Multi-View Reconstruction using Narrow-Band Graph-Cuts and Surface Normal Optimization

Alexander Ladikos Selim Benhimane Nassir Navab

Chair for Computer Aided Medical Procedures
Department of Informatics
Technische Universität München
Boltzmannstr. 3, 85748 Garching, Germany
ladikos@in.tum.de

Abstract

This paper presents a new algorithm for reducing the minimal surface bias associated with volumetric graph cuts for 3D reconstruction from multiple calibrated images. The algorithm is based on an iterative graph-cut over narrow bands combined with an accurate surface normal estimation. At each iteration, we first optimize the normal to each surface patch in order to obtain a precise value for the photometric consistency measure. This helps in preserving narrow protrusions with high curvature which are very sensitive to the choice of normal. We then apply a volumetric graph-cut on a narrow band around the current surface estimate to determine the optimal surface inside this band. Using graph cuts on a narrow band allows us to avoid local minima inside the band while at the same time reducing the danger of taking "shortcuts" and converging to a wrong "global" minimum when using a wide band. Reconstruction results obtained on standard data sets clearly show the merits of the proposed algorithm.

1 Introduction

Reconstructing the shape of an object given a set of calibrated input images is a topic which has been extensively studied in the computer vision community as recently surveyed in [21]. A first class of methods reconstruct the visual hull of the object [26, 17, 5, 18]. These methods exclusively use silhouette information; consequently, they do not recover the concavities of the shape. A better shape reconstruction can be obtained using voxel-coloring or space carving which take the photometric consistency of the surface across the input images into account and allow the recovery of the photo-hull that contains all possible photo-consistent reconstructions [22, 16, 15, 25]. This approach works well in general; unfortunately, if a voxel is wrongly carved, it cannot be restored in later iterations. This can lead to the carving of the whole neighboring region (if not the whole volume). In addition, since space-carving is a greedy approach, it is hard to enforce smoothness constraints on the reconstruction. Another class of methods optimizes the surface integral of a consistency function over the surface shape. One way of minimizing this cost function is the variational/level set formulation [4, 11, 20]. In this formulation

BMVC 2008 doi:10.5244/C.22.15

the surface is iteratively deformed using a gradient descent method. It is also possible to add regularizers to the cost function, such as smoothness constraints. However, being a local method it can fall into a local minimum. A second way of minimizing the surface integral is to use graph cuts [1, 24, 29, 2, 6, 13, 19, 27, 31, 12, 23]. One problem is that the cost function which is minimized is a minimum surface functional and hence the global solution is biased towards smaller shapes. Therefore, the global minimum might not correspond to the actual surface. This is also true for level set methods, but due to the local convergence properties of level sets, the effect is not as strong as in the case of graph-cuts. In practice this leads to the carving of narrow protrusions, since the graph-cut solution prefers shorter cuts over long cuts. This problem has been addressed by incorporating silhouette constraints [24, 27, 6, 23] or adding a ballooning term [29, 12]. However, using silhouette constraints is only viable when exact silhouettes are available and a global ballooning term has the side effect of also pushing out concave regions of the object. In addition to the optimization method, the photometric consistency measure and its accurate computation also play an important role for the quality of the reconstruction results. In fact, an incorrect estimation may lead to overcarving the volume and eliminating protrusions.

In this paper, we propose a method to reduce the bias of graph-cuts for smaller cuts without requiring exact silhouette images or using a ballooning term. To this end, we use iterative volumetric graph-cuts over narrow bands to minimize the influence of shortcuts on the reconstruction coupled with an exact normal optimization for each surface patch which is used to compute an accurate photoconsistency score. It is through this combination that we are able to avoid the overcarving of narrow protrusions. Using each technique by itself does not avoid the overcarving. Using only narrow bands still allows overcarving when the photoconsistency score is incorrectly estimated, while using an exact photoconsistency score in a wide band also has a high probability of overcarving, because of the high variation in possible path lengths (see figure 1(a)). At each iteration, we first compute the visibility of the current surface estimate and optimize the normals to the surface in order to obtain a precise value for the photometric consistency measure. Then, we apply a volumetric graph-cut in order to determine the optimal surface inside a narrow band around the current surface estimate. The band size can be adjusted to achieve a tradeoff between the ability to overcome local minima and the ability to preserve protrusions by discouraging shortcuts. We show that iteratively searching for the correct orientation of the surface and considering narrow bands is able to deal with the graph-cut's inherent bias for smaller shapes without requiring a ballooning term or exact silhouettes. Thereby our algorithm reduces problems with overcarving and preserves protrusions of the surface.

