

3-D Map Reconstruction from Range Data

Daniel Huber, Owen Carmichael, Martial Hebert
{dhuber, owenc, hebert}@ri.cmu.edu

The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Abstract

We present techniques for building models of complex environments from range data gathered at multiple viewpoints. The challenges in this problem are the matching of unregistered views without prior knowledge of pose, the use of very large data sets, and the manipulation of data sets of different resolutions and from different sensors. Our approach is unique in that no prior knowledge of the relative viewpoints is needed in order to register the data. We show results in building maps of interior environment from range finder data, building large terrain maps from ground-based and from aerial data, and from an operational for mapping from stereo data for hazardous environment characterization. The paper summarizes the major results obtained so far in this area. Details of the algorithms can be found in earlier papers as referenced in the text.

1.0 Introduction

The problem of building models from multiple views is critical in various applications, including remote operation, virtual environment building, and construction of object libraries for recognition. In this paper, we consider the problem of registering range data sets from multiple locations in order to build a complete model of an environment. A typical scenario involves many viewpoints over a large area with poor or nonexistent initial estimation of the relative positions of the views. Eventually, our goal is to build maps that cover hundreds of square meters at very high (e.g., sub-centimeter) resolution.

This problem has many facets, including computing the transformations between views, correcting the transformations to compensate for drift and error accumulation, merging the views into a single model once the transformations are computed, and incorporating color and texture information into the final model. Here, we focus on the first problem: reliable registration of large point sets from range data. Examples of view merging are shown only to illustrate the results of registration.

The method described here is based on a general approach

to surface matching [13], which has been applied successfully to object recognition [14] and three-dimensional object modeling. This method has two critical features which address the shortcomings of previous techniques. First, it does not require surface segmentation or feature extraction. Second, it does not require knowledge of the transformation between surfaces prior to registration. Owing to these two features, this surface matching algorithm can be applied to problems of large-scale model building.

The goal of this paper is to show how our matching algorithm can be used in the context of mapping large, cluttered environments. The paper summarizes the major results obtained in this field. More technical details can be found in earlier publications [2][11].

2.0 Point Matching Using Local Signatures

In [9], Johnson introduces the spin-image, a two-dimensional signature describing the local shape of a free-form three-dimensional surface at a point p on that surface. Spin-images encode the positions of points near p in terms of distance along and distance from the approximated normal to the surface at p . By comparing this coordinate information from the spin-images of two different points, we arrive at a measure of local shape similarity between the surfaces surrounding those two points.

To construct the spin-image for an arbitrary point p , we first find the best-fit plane to the nearest neighbors of p and approximate the normal to p as the normal to this plane. We then define a 2-D basis using the normal n and the plane \mathcal{P} perpendicular to n and passing through p . For each point x in the vicinity of p we compute its coordinates (α, β) with respect to this basis; α is the distance from p to x measured in the plane \mathcal{P} , while β is the perpendicular distance from x to \mathcal{P} (Figure 1.)

The (α, β) values are then discretized and accumulated into a 2-D array of bins called a spin-image; each bin in the spin-image corresponds to some range of α and β values.




Figure 1: The signature at point p is computed by recording the distance of all nearby points x from the surface normal n (α) and the distance from x to p along n (β). Corresponding points from different views have similar signatures.

These spin-images are compact signatures describing the local shape of a surface around a particular point; if two points have similar signatures, they are considered to have similar local shape (Figure 1).




Figure 2: Signatures from points on surface meshes from two views. Corresponding points have similar signatures.

Spin-image similarity is determined by simple linear correlation of bin values, supplemented with a confidence metric to take into account the number of empty bins in. Bins containing no points are not considered in the correlation calculation to help minimize the effects of clutter

and occlusion on shape similarity by only considering bins in the spin images that have been filled. The final similarity measure between two points takes into account the number of bins used to compute correlation, so that signatures with the highest amounts of overlap are considered the most similar.

Johnson shows in [9] that through determining the similarity between spin-images of points from two different surfaces, it is possible to recover the transformation that registers the surfaces even though no initial estimate of this transformation is known at the time of registration.

3.0 View Registration

The matching approach described above provides a powerful tool for building large-scale models from sensor scans taken at unspecified locations. In this section we describe the process of using spin-images to register range data taken from a variety of locations. We assume that regardless of the sensor used, range data from it consists of 3-D points in some fixed coordinate frame. Given these points, a triangular mesh is formed by connecting nearest neighbors, and noisy points and edges in this mesh are removed through cleaning and smoothing operations. At this point, we have a high-resolution 3-D representation of the model as seen from a particular viewpoint; for reasons to be explained below, we then simplify this surface using the algorithm found in [12] to obtain a low-resolution version of the surface.

