

Surface Reconstruction from Feature Based Stereo

Camillo J. Taylor
GRASP Laboratory
University of Pennsylvania
Philadelphia, PA, 19104

Abstract

This paper describes an approach to recovering surface models of complex scenes from the quasi-sparse data returned by a feature based stereo system. The method can be used to merge stereo results obtained from different viewpoints into a single coherent surface mesh. The technique proceeds by exploiting the freespace theorem which provides a principled mechanism for reasoning about the structure of the scene based on quasi-sparse correspondences in multiple image. Effective methods for overcoming the difficulties posed by missing features and outliers are discussed. Results obtained by applying this approach to actual images are presented.

1 Introduction

This paper addresses the problem of recovering representations for the surface structure of 3D scenes such as the one shown in figure 1 from a set of images. The proposed scheme takes as input the results of applying a feature based, multi-baseline stereo algorithm to clusters of images taken from various vantage points around the scene. Based on this information, the system constructs a mesh representation of the surface like the one shown in Figure 1b. Importantly, the approach is able to handle scenes that involve multiple disconnected surfaces which may occlude each other in interesting ways as opposed to scenes that can effectively be modeled as 2.5D surfaces.




Figure 1: a. Image of a typical scene b. Surface mesh recovered by the proposed procedure.

When accurate dense range scans are available, several

excellent techniques are available for merging this information into coherent surfaces [12, 15, 2, 7, 16]. Unfortunately, there are situations in which it can be very difficult to obtain the dense, accurate depth maps required by these techniques solely from image data.

Passive image based ranging techniques like stereo proceed by correlating pixels or features between images. In any non trivial scene we can expect to encounter regions which cannot be readily matched between images. This can occur either because certain portions of the structure are occluded in one or more of the viewpoints under consideration or because the images contain homogenous regions which admit multiple possible matches. In either case, dense stereo systems are obliged to deal with these problems by appealing to a priori expectations about scene structure such as smoothness or continuity rather than direct measurements.

It would seem, then, that there is a certain mismatch between the character of the data returned by most stereo systems and the accuracy requirements of the previously mentioned surface reconstruction techniques. This paper argues that, given the nature of the stereo matching process, it may be more appropriate to view each stereo reconstruction as providing information about the *freespace* in the scene rather than measurements of surface structure. In this framework, the stereo results obtained from various vantage points can be combined by considering the union of the freespace volumes they induce. This procedure can be made to work even when the individual reconstructions contain interpolation artifacts.

The proposed scheme proceeds from the position that our multi-baseline stereo system can be expected to find accurate correspondences with high confidence in areas where there is sufficient intensity variation in the image. These regions typically correspond to intensity edges or textured regions in the scene. Although the stereo system may only obtain depth information for a small fraction of the pixels in the input images, these results can be used to reason about the scene freespace in a principled manner and, hence, to construct an appropriate 3D model.

The model of the stereo process used in this work is somewhat extreme since it is usually possible to produce

reasonable disparity values for a fairly large fraction of the pixel locations. The point of choosing such an impoverished model is to demonstrate that it is still possible to derive useful information about the structure of the scene from a small number of high confidence correspondences. When more correspondences are available they will only serve to improve the quality of the reconstructions.

Several techniques have been proposed to construct 3D models from the sparse 3D data returned by feature based structure from motion algorithms. Faugeras, Lebras-Mehlman and Boissonnat [4] describe a technique for building 3D models by triangulating a set of 3D point features and eliminating surfaces which contradict feature visibility constraints. Manassis et al [9] proposed a related method for triangulating sparse feature sets. Morris and Kanade [11] describe a scheme that searches for reasonable triangulations of a point set by considering the implied motion of textures in the image. The method described in this paper differs from these algorithms in that it generates a surface description by considering the freespace volumes associated with the stereo results instead of tessellating the recovered 3D points.

Several effective techniques for recovering the structure of a scene based on the voxel carving algorithm proposed by Kutulakos and Seitz [8] have been proposed. The method proposed in this paper differs from these schemes in two important ways. Firstly it does not make use of the viewpoint consistency constraint directly, it makes determinations about the structure of the scene based on the 3D results returned by the stereo system. Secondly the proposed approach does not involve discretizing the scene into voxels.