2 Related Work

The use of graph-cuts on narrow bands was proposed by Xu *et al.*[30] in the context of image segmentation. Hornung *et al.*[13] suggested the use of hierarchical iterated graph-cuts for 3D reconstruction. However, using graph-cuts on narrow bands alone does not necessarily preserve protrusions. Another important aspect is the exact computation of the surface consistency score. This is typically performed by computing the NCC of a small patch in the tangent plane of the surface over several images. A critical aspect in this computation is the choice of the normal. The current surface estimate given by

the graph-cut is not well-suited to determining accurate normals to the surface due to the discretization. In addition, the current surface estimate can be quite different from the actual surface. Therefore it is not sufficient to only evaluate the consistency measure using the normal given by the current surface. Instead it has to be optimized in order to find the most consistent normal over all input images in which the surface point is visible. This helps significantly in determining correct consistency values even for high curvature regions which cannot be well represented on the voxel grid. The concept of optimizing the surface normal appeared in other work related to 3D reconstruction [8, 7, 9, 10, 32]. Habbecke and Kobbelt [9, 10] use a normal optimization in order to find the orientation of 3D patches which are combined to approximate the surface. A similar idea is used by Furukawa *et al.* [7]. Both of these assume a reference patch with respect to which the normal is optimized. This is in contrast to our work which does not assume a reference view. None of these consider the use of normal optimization in conjunction with narrow band graph-cuts.

Prior work related to reducing the minimum surface bias in graph-cut reconstructions uses either a ballooning term [29, 2, 28, 12] or includes silhouette constraints [27, 6, 23]. The use of a global ballooning term as suggested in [29] has the drawback of also affecting concavities. It is often not possible to find one single value which achieves the desired effect of preserving protrusions and not affecting concavities. Hernandez *et al.* [12] use an intelligent ballooning term which is based on the evidence of regions being inside or outside the volume. This yields improved results over using a global ballooning term. Boykov *et al.* [2] use the photoflux as an intelligent ballooning term to drive the reconstruction towards the object boundaries, thereby allowing the recovery of thin structures. The results they achieve on the gargoyle data set are comparable to ours. However, we achieve this by using the normal optimization which does not require to add extra links to the graph for incorporating the intelligent ballooning term. The use of silhouette constraints is only viable when exact silhouettes are available. This is not always the case. In addition, the silhouettes only constrain the shape of the surface on the rims. The method proposed by Yu *et al.* [31] uses graph-cuts on surface distance grids. The drawback of their method is that they assume, that the initial estimate is already quite close to the final result.

3 Volumetric Graph-Cuts

3.1 Consistency Measure

One of the most common measures for evaluating the consistency of a surface patch in the input images is the NCC due to its invariance to linear illumination changes. Similar to [6, 23], we compute the NCC for a surface patch over the image projection of its tangent plane. The tangent plane is sampled with a uniform grid in 3D and each point is projected into the images in which the surface patch is visible. Since the NCC is only defined for image pairs, it has to be extended to multiple images in order to get one consistency score. One way of doing this is to select a reference image and to compute the NCC between this image and all other images. Using this approach, the reliability of the score depends very much on the reference image: If the reference image is of bad quality or if it contains occlusions which have not been captured by the model, the correlation will be low. Therefore, we choose to compute the mean of the pairwise score of all image pairs.

If the NCC score between images i and j of the voxel $\mathbf{x}$ is $\eta_{ij}(\mathbf{x})$, the photo-consistency score is given by

$$\rho(\mathbf{x}) = \Psi\left(\frac{2}{n(n-1)}\sum_i\sum_{j>i} a_{ij}\eta_{ij}(\mathbf{x})\right) \tag{1}$$

where $\Psi(x) = \min(1-x, 1)$ normalizes the score to the interval $[0,1]$. n is the number of images and a_{ij} is a weight depending on the angle between the two cameras. Since this consistency measure is based on the tangent plane it is important to obtain a good normal to the surface. A wrong normal can result in a low consistency score which in turn will lead to overcarving. Therefore, in our algorithm, we use the normal which results in the highest photo-consistency score as described in section 4.1.