Given two low-resolution meshes, we register them using signature matching as follows. First, a fixed fraction of points is selected at from both surfaces. The signatures are produced for these points, and all the signatures from one surface are compared to all the signatures from the other using the similarity measure described above. When a pair of signatures are found to have high similarity, the points that produced them are considered to correspond to each other (Figure 2); after all signature comparisons have taken place, we are left with a set of point matches between the two surfaces. Finally, an estimate of the transformation that aligns the two surfaces is computed from the point matches.

Based on this estimate an ICP algorithm is applied to the approximately-aligned high-density surfaces. The resulting transformation aligns the full range data sets with acceptable accuracy.

One feature of this procedure is that the only step in the process which depends on the user's choice of range sensor is the acquisition of range data; the manipulation of meshes, registration of surfaces, ICP, and merging are completely independent of the range sensor used. In fact, this technique has been used with eleven different range

sensors to data; examples with three sensors are included at the end of the paper.

4.0 Surface Matching with Large Data Sets

The basic procedure outlined above is sufficient for data sets of moderate size, for example, individual objects. When dealing with data sets that cover a large area, two issues must be addressed carefully. First, the large variation of data resolution across the area scanned complicates surface matching. Second, the size of the data may render the matching procedure computationally impractical.

4.1 Variable Data Resolution

In the representation described thus far, the signature images are computed by histogramming the vertices of the model or scene meshes. As a result, the distribution of the vertices on the mesh directly affects the signatures. In fact, two meshes with different vertex distributions may generate very different signatures at the same basis point. Therefore, in order for the surface matching algorithm to work properly, some constraint has to be enforced on the distribution of vertices on the meshes. Specifically, it can be shown that the signatures remain stable as long as the vertices are uniformly distributed on the surface.

Although this approach works well in practice, it has several problems. First of all, there are cases in which the data simply cannot be made uniform without losing a great deal of information because the variation of resolution in the input sensor data is too large. A typical example is terrain data taken from a forward-looking sensor. The sensor data varies from high-resolution at close range to quadratically decreasing resolution as the range from the sensor increases. Variations in data point spacing of as much as 1:10 are routinely observed on terrain data. The second problem is that the uniform decimation requires on the same order of computation time as the matching itself, even though much faster decimation and filtering algorithms do exist.

The solution to those problems is to compute the signatures by integrating over the entire surface rather than by computing α and β values at the vertices only. Essentially, this requires interpolating the spin-image values "in between" the mesh vertices. The simplest way of achieving this is to raster scan each triangle of the mesh (Figure 3) and to compute the (α, β) coordinates of each point inside the face. The corresponding location in the signature is incremented by a constant amount for each new point. This algorithm can be made efficient by using a fast geometric test in order to determine whether a face is inside the region of influence of the basis point and is within the boundary of the spin-image space.

This approach is still an approximation because it uses a discrete sampling of the surface. In particular, although the signatures are less sensitive to the distribution of vertices, they are still sensitive to the choice of the sampling rate used for interpolation.




Figure 3: Discrete vertex interpolation: each face is raster-scanned before mapping to spin-image. The sample points used for the scanning are shown in green in the neighborhood of the basis point shown in red.

The second approach is exact in that it computes the spin-images by integration over the whole surface without additional sampling. In this approach, the boundary of each triangle is mapped into $\alpha\beta$ -space, as shown in Figure 4 (a). Each edge of the face maps to a segment of hyperbola. The hyperbolic segments computed in the projection are then used for determining which cells of the spin-image may contain some portion of the triangle. Figure 4 (b) shows the portion of the spin-image that contains a portion of the triangle based on the segments of Figure 4 (a). Finally, each cell in the spin-image is incremented by the area of the part of the triangular face that is mapped to that cell in $\alpha\beta$ -space.




Figure 4: Continuous interpolation: each face is mapped to $\alpha\beta$ space by mapping its edges (a); cells inside the mapped region are incremented by the area of the intersection of the face and the volume in space corresponding to the cell (b).

This last step is illustrated geometrically in Figure 4 (b). The region of 3D corresponding to the cell (α, β) is an annulus of height $\Delta\beta$ and thickness $\Delta\alpha$. The cell is incremented by the surface area of the intersection between the triangle and this annulus.

Because it uses the actual surface area for incrementing

the signature cells, this algorithm computes an “exact” mapping of the surface to the signatures, given a mesh discretization of the surface. Figure 5 shows the difference between a signature computed with and without interpolation.