An interesting and effective technique for recovering surface models from multiple images was proposed by Faugeras and Keriven [5]. This work reformulates the stereo reconstruction process as a surface evolution problem and solves it using variational principles. In contrast, the approach taken in this paper directly recovers scene occupancy from the available disparity measurements.

Another important class of techniques for recovering surface models from image data draw inspiration from the work of Baumgart [1]. These methods work by considering the intersection of the volumes obtained by extruding the object silhouette into the working volume. Variants of this scheme have been proposed by Sullivan and Ponce [17], Wong and Cipolla [18] and Matusik et al [10]. The method proposed in this paper is in some sense a dual to these techniques in that it considers the union of the freespace volumes associated with each image as opposed to the intersection of the volumes associated with each silhouette. It is also important to note that the technique does not involve the use of distinguished contours like silhouettes which can be difficult to obtain.

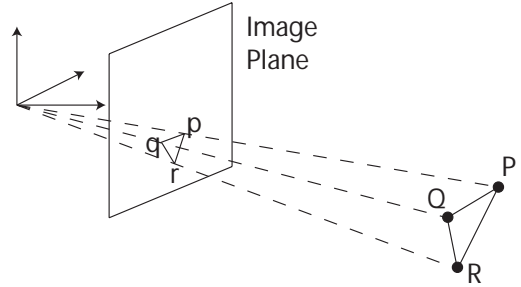


Figure 2: If the space triangle $\triangle PQR$ was occluded by another surface in the scene then the corresponding triangle in the image, $\triangle pqr$, would contain points corresponding to the boundary of the occluding surface which would preclude the triangle $\triangle pqr$ from the Delaunay triangulation of the image.

2 Reasoning About Scene Structure

Figure 8 shows the typical input provided to the surface reconstruction algorithm by the feature based stereo algorithm. In this case the results were obtained from a set of five images taken from slightly different vantage points. Note that the stereo system only returns depth information in regions of the image corresponding to intensity discontinuities. Since this particular scene contains numerous homogenous regions, the fraction of pixels with depth values is very small. The disparities at other locations in the image can be estimated by linearly interpolating the disparity values of the recovered features based on a Delaunay triangulation of their projections in the image. Typically, the vast majority of the triangles produced by this scheme will correspond to actual surface regions in the scene. However, at occlusion boundaries the scheme will naively produce triangles that connect foreground surfaces to background surfaces. While it is true that these interpolated disparity maps may not yield accurate information about the surfaces in the scene they do yield reliable information about the freespace in the scene. The following theorem states this observation more concisely:

Freespace Theorem : Suppose that three recovered space points P , Q , and R project to the pixel locations p , q , and r respectively in a given input image, and suppose further that $\triangle pqr$ is one of the Delaunay triangles formed in that image. Then the tetrahedron formed by the camera center and the space triangle $\triangle PQR$ must consist entirely of free space, i.e. no surface in the scene intersects this tetrahedron.

Proof: Assume, to the contrary, that there is a surface within this tetrahedron. Then the surface cannot occlude any of the points P , Q , or R (or else the occluded point would not be visible in the image). Therefore the boundary of this surface must lie at least partially inside the tetrahedron. Points from the boundary would then show up as

edges inside the image triangle $\triangle pqr$ (see Figure 2). This contradicts our assumption that p , q , and r are part of a single Delaunay triangle, since the interiors of Delaunay triangles cannot contain other vertices¹. Hence, the tetrahedron must consist entirely of free space. ♣²

Note that this argument rests on two assumptions. Firstly it assumes that we have a procedure that accurately reconstructs the depths of all of the salient edges visible from a given vantage point. If the reconstruction system only returned the depths of isolated corner features in the image, this property would not hold. Secondly the argument presumes that the surface in question can be adequately represented by a polygonal mesh where the vertices correspond to intensity discontinuities on the surface. Note that this does not preclude the reconstruction of curved surfaces so long as there are sufficient surface markings to provide vertices for an approximating mesh. It would, however, pose a problem in the case of curved objects with no surface markings. Consider, for example, the problem of applying a feature based stereo algorithm to images of a featureless cylinder. In this case, the only structural features that could be recovered would correspond to points along the occluding edges of the cylinder. Triangulations that simply connected these points together in the images would provide inappropriate approximations of the object. In the sequel we will describe how both of these assumptions can be relaxed to account for problems that may occur in practice.