In our implementation, we sample the tangent plane with a 5×5 grid. The spacing between the grid points is chosen so that no pixels in the image are missed. We only use views forming an angle of less than 60 degrees with the surface normal. The weights a_{ij} are chosen as the cosine of the angle between the viewing directions of camera i and j thereby penalizing big angles.

3.2 Surface Optimization with Graph-Cuts

3D reconstruction can be cast as an optimization problem to find the minimum cost surface where the cost of the surface S is modeled as the surface integral over a consistency score $\rho(\mathbf{x})$ for each surface patch $\mathbf{x}$

$$E = \iint_S \rho(\mathbf{x}) dA \tag{2}$$

Volumetric graph-cuts provide a way to find an approximately optimal solution to this minimization problem. This is done by converting the continuous problem into a discrete formulation over a voxel grid [1].

Given an initial estimate of the surface, a band is constructed around it and the best surface inside this band is found using a graph-cut [3, 14]. For every voxel $\mathbf{x}$ in the band the consistency score $\rho(\mathbf{x})$ is computed, where a lower cost signifies a better consistency. In order to be able to compute the consistency for a voxel, it is necessary to know the visibility and depending on the consistency measure also the normal of the voxel. This information is propagated from the original surface by assigning the visibility and the normal from the closest voxel in the original surface. Alternatively, it can also be estimated for each layer in the band using the surface given by this layer.

Every voxel in the band is represented as a node in the graph. The source of the graph is placed inside the object while the sink is placed outside. All voxels on the inner boundary of the band are connected to the source and all voxels on the outer boundary are connected to the sink by edges of infinite weight. Neighboring voxels $\mathbf{x}_k$ and $\mathbf{x}_l$ are connected by edges of weight

$$w_{kl} = c_{kl}\frac{\rho(\mathbf{x}_k) + \rho(\mathbf{x}_l)}{2} \tag{3}$$

where the term c_{kl} is a weight proportional to the distance between the two voxels [1]. Typical neighborhood systems are the 6- and the 26-neighborhood. All voxels still connected to the source after applying the graph-cut are part of the optimal volume.

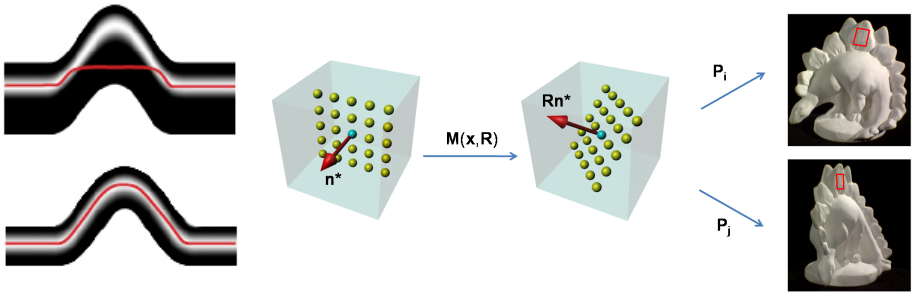

(a) (b)

Figure 1: Figure (a) shows the advantages of using a narrow band. When using wide bands protrusions can be totally filled out. The graph-cut will then prefer the shortcut through the high-cost region, because its accumulated cost is less than that of the long path through the low-cost region. When using a narrow band, the graph-cut will take the low-cost path because there is no shortcut and any deviation from the low-cost path will incur a high cost. Figure (b) shows the photo-consistency computation. We find the plane which best approximates the surface passing through a voxel.

In our implementation, we use the 26-neighborhood of the voxel instead of the 6-neighborhood when constructing the graph. This helps to avoid some of the discretization artifacts associated with graph cuts.

4 Proposed Reconstruction Method

We use an iterative graph-cut approach to recover the shape of the object. In each iteration, first the surface visibility is computed. Using this information the normals to the surface are optimized and used to compute a reliable consistency score. These scores are used in a volumetric graph-cut over narrow bands around the current surface estimate.

4.1 Surface Normal Optimization

As explained in section 3.1, we need the normal to the surface to compute the consistency score. The accuracy of the normal plays an important role in obtaining good reconstruction results. If the used normals are inaccurate the computed consistency score will be low which in turn will lead to overcarving. This is especially true for surface regions with high curvature, because a small change in the normal can have a big effect on the consistency score. We therefore find the optimal normals which lead to the highest consistency score.