Figure 5: The vertex-based signature for the basis point of Figure 3 (left) is degraded when compared with the corresponding face-based signature (right).

4.2 Fast Access Data Structures

This part of the work addresses the practical use of the matching techniques, in particular using the more advanced surface integration, for the very large data sets that one expects to encounter in applications such as building terrain models, or virtual models of large interior environments.

The main potential obstacle to the practical use of signature techniques is that the computation of the signatures may become prohibitively expensive if the data set is very large. In particular, the amount of computation needed from each face in order to compute the signature at a given basis point is much more substantial than in the standard method. Therefore, it becomes especially important that only the points that are inside the region of influence of a basis point be used for computing the corresponding signature.

The standard approach to this type of problem is the use of variants of the K-D tree structure designed for fast access in multidimensional spaces. After evaluation of several implementations of similar geometric data structures, the best design turned out to be a regular hierarchical data which is similar to octrees, except that, because we are working with 2-D surfaces, the tree is sparse and access can be efficiently implemented by a fast hashing method. The graphs below illustrate the improvement in signature computation speed obtained using this technique.

Because of the overhead involved in computing the data structure, and in computing the hashing function and retrieving points from the data structure, this technique is really beneficial only for large data sets. In fact, the computation is slower for data size of moderate size. The graphs show that the crossover point is at approximately 8000 points (Figure 6.) This technique should not be used for smaller data sets.




Figure 6: Computation time as a function of the number of points without indexing data structure (red), and with indexing data structure (green).

5.0 Indoor Environment Reconstruction

As a first example, the surface registration technique has been used for building a model of a large warehouse space composed of two adjacent rooms. The building measured roughly 60 meters long by 20 meters wide by 10 meters high, and was filled with an assortment of clutter and debris as shown in Figure 7. We used a K2T/Z+F laser

View A:



View B:



Figure 7: Two reflectance images of different sections of the warehouse area. Only part of the panoramic scan is shown for each image.

range finder [5][6] mounted atop a minivan to collect range data at 32 different points in the warehouse; each scan of the sensor covered a 360-degree field of view with 20 degree depression and a maximum range of about 40 meters. The resolution of the sensor is rated in the millimeter region.

Each scan of the range finder returned 1.8 million 3-D points; through naive subsampling this set was reduced to 65000 points. The resulting point cloud was converted into a mesh and resampled to roughly 5000 points for registration.

Surface meshes produced from each of the 32 data sets were successfully registered to the meshes of adjacent

scans. Figure 8 shows the registration results for two of the meshes. With the transformations between neighboring




Figure 8: (Top) Registered range data acquired from two separate sensor locations. One data set is displayed in black wireframe, the other in solid grey. (Bottom) Boxed section enlarged to show detail.

meshes known, it was then possible to align all meshes in a common coordinate system so that they could be unified into a single mesh that covered the entire warehouse space. Because of limitations in our implementation of the merging algorithm, only 20 of the views were merged into the final model, although all 32 data sets were registered. The complete mesh of the warehouse contained 138000 points at full resolution, while a low-resolution version of the model contained 25000 points. The resolution of the final mesh was low compared to that of the constituent high-resolution views due to limitations of the mesh merging implementation; otherwise, it would have been possible to produce a final mesh with no resolution loss. Statistics for each of the different types of surface mesh-- high-resolution single-viewpoint, low-resolution single-viewpoint, high-resolution complete, and low-resolution complete-- are summarized in Table 1. Resolution is measured as the average length of edges in the mesh, in millimeters.

TABLE 1. Mesh Statistics

Mesh Type	# Points	# Edges	Resolution
High Res.	65000	160000	70
Low Res.	5000	12000	275
Complete High Res.	138000	400000	150
Complete Low Res.	25000	65000	320

Figure 9 shows a bird's-eye view of the final model, while Figure 10 shows the positions of the sensor at viewpoints A, B, and C from Figure 4 along with views of the final model from these vantage points. As mentioned above, we did not attempt to refine the final merged model; more sophisticated merging algorithms can certainly be used

once the transformations are computed.




Figure 9: Bird's-eye view of the building model.




Figure 10: Views of the final model from viewpoints A and B of Figure 4. The wireframe are shown using the low-resolution data.

6.0 Large-Scale Terrain Mapping

A second example of map reconstruction is map building from a ground-based range sensor and from an autonomous helicopter. The ground-based sensor is a K2T/Z+F sensor [5] which was configured to generate range and reflectance images of size 6000 x 300, corresponding to angular resolutions of 0.06° and 0.1° in the horizontal and vertical directions respectively. The initial data is subsampled by a factor of 5 horizontally and 3 vertically and then converted to a mesh.