The practical consequence of the freespace theorem is that we can think of each interpolated disparity map as defining a star shaped region of freespace and we can reason about the 3-D structure of the scene by reasoning about the union of these freespace volumes as shown in Figure 3.

Given the coordinates of a point in space, P , we can easily test whether that point is in the union of the freespace volumes by projecting the point into each of the original images and determining whether the depth of the point P with respect to the image center is less than the depth of the corresponding entry in the interpolated depth map for that image. This operation can be performed in constant time by dividing each image into rectangular regions and caching references to the triangles that intersect those regions.

This freespace procedure can be represented by an indicator function $\Phi(P) : R^3 \rightarrow [0, 1]$ which returns 0 whenever the point P lies within the freespace volume and 1 otherwise. This function can be thought of as an *implicit* representation of the total freespace volume.

Given this freespace indicator function $\Phi(P) : R^3 \rightarrow [0, 1]$, it is relatively straightforward to construct an explicit

¹In fact, the defining property of Delaunay triangulations is that the interior of the *circumcircle* of every triangle does not contain any other vertices; however, we do not require such a strong property for our argument.

²This argument has appeared in a previous publication it is repeated here for ease of reference [6].

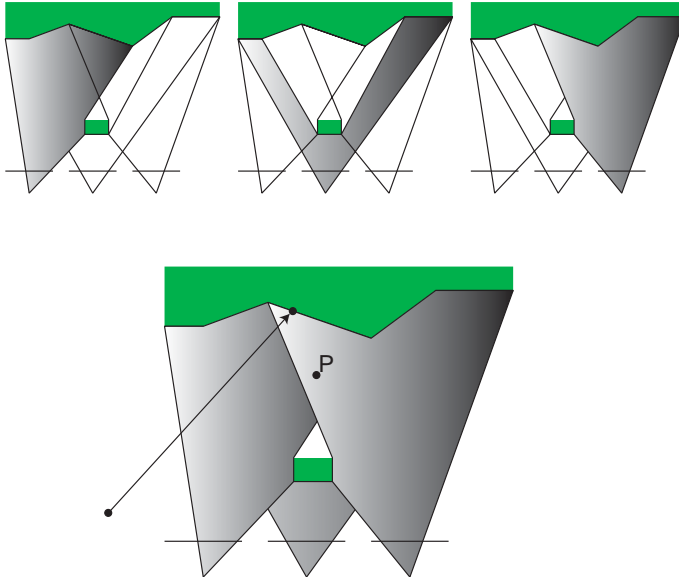


Figure 3: The first row of figures depicts the freespace volumes associated with the triangulations of each of the 3 input images. The union of these volumes provides a more accurate approximation of the 3D structure of the scene than any of the original depth maps. It is a simple matter to construct a function $\Phi(P)$ which indicates whether a particular point P lies within the union of the freespace volumes.

representation for the freespace boundary in the form of a triangular mesh by invoking an isosurface algorithm such as Marching Cubes [14]. The boundary of the freespace volume can be thought of as a fair approximation for the surface geometry in the sense that it will correspond to a surface which is consistent with all of the available information. Notice, however, that the actual surface may contain regions that are not visible from any of the camera positions. The freespace union method will still produce reasonable results in these situations.

2.1 Handling Missing Features

The procedure described in the previous subsection will work perfectly so long as the feature based stereo algorithm can be relied upon to produce depth values for all of the salient edges in the scene. Unfortunately, there are cases in practice where the feature based stereo may fail in this regard. Problems can arise in situations where a feature is occluded in one or more of the images in the stereo cluster or when a salient feature does not have sufficient contrast to produce a reliable correspondence.

