Surface Reconstruction by Propagating 3D Stereo Data in Multiple 2D Images

Gang Zeng, Sylvain Paris, Long Quan, Maxime Lhuillier

► To cite this version:

Gang Zeng, Sylvain Paris, Long Quan, Maxime Lhuillier. Surface Reconstruction by Propagating 3D Stereo Data in Multiple 2D Images. The 8th European Conference on Computer Vision, 2004, France. p. 163-174. hal-00118590

HAL Id: hal-00118590

<https://hal.science/hal-00118590>

Submitted on 5 Dec 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Surface Reconstruction by Propagating 3D Stereo Data in Multiple 2D Images

Gang ZENG¹, Sylvain PARIS², Long QUAN¹, and Maxime LHUILLIER³

¹ Dep. of Computer Science, Hong Kong University of Science and Technology,
Clear Water Bay, Kowloon, Hong Kong

² ARTIS / GRAVIR-IMAG, INRIA Rhône-Alpes,
38334 Saint Ismier, France

³ LASMEA, UMR CNRS 6602, Université Blaise-Pascal,
63177 Aubière, France

Abstract. We present a novel approach to surface reconstruction from multiple images. The central idea is to explore the integration of both 3D stereo data and 2D calibrated images. This is motivated by the fact that only robust and accurate feature points that survived the geometry scrutiny of multiple images are reconstructed in space. The density insufficiency and the inevitable holes in the stereo data should be filled in by using information from multiple images. The idea is therefore to first construct small surface patches from stereo points, then to progressively propagate only reliable patches in their neighborhood from images into the whole surface using a best-first strategy. The problem reduces to searching for an optimal local surface patch going through a given set of stereo points from images. This constrained optimization for a surface patch could be handled by a local graph-cut that we develop. Real experiments demonstrate the usability and accuracy of the approach.

1 Introduction

Surface reconstruction from multiple images is one of the most challenging and fundamental problems of computer vision. Although effective for computing camera geometry, most recent approaches [6, 9] reconstruct only a 3D point cloud of the scene, whereas surface representations are indispensable for modelling and visualization applications. Surface reconstruction is a natural extension of the point-based geometric methods. Unfortunately, using 3D data from such a passive system in the same way as range scanner data is often insufficient for a direct surface reconstruction method. The main difficulty is that the 3D points are sparse, irregularly distributed and missing in large areas. On the other hand, most image-based surface reconstruction approaches equally consider all surface points. They ignore the *feature points* although these points can be precisely matched between multiple images and therefore lead to accurate 3D locations. These shortcomings motivate us to develop a new approach to constructing a surface from stereo data [17], while using extra image information that is still available from a passive system.

Surface reconstruction from 3D data Surface reconstruction from scanned data is a traditional research topic. Szeliski et al. [27] use a particle-based model of deformable surfaces; Hoppe et al. [11] present a signed distance for implicit surfaces; Curless and Levoy [4] describe a volumetric method; and Amenta et al. [1] develop a set of computational geometry tools. Recently, Zhao et al. [32] develop a level-set method based on a variational method of minimizing a weighted minimal surface. Similar work is also developed by Whitaker [30] using a MAP framework. Surface reconstruction from depth data obtained from stereo systems is more challenging because the stereo data are usually much sparser and less regular. Fua [8] uses a system of particles to fit the stereo data. Kanade et al. [20] propose a deformable mesh representation to match the multiple dense stereo data. These methods that perform reconstruction by deforming an initial model or tracking discretized particles to fit the points are both topologically and numerically limited. Compared with our method, Hoff and Ahuja [10] and Szeliski and Golland [28] handle the data in an unordered way. Tang and Medioni [29] and Lee et al. [16] formulate the problem under a tensor voting framework.