The helicopter range data was preprocessed by the CMU Autonomous Helicopter group, but we will briefly describe the steps used. The combination of the line scanning sensor, helicopter motion, and position uncertainty causes sequential scan lines to overlap unpredictably. Therefore it is not possible to simply connect the data points as with the image based scanner. Instead, the points

from a series of scans are projected orthographically into a grid, and the range is calculated using the closest point in each bin. Then the bins are treated as a range image and converted to a mesh using the same method as with the ground-based sensor.

The registration algorithms were tested on several large data sets, two of which are shown here. The first data set was acquired with the sensor mounted on the roof of a van at a slag heap located near CMU. The data is a sequence of 13 range images obtained at 3-5 meter intervals along a road leading between two hills. After the 7th data set, the road curves to the left. The registered data sets were merged into a single, global map shown in Figure 8.

The second data set was collected by the Autonomous Helicopter group during field tests at Haughton crater in the Canadian arctic as part of NASA's Haughton-Mars project. The integrated terrain map in Figure 11 was formed from three passes of the helicopter along the boundary of a 20-meter cliff and covers a 260 x 166-meter area.

7.0 Conclusion

Building environment models from unregistered views is challenging because of the difficulty in extracting and matching surfaces, the size of the data sets, and the variation in resolution and accuracy among views and sensors. The approach presented here addresses those challenges

and was successfully demonstrated on large 3-D data sets. Beyond the two examples presented here, the system has been exercised with eleven different range sensors, from active laser range finders to passive stereo systems.

Now that we have a reliable algorithm for building maps, we are beginning to analyze the limitations of the approach. The next step is to investigate how terrain shape affects the algorithm's performance. Two areas of investigation are of particular interest: intelligent selection of points and dynamic estimation of overlap between views.




Figure 11: Top-view integrated map built using data from the autonomous helicopter (top) and perspective views of the map (bottom). Grid lines are 5 meters apart.

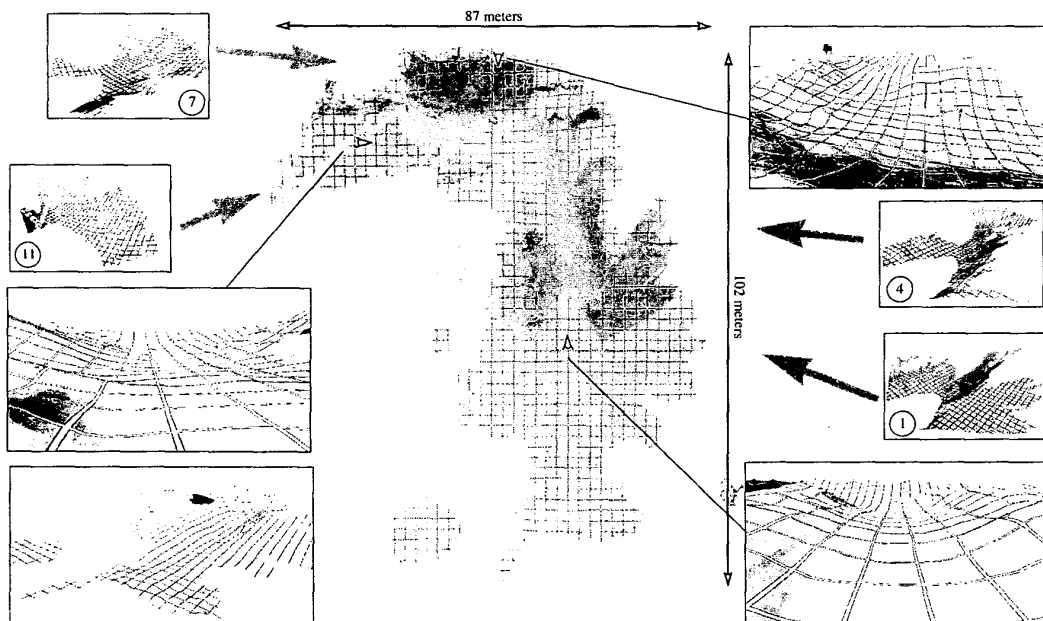


Figure 8: Top view of the integrated terrain map for the eleven data sets in the mesa sequence (center). The numbered insets illustrate individual data sets, and two sets (1 and 11) are overlaid on the top-view. The larger insets show various perspective views of the map, and the white arrows indicate the location and direction of the viewpoint. Grid lines are two-meters apart on the combined map and one meter apart on the individual data sets.