Figure 4 shows a situation in which the stereo results obtained from stereo cluster 1 do not reflect the presence of the point P . This means that the freespace volume constructed

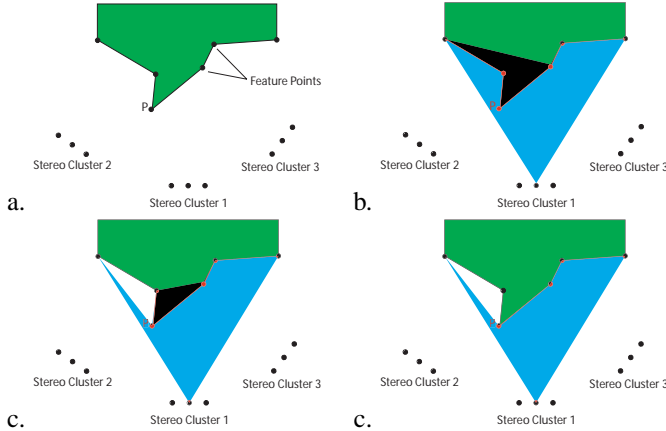


Figure 4: (a) This figure shows a situation where the feature labeled P is not recovered by stereo cluster 1. (b) Without feature point P, the freespace volume obtained by triangulating the remaining points will eliminate a significant portion of the surface (c) Since the point P was among the points reconstructed by stereo cluster 2 one could consider combining the two stereo results by simply adding all of the feature points that appear inside the freespace volume to the mix and recomputing the triangulation for stereo cluster 1. Unfortunately this policy is too liberal and the resulting surface can still have problems as shown here. (d) These problems can be overcome by adding the points to the triangulation incrementally based on their height above the current triangulated surface.

for that cluster will eliminate a significant portion of the structure. This is clearly undesirable.

In practice, it is usually the case that multiple stereo clusters are used to reconstruct a given scene and these are typically arranged with a significant region of overlap. This means that features that may be missing from one stereo reconstruction are often recovered by a neighboring stereo system. In the example shown in Figure 4 we assume that the point P was recovered by stereo cluster 2. Since the reconstruction system is aware of this point, P, it can determine that its initial estimate for the freespace volume associated with cluster 1 is flawed since it includes a known feature point, P. This observation that the freespace volume cannot contain any of the recovered feature points is referred to as the *minimum disparity constraint* since it effectively constrains the freespace volumes associated with each vantage point.

The most straightforward approach to fixing the freespace volume associated with stereo cluster 1 so that it satisfies this constraint would be to add all the feature points that violated the constraint to the mix and recompute the triangulation associated with stereo cluster 1. Figure 4c shows the results of such a process. The resulting freespace

volume contains all of the known feature points but it still includes regions of the object that should not have been included.

This problem can be solved by modifying the update procedure to add features to the triangulation in order based on their height above the current estimate for the freespace surface. On each iteration the update procedure would project each of the recovered scene points into the viewpoint of cluster 1 and determine its height above or below the current triangulated disparity surface associated with that viewpoint. The point with the greatest height above the disparity surface would be used to update the Delaunay triangulation incrementally. The process would continue until all of the feature points were safely contained within the triangulated surface associated with the cluster. Using this procedure the system would update its original flawed estimate for the freespace surface by adding point P and retriangulating. At this point it would determine that there were no further illegal points and terminate.

This procedure for updating the triangulation to include all of the recovered points can be implemented efficiently by modifying a standard incremental Delaunay triangulation algorithm such as the one described in [3]. Typically, only a small fraction of the points need to be added to fix the triangulation. Furthermore, the process of finding the point with the greatest height above the triangulation can be optimized very effectively by caching references to the Delaunay triangles that enclose the projected feature points. The worst case computational complexity of this update procedure would be $n^2 \log n$ where n denotes the total number of feature points, this corresponds to the case where all n features need to be added to the triangulation.

This update procedure also provides an avenue for dealing with featureless curved surfaces such as the one considered in section 2. Figure 5 shows the case where multiple stereo clusters are viewing a featureless cylinder, each of the stereo clusters locates two feature points corresponding to the occluding contour visible from that vantage point. By applying the minimum disparity constraint described previously, the reconstruction system would recover the polygonal estimate for the curved surface shown in Figure 5b. This estimate contains all of the recovered feature points. Clearly the fidelity of such an approximation would depend upon the number of stereo clusters used to reconstruct the scene and their distribution in space. Nonetheless, even in this worst case scenario it is possible to recover some information about the scene structure from the sparse data provided by a feature based stereo system.