Surface reconstruction from 2D images Recently, several volumetric algorithms [5, 14, 15, 25] simultaneously reconstruct surfaces and obtain dense correspondences. The method of *space carving* or *voxel coloring* [15, 25] works directly on discretized 3D space, voxels, based on their image consistency and visibility. These methods are purely local and therefore rely either on numerous viewpoints or on textured surfaces to achieve satisfying results. Kolmogorov and Zabih [14], Roy [23], and Hishikawa and Geiger [12] propose direct discrete minimization formulations that are solved by graph-cuts. These approaches achieve disparity maps with accurate contours but limited depth precision. The latest improvement by Boykov and Kolmogorov [3] overcomes this point but is restricted to data segmentation. Paris et al. [22] also propose a continuous functional but is restricted to open surfaces. Faugeras and Keriven [5] propose a method implemented by level-sets, which is intrinsically a multiple view and naturally handles the topology and occlusion problems. However, it is not clear under what conditions their methods converge as the proposed functional seems non-convex.

Some ideas in our approach are inspired by these methods, but fundamentally these methods either solely operate on 3D data or on 2D data. Our approach is more similar to the work of Lhuillier and Quan [18], which integrates 3D and 2D data under a level-set framework and suffers from the same limitations. Our strategy to achieve this goal is to perform a propagation in 3D space starting from reliable feature points. The propagation is driven by image information to overcome the insufficiency of 3D data. Among possible types of image information, *cross-correlation* considers the local texture information and gives the surface location with great precision – especially its zero-mean normalized version (ZNCC) that is robust to lighting variations, but error matches occur when coherence constraints are not taken into consideration. We define a propagation technique as a local optimization to combine coherence constraints and image information. We make the surface grow patch by patch and we show that with an appropriate graph-cut technique, each patch is optimal under our hypotheses.

Fig. 1. Our framework is centered on the propagation loop that progressively extends the surface. Each iteration of the loop picks a seed, builds an optimal patch surrounding it, updates the surface and selects new seeds for further propagation. This process is initialized with 3D points computed by a stereoscopic method.

Our contributions Our approach explores the integration of both stereo data and calibrated images. It fully exploits the feature points to start the reconstruction process to avoid potential ambiguities and build a precise surface. This is achieved through a new and innovative propagation framework controlled by 2D images. The graph-cut optimization is used locally to guarantee that the propagation is optimal under our hypotheses. Theoretical justifications are given for a formal understanding of the process and technical issues are exposed.

2 Problem statement and formulation

Given a set of calibrated images $\{C_j\}$, a set of stereo points $\mathcal{Q} = \{\mathbf{q}_i\}$ derived from the given images, the goal is to reconstruct a surface $\mathcal{S}$ of the objects in the scene. The problem is different to surface reconstructions from multiple images as addressed in [5, 15, 25] in which only 2D images are used without any 3D information. It is also different to surface reconstructions from scanned 3D data without 2D image information [4, 8, 11, 20, 27, 32], as the 3D stereo data are often insufficient in density and accuracy for a traditional surface reconstruction.

Futhermore, most of the existing surface reconstruction techniques characterize their results by a global functional. However, Blake and Zisserman [2] demonstrate that their optimization scheme follows local rules. For instance, it is shown that a discontinuity has a local influence on the result and that this influence disappears beyond a given distance. Examining the global approach of Faugeras and Keriven [5] leads to the same remark: even if the initial formulation relies on a surface integral, the practical resolution is performed with local differential operators. This motivates our approach (Figure 1): the surface is built patch by patch, each one being locally optimal for a global criterion.

Since the patch is local, it is reasonable to establish the local coordinates system (x, y, z) where xy is the tangent plane to the object surface and z is the normal, to parameterize the patch as a single valued height field $z = h(x, y)$.

Then, we look for a surface patch $\{(x, y, h(x, y)) | (x, y) \in \mathcal{D}\}$ by minimizing a functional of type:

$$\iint_{\mathcal{D}} c(x, y, h(x, y)) dx dy, \quad (1)$$

where $c(x, y, h(x, y))$ is a cost function accounting for the consistency of the surface point $(x, y, h(x, y))$ in multiple images, for instance, either photo-consistency or cross-correlation.