Acknowledgments

Work on the matching and registration algorithms was conducted by Andrew Johnson, Owen Carmichael, and Daniel Huber. Ryan Miller and John Hancock provided valuable assistance in collecting and preprocessing the laser range finder data.

References

- [1] P. Besl and N. McKay. A method of registration of 3-D shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 14, no. 2, pp. 239-256, February, 1992.
- [2] O. Carmichael and M. Hebert. Unconstrained registration of large 3D point sets for complex model building. *IEEE/RSJ International Conference On Intelligent Robotic Systems*. Vancouver, September 1998.
- [3] Y. Chen and G. Medioni. Surface description of complex objects from range images. *Proc. IEEE Computer Vision and Pattern Recognition (CVPR '94)*, pp. 153-158, 1994.
- [4] C. Chua and R. Jarvis. 3-D free-form surface registration and object recognition. *Int'l Jour. Computer Vision*, vol. 17, no. 1, pp. 77-99, 1996.
- [5] Froelich, C., M. Mettenleiter, F. Haertl. Imaging laser radar (LIDAR) for high-speed monitoring of the environment. *SPIE Proceedings of the Intelligent Transportation Systems Conference*, October 1997.
- [6] Hancock, J., E. Hoffman, R. Sullivan, D. Ingimarson, D. Langer, M. Hebert. High-performance laser range scanner. *SPIE Proceedings of the Intelligent Transportation Systems Conference*, October 1997.
- [7] M. Garland and P. Heckbert. Surface simplification using quadric error metrics, in *Proc. SIGGRAPH 97*, pp. 209-16, 1997.
- [8] D. Gennery. Visual terrain matching for a Mars rover, in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR '89)*, pp. 483-91, 1989.
- [9] J. Hancock, D. Langer, M. Hebert, R. Sullivan, D. Ingimarson, E. Hoffman, M. Mettenleitner, C. Froehlich. Active Laser Radar for High Performance Measurements, in *Proc. IEEE International Conference on Robotics and Automation*, pp. 1465-70, 1998.
- [10] B. K. P. Horn and J. G. Harris. Rigid body motion from range image sequences, *CVGIP: Image Understanding*, vol.53, no.1, pp. 1-13, January, 1991.
- [11] D.F. Huber and M. Hebert. A New approach to 3-D terrain mapping. *Proc. of the 1999 IEEE/RSJ International Conference on Intelligent Robotics and Systems*. 1999.
- [12] A. Johnson and M. Hebert. Surface matching for object recognition in complex three-dimensional scenes, *Image and Vision Computing*, vol.16, no.9-10, p. 635-51
- [13] A. Johnson and M. Hebert. Using Spin-Images for Efficient Object Recognition in Cluttered Scenes. *IEEE Trans. Pattern Analysis and Machine Intelligence*. Vol. 10 No. 5. May 99.
- [14] A. Johnson and M. Hebert. *Control of Polygonal Mesh Resolution for 3-D Computer Vision*, in *Graphics Modeling and Computer Vision*. 1998.
- [15] A. Johnson and S.B. Kang. Registration and integration of textured 3-D data. *International Conference on Recent Advances in 3-D Digital Imaging and Modeling (3DIM '97)*, pp. 234-241, Ottawa, Ontario, May 12-15, 1997.
- [16] I. S. Kweon and T. Kanade. High-resolution terrain map from multiple sensor data, *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 278-292, February, 1992.
- [17] R. Miller and O. Amidi. 3-D site mapping with the CMU autonomous helicopter, in *Proc. Intl. Conf. on Intelligent Autonomous Systems (IAS-5)*, June, 1998.
- [18] C. Olson and L. Matthies. Maximum likelihood rover localization by matching range maps. In *Proc. IEEE International Conference on Robotics and Automation*, pp. 272-277, 1998.
- [19] H.Y. Shum, R. Reddy, M. Hebert, K. Ikeuchi. An integral approach to free-form object modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. December 1997.
- [20] G. Turk and M. Levoy. Zippered polygonal meshes from range images. *Proc. Computer Graphics (SIGGRAPH '94)*, pp. 311-318, 1994.
- [21] M. Wheeler. Automatic Modeling and Localization for Object Recognition. Ph.D. Thesis, School of Computer Science, Carnegie Mellon University, October 1996.
- [22] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *Int'l Jour. Computer Vision*, vol. 13, no. 2, pp. 119-152, 1994.