Since we use a volumetric representation, we first compute the gradient in the volume for all surface voxels, giving us the initial normals. However, the normals obtained this way are usually not very accurate due to the discretization and because the current surface estimate can be wrong. In particular, there are problems in surface regions with high curvature. Therefore, we only use these estimated normals as a starting point for an

optimization over the surface orientation. The optimization allows us to obtain a better normal estimate and consequently a better consistency measure.

We represent the orientation of the surface through the rotation $\mathbf{R}$ of a reference plane with normal $\mathbf{n}^* = [0,0,1]^\top$ (see figure 1(b)). The rotation is represented using a (3×3) rotation matrix $\mathbf{R}$. The consistency score for a voxel $\mathbf{x}$ between images i and j depending on the rotation $\mathbf{R}$ is given by

$$\eta_{ij}(\mathbf{x},\mathbf{R}) = NCC\left(I_i(\mathbf{P}_i\mathbf{M}(\mathbf{x},\mathbf{R})\mathbf{X}), I_j(\mathbf{P}_j\mathbf{M}(\mathbf{x},\mathbf{R})\mathbf{X})\right) \tag{4}$$

where $\mathbf{P}_i$ and $\mathbf{P}_j$ are the projection matrices for images I_i and I_j respectively. $\mathbf{X}$ are 3D points (expressed in homogeneous coordinates) defined on the reference plane. The matrix $\mathbf{M}(\mathbf{x},\mathbf{R})$ is a (4×4) matrix that describes the transformation of these points given the orientation $\mathbf{R}$ and the voxel center $\mathbf{x}$. It is defined as

$$\mathbf{M}(\mathbf{x},\mathbf{R}) = \begin{bmatrix} \mathbf{R} & (\mathbf{I}-\mathbf{R})\mathbf{x} \\ \mathbf{0} & 1 \end{bmatrix} \tag{5}$$

The optimal normal and the highest consistency measure are found by determining the optimum of the cost function

$$\rho(\mathbf{x}) = \min_{\mathbf{R}} \left\{ \Psi\left(\frac{2}{n(n-1)} \sum_i \sum_{j>i} a_{ij}\eta_{ij}(\mathbf{x},\mathbf{R}) \right) \right\} \tag{6}$$

where Ψ is the normalization function defined in section 3.1. In practice we optimize this cost function by discretely sampling a dense set of normals in the hemisphere into which the initial normal is pointing. The optimization over the rotation $\mathbf{R}$ is restricted such that it provides a normal vector to the tangent plane inside a cone defined by the initial normal and an opening angle θ. In our experiments, we set θ to 60 degrees. Empirically we found that the normal optimization is a very important factor in achieving good reconstruction results. Without it protrusions and thin structures are carved much more often.

4.2 Narrow Band Graph-Cut

Existing volumetric graph-cut-based methods usually perform only one cut in a fairly big band around the initial surface estimate. By doing this, they assume that the maximum displacement of the true surface from the initial estimate is known. This is in particular not true if only rough silhouette estimates are available. Another disadvantage of one-shot graph-cuts over big bands is that it is hard to reason about visibility for voxels far from the initial surface estimate. Therefore, the consistency score computed for these voxels can be wrong. Another problem is setting the normals for the voxels in the band. Typically, either the normal of the closest surface point is chosen or the normal is computed using the surface given by the layer. However, this approximation can quickly become wrong when going further away from the original surface. Moreover, a big band can completely fill out protrusions and narrow parts of the object. As graph-cuts prefer short cuts, these protrusions are easily carved away. In addition, the danger of overcarving increases since, in a big band, there are more possibilities to find a "shortcut" than in a narrow band.