We also require that the optimized surface patch goes through the existing 3D points if they do exist in the specified neighborhood. We are therefore looking for an interpolating surface through the given 3D points. That is, to minimize

$$\iint_{\mathcal{D}} c(x, y, h(x, y)) dx dy \quad \text{with} \quad h(x_i, y_i) = z_i, \quad (2)$$

where (x_i, y_i, z_i) is the local coordinate of a given 3D point $\mathbf{q}_i$. This functional can either be regularized as a minimal surface $\iint c(x, y, h(x, y)) ds$ leading to a level-set implementation [5, 13, 26]. We do not follow this path as it very often results in an over-smoothed surface [18] due to high order derivatives involved in the dynamic surface evolution. Instead, we apply a recent approach that reaches sharper results [22].

We use first derivatives as smoothing term

$$s(x, y, h(x, y)) = \alpha \left(\left| \frac{\partial h}{\partial x}(x, y) \right| + \left| \frac{\partial h}{\partial y}(x, y) \right| \right) \quad (3)$$

with α controlling the importance of $s(\cdot)$, to minimize the following functional

$$\iint_{\mathcal{D}} (c(x, y, h(x, y)) + s(x, y, h(x, y))) dx dy \quad \text{with} \quad h(x_i, y_i) = z_i. \quad (4)$$

The advantage of this formulation is that this continuous functional can be discretized, then optimized by a graph-cut algorithm. This is an adaption of the graph-cut approach, with a simplified smoothing term $s(\cdot)$ to a local environment that we call a *local graph-cut*. The choice of $c(\cdot)$ is purely independent of the general framework. Recently researchers mainly focus on *photo-consistency* and *cross-correlation*. In this paper we present an implementation of the propagation framework with $c(\cdot) = g(zncc(\cdot))$ with $g(\cdot)$ being a decreasing function like $x \mapsto (1 - x)$ to fit the minimization goal.

This local formulation relies on the surface orientation to choose the parameterization of the local patch, the normal direction of a given 3D point has to be estimated from its neighborhood. We define a seed point as a couple $(\mathbf{p}_i, \mathbf{n}_i)$ where $\mathbf{p}_i$ is the 3D position and $\mathbf{n}_i$ the surface normal direction at $\mathbf{p}_i$. This is dependent on the reconstruction algorithms used, which is further discussed in Section 4. The above analysis finally leads to a surface that optimally interpolates the given set of 3D stereo points by 2D image information defined in the functional (4).

3 Algorithm description

We propose here an algorithm to apply the previously discussed framework. This algorithm follows the general organization presented in Figure 1. First, using the technique discussed in Section 4, 3D stereo points are computed. The local surface orientation at these points is estimated to attach with the points to form the initial seeds. Then the information is propagated from the most reliable data to grow the surface.

3.1 Initialization from 3D stereo points

To start the algorithm, we initialize the list of seeds, from which the surface is propagated gradually. A high number of 3D points are robustly computed from the given set of images by stereo and bundle-adjustment methods (Section 4). Figure 4 shows these points for different data sets. Compared with standard sparse points, these points are highly redundant and well distributed in several parts. This redundancy makes it possible to evaluate the surface orientation. It is important to address here that these input points are regarded as a “black box” without assuming any property. Other preprocessing methods are also possible as long as the surface orientation can be estimated at each given position.

For each 3D point $\mathbf{p}_i$, the surface orientation $\mathbf{n}_i$ is provided by the symmetric 3×3 positive semi-definite matrix $\sum_{\mathbf{y} \in \mathcal{B}_r(\mathbf{p}_i) \cap \mathcal{Q}} (\mathbf{y} - \mathbf{p}_i) \otimes (\mathbf{y} - \mathbf{p}_i)$, where $\mathcal{B}_r(\mathbf{p}_i)$ denotes the ball of radius r . Among the eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ respectively associated to the eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$, we choose $\mathbf{n}_i$ to be either $\mathbf{v}_3$ or $-\mathbf{v}_3$. The sign depends on the cameras used to reconstruct $\mathbf{p}_i$. The confidence is indicated by the ratio between the lowest eigenvalue and the larger ones.

$\mathcal{B}_r(\mathbf{p}_i) \cap \mathcal{Q}$ may contain very few points, baffling the orientation estimation. And in dense regions, a large radius results in an over-smoothed estimation whereas a small radius makes the estimation sensitive to noise. Therefore r is defined as a function of $\mathbf{p}_i$: in dense regions, r is fixed to a reference value r_{dense} representing the minimum scale. In the diluted regions, the radius is increased so that $\mathcal{B}_r(\mathbf{p}_i)$ contains at least k 3D stereo points. From many experiments, a good compromise is to define r_{dense} to be the radius of local patches and k to be 15-20.