2.2 Recovering Surfaces

Once the minimum disparity constraint has been applied to update the freespace volumes associated with each of the

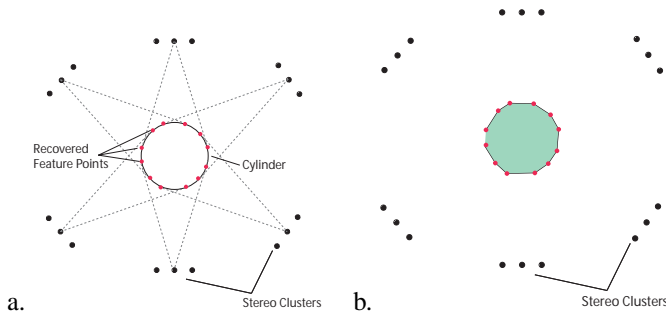


Figure 5: In this worst case scenario we consider a set of stereo clusters viewing a featureless cylinder. a. Each of the stereo clusters reconstructs a pair of feature points corresponding to points on the occluding contour of the cylinder. b. The approximation constructed based on these features by applying the minimum disparity constraint.

stereo clusters, the only remaining task is to recover an explicit, triangulated representation for the boundary of the freespace defined by the implicit function $\Phi(P)$. Actually, the reconstruction procedure recovers the boundary of a slightly different indicator function, $\Phi'(P)$ which evaluates to 1 iff P is inside the freespace defined by $\Phi(P)$ or outside the convex hull of the points. This modification reflects the fact that the reconstructed surface should not extend beyond the convex hull of the recovered feature points. It is a very simple matter to determine whether a point, P , lies within the convex hull of our point set. One need only evaluate a set of linear equations which correspond to the facets of the convex hull.

The boundary of this function, $\Phi'(P)$ can be obtained by invoking a modified version of the well known marching cubes algorithm³. Recall that the marching cubes procedure finds isosurfaces by sampling the given function at a set of grid points in the specified volume (see Figure 6). When the algorithm detects a transition in the function value between two neighboring grid points it inserts a vertex along that edge in the grid. In this situation it is possible to refine the location of the transition along the segment between the two grid points by applying standard bisection search [13] to the freespace indicator function $\Phi'(P)$. This makes it possible to use a fairly large sampling distance between the grid points while still accurately determining the location of surface vertices.

One issue to be mindful of is that isolated outliers can have a significant impact on the recovered surface since in certain situations. Problems with outliers can typically be overcome by insisting that the recovered features have a requisite amount of support in a local neighborhood of the image.

³The same arguments can also be applied to modify the marching tetrahedra algorithm

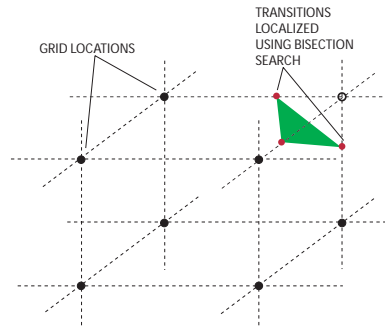


Figure 6: The marching cubes algorithm proceeds by sampling grid points in the specified volume in search of transitions in the given implicit function. Once a transition has been detected it can be further localized by applying standard bisection search techniques.




Figure 7: Two of the 30 images acquired of this scene

The marching cubes procedure also provides a certain amount of robustness to isolated outliers since these types of features usually induce thin spikes in the freespace volume. These thin volumes typically do not contain many grid points and, hence, will not induce many surface vertices. This amounts to a form of “structural aliasing” which works to our advantage in this case.

3 Experimental Results

In order to test the effectiveness of the proposed reconstruction procedure, a series of image data sets was obtained. Each data set consisted of 30 images divided into 6 stereo clusters. The stereo clusters were positioned at evenly spaced intervals all around the target scene. The calibration grid visible in the images shown in Figure 7 was used to determine the intrinsic and extrinsic parameters associated with each of the images in the data set.