We construct a narrow band by expanding the current surface estimate to the outside and the inside. The inside layers carve inconsistent voxels while the outside layers allow us to correct segmentation errors in the initial silhouette images. This iterative approach

makes the visibility reasoning and the normal propagation much more stable as we stay close to the current surface estimate. Using narrow bands makes it possible to avoid the carving of narrow parts and protrusions since they are not totally filled out. Assuming that there is a clear low-cost path through the band (if there is one the normal optimization will find it), the graph-cut does not deviate from this path. In fact, the differences in path length through the band are so small, that they cannot outweigh the penalty of cutting through a high-cost region. Using narrow bands also allows us to recover concavities. Indeed, we are able to recover even deep concavities, as can be seen from the results we obtained on real-world data sets. While we cannot guarantee to reach a global minimum (which might not be useful due to the minimal surface bias of the graph-cut), we overcome local minima inside the band, which is sufficient for most practical scenes. This is an advantage over using level sets since by choosing the size of the band we can overcome local minima within the band and achieve a better solution than level sets.

In our implementation, we apply the graph-cuts iteratively until the consistency score of the surface converges. The number of iterations depends on how close the initial estimate approximates the true surface. In most cases, less than 20 iterations are needed to converge to a stable reconstruction.

5 Results

We present the evaluation of our method on two standard data sets and one custom data set that illustrate the advantages of our algorithm. The first data set is the *gargoyle* set provided by K. Kutulakos. It shows a stone sculpture and contains 16 images with a resolution of 719×485 pixels. The difficulty in this data set is that it contains holes which introduce a fair amount of self-occlusion. The input images were roughly segmented to obtain an initial estimate of the visual hull (see figure 2(a)). The resolution used is 200^3 voxels. We used a 5×5 grid to sample the plane used for estimating the voxel score. As in all of our experiments, we allowed a movement of one voxel layer outside and inside the current estimate for every iteration of the graph-cut. Therefore the width of the band is three layers. The width of the band stays unchanged during the optimization. The reconstruction time was 60 minutes. As shown in figure 2(a), we obtain an accurate reconstruction of the gargoyle. In particular, we correctly reconstruct protrusive parts of the object like the ears and the nose without the need of an 'ad-hoc' ballooning term. At the same time, we recover the concavities of the object, for instance around the eyes and in particular in the region between its body and the stick it is holding. The reconstruction converges after 19 iterations, while most of the surface details are already recovered after 10 iterations. Only the deep concavity between the belly and the stick requires more carving cycles. This shows that our method is able to recover even deep concavities. Our reconstruction is of at least the same quality as the one presented in [2] where the photoflux is used to preserve thin structures.

The second data set we used our method on is the *dinoSparseRing* data set provided by [21]. This set contains 16 images of a plaster dinosaur model with a resolution of 640×480 pixels. The reconstruction time was 62 minutes. The particular difficulty of this data set is the lack of texture. We used a resolution of $200 \times 234 \times 200$ and a 5×5 grid for this reconstruction. The algorithm converged after 21 iterations and recovered most of the surface details as seen in figure 2(b). The concavities between the legs and