3.2 Propagation loop

Starting from the seed list initialized with the 3D stereo points, this step makes the surface grow until it reaches the final result. Each seed is handled one by one to generate a surface patch and create new seeds by the local graph-cut technique. It is divided into the 3 following issues:

1. The selection of the next seed from the current seed list.
2. The generation of an optimal patch from this seed.
3. The creation of the new seeds on this patch to add in the list.

3.3 Selection of the next seed

To select a new seed $(\mathbf{p}, \mathbf{n})$ for propagation, we need a criterion Π to evaluate how “good for propagation” a seed is. Once we have such a criterion, we follow a classical *best-first strategy* to ensure that the most reliable seed is picked each time. This choice directly drives the propagation because it indicates where the growing regions are. It is our global control over the surface reconstruction.

First of all, the initial seeds (*i.e.* 3D stereo points) are regarded as the reliable 3D points on the surface. Therefore, they are always selected before the generated seeds. The algorithm ends when there is no seed left in the list.

Selection criterion for 3D stereo points The criterion of the stereo points is defined as the confidence of the orientation estimation described in Section 3.1. This confidence is related to the local planarity of the stereo points and therefore estimates the accuracy of the normal computation. This leads to $\Pi = \frac{\lambda_2}{\lambda_3}$. The visibility of a stereo point is given by the cameras used to reconstruct it.

Selection criterion for generated seeds For a generated seed, we use the ZNCC correlation score Z by its two most front-facing cameras, since a strong match gives a high confidence. This strategy ensures that the surface grows from the part which is more likely to be precise and robust. Thus: $\Pi = Z$. Correct visibility is computed thanks to this propagation criterion. If the criterion is computed from occluded cameras, the local textures in both images does not match and the ZNCC value is low. Therefore a seed without occlusion is processed before a seed with occlusion. The occluded parts “wait” until other parts are reconstructed. The visibility of the processed seed is classically determined by the current propagated surface using a ray-tracing technique.

3.4 Generation of an optimal patch from a given seed

Given a seed $(\mathbf{p}, \mathbf{n})$ with its visibility, a patch is grown in its neighborhood to extend the existing surface. Since it is a local representation of the surface, it is parameterized by a height field $z = h(x, y)$ relatively to the tangent plane of the surface. Therefore, for the local coordinate (x, y, z) , the x and y axes are arbitrary chosen orthogonally to $\mathbf{n}$, and the z axis is parallel to $\mathbf{n}$ with the coordinate origin at $\mathbf{p}$. A surface patch is defined by $\{(x, y, h(x, y)) | (x, y) \in \mathcal{D}\}$ that is minimal for functional (4). This patch is also enforced to pass through the stereo points and the previously computed surface.

The graph-cut technique described in [22] is then applied. This technique reaches an exact minimum of the functional (4) up to any arbitrary discretization. This involves a graph whose cuts represent the set of all possible surfaces $z = h(x, y)$. Thus the classical max-flow problem [7] is equivalent to an exhaustive search leading to a minimal solution of (4). Moreover, it allows us to constrain $h(x, y)$ to a specified z -range.

To use this technique, the local domain is discretized as a regular rectangular grid $\{X_1, \dots, X_{n_x}\} \times \{Y_1, \dots, Y_{n_y}\} \times \{Z_1, \dots, Z_{n_z}\}$ separated by $\Delta x, \Delta y, \Delta z$. The functional (4) is discretized into the following form:

Fig. 2. Correspondence between the graph and the voxel $(x, y, h(x, y))$.

