspondences. This has been demonstrated in a surface based object recognition application.

The representation inherits many of the advantages of the original pairwise geometric histogram descriptor [1]. By careful selection of the measurements used to construct the histogram, the descriptor is invariant to rigid transformations of the surface data and in combination with its compactness promotes efficient matching. A reasonable criticism would be the large number of histograms needed to describe a particular surface, between 1000 and 2000 for the surfaces used here. This problem can be minimised by developing improved algorithms for segmenting surfaces into a small number of planar surface facets.

In this paper the search for global surfaces correspondences has been conducted using a variant on RANSAC in contrast to the Probabilistic Hough Transform used previously [2]. The random element of RANSAC means that there is a small chance of not finding the best alignment but this is minimised by running for a large number of trials. The algorithm finds the solution significantly faster however.

A Database of 3D Shape Models

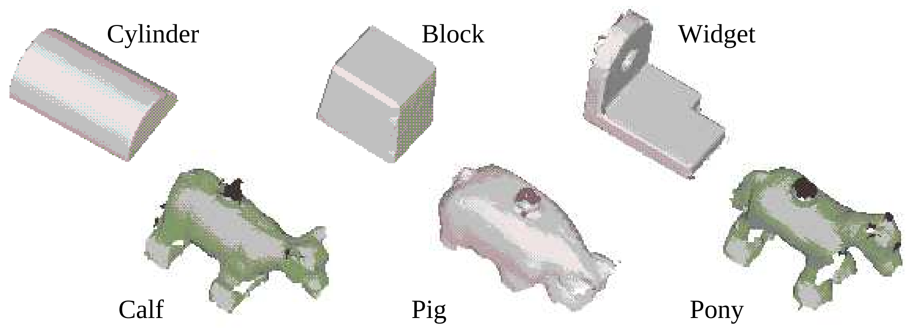

Figure 4: 3-Dimensional surface models.

References

[1] Evans, A. C., "Geometric Feature Distributions for Shape Representation and Recognition", Submitted for Ph.D. at the University of Sheffield, 1994.

[2] Ashbrook, A. P., Fisher, R. B., Robertson, C. and Werghi, N., "Finding Surface Correspondence for Object Recognition and Registration using Pairwise Geometric Histograms", To appear at ECCV98, Freiburg, Germany.

[3] Fischler, M. A. and Bolles, R. C., "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography", Communications of the ACM, Vol. 24, pp381, 1981.

[4] Hoppe, H., DeRose, T., Duchamp, T., McDonald, J. and Stuetzle, W., "Surface Reconstruction from Unorganised Points", Computer Graphics, 26(2), pp71-78, 1992.

[5] Garland, M. and Heckbert, P. S., "Surface Simplification using Quadric Error Metrics", SIGGRAPH97, pp209-216, 1997.

[6] Besl, P. J. and McKay, N. D., "A method for registration of 3-D shapes", IEEE PAMI, 14(2), pp 239-256, 1992.