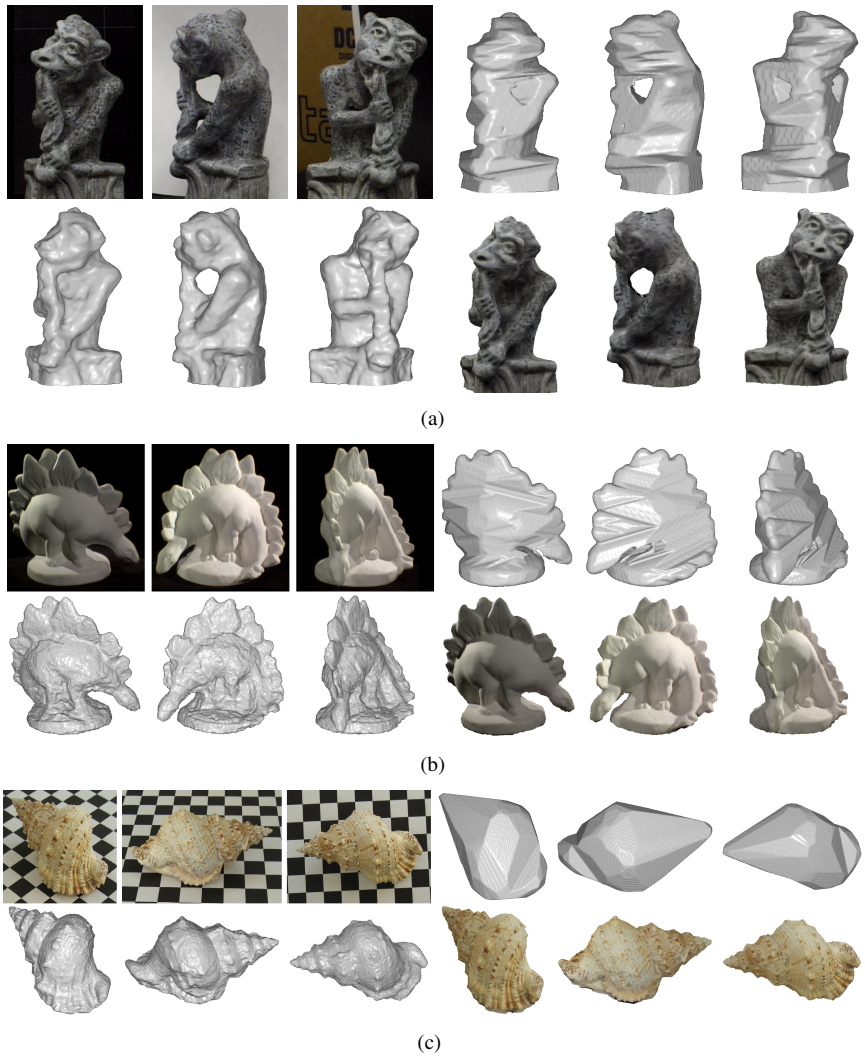

Figure 2: Reconstruction results. The first row of each result shows three input images and three views of the visual hull which is used as a starting point for the reconstruction. The second row shows the untextured and the textured reconstruction results.

the scales on the back were both recovered. The surface on the back of the model is fairly smooth while the surface on the front is a little rough due to the very uniform intensity. When compared to the reconstruction result of [29] on the same data set, we can see a clear improvement especially around the tail area and around the legs. Compared to [27], we obtain a more accurate surface estimate with more details especially around the legs. Our reconstruction achieved an accuracy of 0.89mm (defined as the distance that brings 90% of the reconstruction result within the ground-truth surface) and a completeness of 95.0% (defined as the percentage of the ground-truth that lies within 1.25mm of the recon-

struction results) in the Middlebury multi-view evaluation. These results reflect that our method provides accurate reconstructions which surpass the results of [29] (using a ballooning term) (1.18mm/90.8%) and [27] (using silhouette constraints) (1.26mm/89.3%) in both accuracy and completeness.

The third data set consists of 24 images of a shell at a resolution of 1600×1200. The reconstruction time was 112 minutes. Since the shell has many small protrusions and concave regions, it is a perfect test object to show the properties of our reconstruction method. We roughly segmented the shell by drawing a bounding polygon around the object. We used a voxel grid of size $415 \times 276 \times 200$ and sampled the tangent plane with an 11×11 grid due to the bigger image resolution. The reconstruction converged after 35 iterations because we started far from the true surface. This also shows that the method is not critically dependent on the initial surface estimate. Figure 2(c) shows our reconstruction results. We correctly recover the concavities and the small protrusions on the shell. The reconstruction also recovers the undulations on the base of the shell. This shows that our method can successfully preserve protrusions and at the same time recover concavities without relying on a ballooning term or on exact silhouette images.

6 Conclusion

We presented a novel iterative reconstruction algorithm designed to overcome problems resulting from the graph-cut inherent bias for shorter cuts. At each iteration, we optimize the surface normals of the current surface and apply a volumetric graph-cut over narrow bands around the current surface estimate.

Experimental results obtained on standard and custom data sets show that our method preserves protrusions and at the same time recovers concavities. We applied the proposed algorithm on ground-truth data and compared the obtained results with the results of existing volumetric graph-cut-based methods that rely on a global ballooning term [29] and on silhouette constraints [27]. We obtained a reconstruction with higher accuracy and completeness.

References

[1] Y. Boykov and V. Kolmogorov. Computing geodesics and minimal surfaces via graph-cuts. In *ICCV*, 2003.

[2] Y. Boykov and V. Lempitsky. From photohulls to photoflux optimization. In *BMVC*, 2006.

[3] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph-cuts. *IEEE PAMI*, 23(11):1222–1239, 2001.

[4] O. Faugeras and R. Keriven. Variational principles, surface evolution, pde's, level set methods, and the stereo problem. *IEEE Transactions on Image Processing*, 7(3):336–344, 1998.

[5] J. S. Franco and E. Boyer. Exact polyhedral visual hulls. In *BMVC*, 2003.