Fig. 3. Only the center part of the optimal solution is actually used to extend the existing surface (shaded area).

$$\sum_i \sum_j (c(X_i, Y_j, h(X_i, Y_j)) \Delta x \Delta y + |h(X_{i+1}, Y_j) - h(X_i, Y_j)| \Delta y + |h(X_i, Y_{j+1}) - h(X_i, Y_j)| \Delta x) \quad \text{with the constraint } h(x_i, y_i) = z_i. \quad (5)$$

For any $(x, y) \in \mathcal{D}$ where the surface has been already computed (including the 3D stereo points), the h value is restricted to the only corresponding z value. Then for each remaining (x, y, z) position, a ZNCC value is computed using the two most fronto-parallel visible cameras to limit the perspective distortion. These ZNCC values are transformed into $c(\cdot)$, resulting in a dense 3D sampling.

Based on the above discrete form (5), a graph is embedded into a 3D grid superimposed on the voxels with correspondence shown in Figure 2. Finally, this graph-flow problem is solved and the resulting minimum cut is equivalent to the optimal patch. Graph design and proof of optimization are in [22]. The confidence of the patch is given by the quantity of the maximum flow F . A smaller value gives a higher confidence, while the large enough value indicates that this patch should be discarded.

Even if the patch is trustworthy, the borders of the patch might not be suitable for the final surface. The border points have a truncated neighborhood and potentially ignore some neighboring data. To avoid this caveat, only the center part of the optimal solution is kept (see Figure 3). The discarded border only provides a complete neighborhood to the center part. Finally, this center part is used to extend the existing surface.

3.5 Creation of the new seeds to add in the list

We have grown a surface patch for the selected seed in the previous section. To continue the propagation, new seeds are created from this patch. These new seeds will be selected later to drive further propagation. The location of the new seeds is driven by several aspects.

Patch quality First of all, the quantity of the maximum flow F indicates the confidence of the optimal patch. With low enough confidence, the surface patch is discarded and no seed is created.

Match quality A point with a high ZNCC value Z is more likely to provide a robust starting point for further propagation.

Surface regularity A singular point does not represent accurate properties of the patch. Using the principal curvatures κ_1 and κ_2 , points with high curvature $K = \kappa_1^2 + \kappa_2^2$ are therefore to be avoided.

Propagation efficiency To ensure a faster propagation, distant points are preferred. This relies on the distance D between the patch center and the potential new seeds.

A value Λ is computed for each potential location of a new seed to represent its appropriateness relative to these objectives.

$$\Lambda = \frac{Z^{\gamma(Z)} \cdot D^{\gamma(D)}}{F^{\gamma(F)} \cdot K^{\gamma(K)}} \quad (6)$$

where $\gamma(\cdot)$ are non-negative weights to balance the different criteria. In our algorithm, we use $\gamma(Z) = \gamma(D) = \gamma(F) = \gamma(K) = 1$, while further study is needed to evaluate the importance of each criterion.

The number of new seeds created is inspired by the triangle mesh configuration. From the Euler property, the average number of neighbors of a vertex is 6 and the average angular distance between two neighbors is $\frac{\pi}{3}$. Thus, the directions of the new seeds in relation to the patch center are selected so that the angular distance between two neighbor seeds lies in $[\frac{2\pi}{5}, \frac{2\pi}{7}]$. In each direction, the location $\mathbf{p}'$ with the highest Λ is selected and the normal $\mathbf{n}'$ at $\mathbf{p}'$ is attached to form a new seed.

4 Implementation and experiments

4.1 Acquisition of 3D stereo data

3D stereo data can be obtained using different approaches. A traditional stereo rig [21, 24] could deliver quite dense points, but it only gives a small part of the object. Automatically merging partial stereo data into a complete model is not easy, except that multiple stereo rigs are calibrated off-line. We choose a more general reconstruction method from an uncalibrated sequence [6, 9, 31]. Our quasi-dense implementation [17] is similar to the standard uncalibrated approaches with the main difference that we use a much denser set of points re-sampled from the disparity map instead of points of interest. To model a complete object, we usually make a full turn around the object by capturing about 30 images to compute the geometry of the sequence.

4.2 Efficiency consideration

To select a seed to propagate, we use selection criterion Π to extend the surface from its most robust regions (Section 3.3). However, it may cause inefficiency, because the neighboring surface may have already been created from other seeds.