Each stereo cluster consisted of 5 images taken in roughly the pattern of a cross as shown in Figure 9. The images were taken with a Foveon SD9 Digital SLR camera outfitted with a 50mm lens. For each stereo cluster two trinocular stereo procedures were performed. The first procedure used the three images corresponding to the horizontal axis of the cross while the second used the three images corresponding to the vertical axis of the cross. The cen-

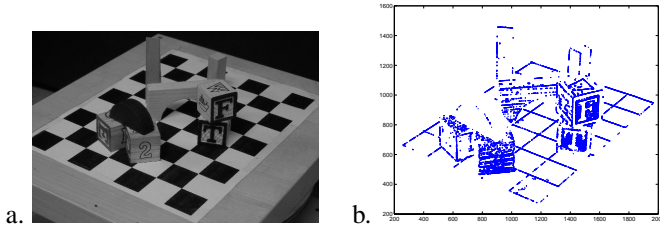


Figure 8: This figure shows typical results obtained from our feature based stereo algorithm on this scene. a. The center image of one of the stereo clusters b. Features recovered in that stereo cluster

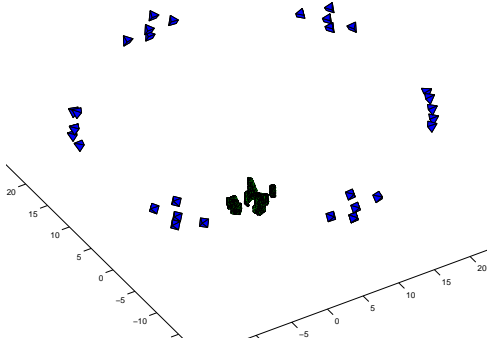


Figure 9: Figure showing the arrangement of viewpoints used in this experiment. The images are divided into 6 stereo clusters each of which consists of five views taken in a cross configuration.

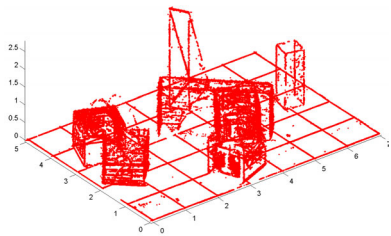


Figure 10: Novel view of the set of feature points recovered by the stereo clusters.

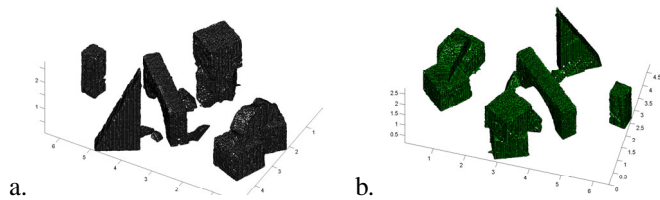


Figure 11: Various novel views of the 3D mesh obtained with the reconstruction algorithm.

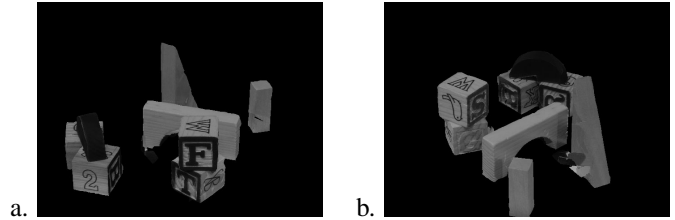


Figure 12: Various novel views of the 3D mesh rendered using a view dependent texture mapping technique

ter image was used twice and served as the base frame of reference for the stereo cluster. This scheme was able to reconstruct edges in the scene at all orientations.

The trinocular stereo algorithm employed in this work rectifies the three images to a common orientation, extracts edge tokens in all of the images, and matches these tokens across the views. Image windows on either side of the extracted edge features are examined in order to verify proposed correspondences. Significantly, this approach is able to deal with occlusion boundaries in the scene which are often mishandled by correlation based stereo methods.

Figure 8 shows an example of the feature points recovered from a stereo cluster projected into the viewpoint of the central image of that cluster. Figure 10 shows all of the recovered feature points viewed from a novel vantage points. Feature points were typically recovered with millimeter precision.

The results of applying the surface reconstruction procedure are shown in Figure 11. They demonstrate that the method is able to recover reasonable approximations for the surface structure of a fairly complex 3 dimensional scene. In this case, some of the volume between the blocks was not visible from any of the vantage points used in the reconstruction. As a result the surface reconstruction procedure produced some “extra” surface elements corresponding to this volume. This problem could be fixed by adding stereo results from additional vantage points.