[6] Y. Furukawa and J. Ponce. Carved visual hulls for image-based modeling. In *ECCV*, 2006.

[7] Y. Furukawa and J. Ponce. Accurate, dense, and robust multi-view stereopsis. In *IEEE CVPR*, 2007.

[8] B. Goldlücke and M. Magnor. Spacetime-coherent geometry reconstruction from multiple video streams. In *IEEE CVPR*, 2004.

[9] M. Habbecke and L. Kobbelt. Iterative multi-view plane fitting. In *VMV*, 2006.

[10] M. Habbecke and L. Kobbelt. A surface growing approach to multi-view stereo reconstruction. In *IEEE CVPR*, 2007.

[11] C. Hernandez and F. Schmitt. Silhouette and stereo fusion for 3d object modeling. *Computer Vision and Image Understanding*, 96(3):367–392, 2004.

[12] C. Hernandez, G. Vogiatzis, and R. Cipolla. Probabilistic visibility for multi-view stereo. In *IEEE CVPR*, 2007.

[13] A. Hornung and L. Kobbelt. Hierarchical volumetric multi-view stereo reconstruction of manifold surfaces based on dual graph embedding. In *IEEE CVPR*, 2006.

[14] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE PAMI*, 26(2):147–159, 2004.

[15] K. Kutulakos. Approximate n-view stereo. In *ECCV*, 2000.

[16] K. Kutulakos and S. Seitz. A theory of shape by space carving. *IJCV*, 38(3):199–218, 2000.

[17] A. Laurentini. The visual hull concept for silhouette-based image understanding. *IEEE PAMI*, 16(2):150–162, 1994.

[18] S. Lazebnik, Y. Furukawa, and J. Ponce. Projective visual hulls. *IJCV*, 74(2):137–165, 2007.

[19] V. Lempitsky, Y. Boykov, and D. Ivanov. Oriented visibility for multiview reconstruction. In *ECCV*, 2006.

[20] J. P. Pons, R. Keriven, and O. Faugeras. Multi-view stereo reconstruction and scene flow estimation with a global image-based matching score. *IJCV*, 72(2):179–193, 2007.

[21] S. Seitz, B. Curles, J. Diebel, D. Scharstein, and R. Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *IEEE CVPR*, 2006.

[22] S. Seitz and C. Dyer. Photorealistic scene reconstruction by voxel coloring. *IJCV*, 25(3):151–173, 1999.

[23] S. Sinha, P. Mordohai, and M. Pollefeys. Multi-view stereo via graph cuts on the dual of an adaptive tetrahedral mesh. In *ICCV*, 2007.

[24] S. Sinha and M. Pollefeys. Multi-view reconstruction using photo-consistency and exact silhouette constraints: A maximum-flow formulation. In *ICCV*, 2005.

[25] G. Slabaugh, B. Culbertson, T. Malzbender, M. Stevens, and R. Schafer. Methods for volumetric reconstruction of visual scenes. *IJCV*, 57(3):179–199, 2004.

[26] R. Szeliski. Rapid octree construction from image sequences. *Computer Vision, Graphics and Image Processing: Image Understanding*, 58(1):23–32, 1993.

[27] S. Tran and L. Davis. 3d surface reconstruction using graph-cuts with surface constraints. In *ECCV*, 2006.

[28] G. Vogiatzis, C. Hernández Esteban, P. H. S. Torr, and R. Cipolla. Multiview stereo via volumetric graph-cuts and occlusion robust photo-consistency. *IEEE PAMI*, 29(12):2241–2246, 2007.

[29] G. Vogiatzis, P. H. S. Torr, and R. Cipolla. Multi-view stereo via volumetric graph-cuts. In *IEEE CVPR*, 2005.

[30] N. Xu, R. Bansal, and N. Ahuja. Object segmentation using graph cuts based active contours. In *IEEE CVPR*, 2003.

[31] T. Yu, N. Ahuja, and W.-C. Chen. Sdg cut: 3d reconstruction of non-lambertian objects using graph cuts on surface distance grid. In *IEEE CVPR*, 2006.

[32] G. Zeng, S. Paris, L. Quan, and F. Sillion. Accurate and scalable surface representation and reconstruction from images. *IEEE PAMI*, 29(1):141–158, 2007.