4 Conclusions

This paper presents a novel scheme for reconstructing the 3D structure of a scene based on the information returned by a feature based stereo algorithm. The approach exploits the fact that even though feature based stereo methods may only return disparity values for a small fraction of the pixels in the input images, the features that they do return tend to correspond to salient structural features in the environment. This allows us to reason about the structure of the scene in a principled manner.

The entire reconstruction procedure is summarized below:

1. Acquire imagery
2. Run feature based stereo on each cluster
3. Construct a triangulated disparity map for each stereo cluster taking into account the minimum disparity constraint
4. Recover an estimate for the surface structure by applying marching cubes to the union of the freespace volumes and the complement of the convex hull of the feature points.

This work deliberately considers a fairly extreme form of the reconstruction problem, the case where the stereo system can only provide information about the depth of edge features in the input images. The intent has been to show that it is still possible to determine quite a lot about the structure of the scene in this case. In many cases, it is reasonable to expect that a stereo system would be able to provide much more information about the structure of the scene and these additional correspondences will only serve to improve the quality of the reconstruction.

It is also interesting to note that for some applications like obstacle avoidance and ray tracing the implicit representation for scene freespace, $\Phi(P)$, is very useful since it provides us with a direct method for answering questions about the occupancy structure of the scene.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-98-75867. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [1] B.G. Baumgart. *Geometric modeling for computer vision*. PhD thesis, Stanford, 1974.
- [2] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of SIGGRAPH 96. In Computer Graphics Proceedings, Annual Conference Series*, pages 31–43, New Orleans, LA, August 4-9 1996. ACM SIGGRAPH.
- [3] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry*. Springer, 2 edition, 1997.
- [4] O. Faugeras, E Lebras-Mehlman, and Jean-Daniel Boissonnat. Representing stereo data with the delaunay triangulation. *Artificial Intelligence Journal*, 44(1-2):41–87, July 1990.
- [5] Olivier Faugeras and Renaud Keriven. Variational principles, surface evolution, pde's, level-set methods, and the stereo problem. *IEEE Trans. Image Processing*, 7(3):336–344, 1998.
- [6] David Jelinek and Camillo J. Taylor. View synthesis with occlusion reasoning using quasi-sparse feature correspondences. In *European Conference on Computer Vision*, volume 2, pages 463–492, May 2002.
- [7] T. Kanade, P.W. Rander, and J. P. Narayanan. Virtualized reality: Constructing virtual worlds from real scenes. *IEEE Multimedia*, 4(1):34–47, 1997.
- [8] K.N. Kutulakos and S. Seitz. A theory of shape by space carving. *International Journal of Computer Vision*, 38:199–218, 2000.
- [9] A. Manassis, A. Hilton, P. Palmer, P. McLauchlan, and X. Shen. Reconstruction of scene models from sparse 3d structure. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, volume 2, pages 666–671, 1999.
- [10] Wojciech Matusik, Christopher Buehler, Ramesh Raskarand Leonard McMillan, and Steven J. Gortler. Image-based visual hulls. In *SIGGRAPH*, 2000.
- [11] D.D. Morris and T. Kanade. Image-consistent surface triangulation. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, volume 1, pages 332–338, 2000.
- [12] Marsahll Bern Nina Amenta and Manolis Kamvyselis. A new voronoi-based surface reconstruction algorithm. In *SIGGRAPH*, pages 415–421, 1998.
- [13] W. Press, B.Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes in C*. Cambridge University Press, 1988.
- [14] J.A. Sethian. *Level Set Methods; Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision and Material Sciences*. Cambridge University Press, 1996.
- [15] F. Bernardini J. Mittleman H. Rushmeier C. Silva and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 5(4):349–359, October/December 1999.
- [16] Ioannis Stamos and Peter Allen. 3-d model construction using range and image data. In *IEEE Conference on Computer Vision and Pattern Recogniton*, 2000.
- [17] Steve Sullivan and Jean Ponce. Automatic model construction and pose estimation from photographs using triangular splines. *IEEE Trans. Pattern Anal. Machine Intell.*, 20(10):1091–1097, 1998.
- [18] K.-Y. K. Wong, , and R. Cipolla. Structure and motion from silhouettes. In *8th IEEE International Conference on Computer Vision*, volume 2, pages 217–222, Vancouver, Canada, July 2001.