Therefore, the *area criterion* A is added to define $\tilde{\Pi} = A \cdot \Pi$. The definition of A represents the efficiency of the patch generation. It is based on the total patch area a_P and the covered area a_c : $A = 1 - \frac{a_c}{a_P}$. Since this value evolves during the process, it is stored with the seed and updated in a “lazy” way: When the best seed is selected by $\tilde{\Pi}$, the corresponding A is updated according to the current surface. If the value changes to a non-zero value (it always decreases because the surface is growing), $\tilde{\Pi}$ is updated and the seed is put back in the list. If it changes to zero, the seed is discarded because it would provide no surface extension. This case mainly occurs for 3D stereo points which are redundant. Otherwise (*i.e.* A does not change), the seed is selected for further processing.

4.3 Representation of the growing surface with a distance field

We use a distance field to register the growing surface. The surface is represented by the signed distance function d in its *narrow band*, which is Euclidean distance with the sign indicating either inside or outside. When a surface patch is generated, the distance field is updated in its narrow band. This results in the merger of this patch with its neighboring ones. Finally, the zero set $\mathcal{Z}(d)$ is our estimation for $\mathcal{S}$, which is extracted by the *marching cube* method [19].

4.4 Experimental results

Three typical modelling examples are shown in Figures 4. We usually acquire about 30 images with a hand-held camera around the objects. The first “toy” example is used to examine the correctness and robustness of our algorithm. The “toy” is difficult since fur is traditionally hard for surface reconstruction. We have made a tradeoff between the local features and the orientation robustness based on the local coherence. As we can see from the reconstructed shape, our algorithm handles occlusion correctly. Two “face” examples illustrate the accuracy of our algorithm. One face has more texture, while the other has less. The details around the eyes, the noses and the ears can be clearly seen thanks to the exploitation of the feature points. The shape is also satisfactory in the other regions, such as the foreheads and the cheeks because of the new propagation technique. It is important to address that the final shape optimally interpolates the given set of 3D stereo points by 2D image information defined in the functional (4).

The results in the above two face examples are of almost same quality. We may notice that the 3D feature points are good surface points, but their quantity is not an important factor for reconstruction. Since the derived 3D points are highly redundant, actually many of the points provide limited information. Future study is needed for the quantity and the positions of the initial seeds.

Fig. 4. Each row shows the results for one example: (a) One of the input images; (b) Reconstructed 3D points; (c,d) Surface shape at two different viewpoints; (e) Surface shape with color.

Fig. 5. Comparative results: (a) Space carving only gives a rough estimation; (b) Level set method yields an over-smoothed result; (c) The propagation gives the most detailed surface geometry.

Figure 5 shows a comparative study among the space carving method [15], the level set method [18] and our propagation approach. The space carving method only uses image information, and misses a lot of details on the nose, the hair and the ears due to its photo-consistency criterion, which often yields an over-estimation on the intensity-homogeneous regions. The level set method and our approach are both based on the stereo points and the image information. However the level set method tends to over-smooth the surface by losing many geometric details.

5 Conclusions

We have proposed a novel approach for surface reconstruction by propagating 3D stereo data in multiple 2D images. It is based on the fact that the feature points can be accurately and robustly determined. These points give important information for surface reconstruction. On the other hand, they are not sufficient to generate the whole surface representation. The major motivation of this study is to improve the insufficiency of 3D stereo data by using original 2D image information.

Our strategy to achieve this goal is to perform a propagation in a 3D space starting from reliable feature points. We have introduced a new functional (4) integrating both 3D stereo data points and 2D image information. With this functional, the surface grows patch by patch from the most reliable regions by the local graph cut technique using a best-first strategy. Finally, the shape optimally interpolates the given set of 3D stereo points by 2D image information defined in the functional (4). This approach have been extensively tested on real sequences and very convincing results have been shown.

Acknowledgements

This project is partially supported by the Hong Kong RGC grant HKUST 6188/02E. Also thanks to Yichen Wei and Ajay Kumar for allowing us to use their head images.

References

1. N. Amenta, S. Choi, and R. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, 2001.
2. A. Blake and A. Zisserman. Visual reconstruction. *MIT Press*, 1987.
3. Y. Boykov and V. Kolmogorov. Computing Geodesics and Minimal Surfaces via Graph Cuts. *ICCV'03*.
4. B. Curless and M. Levoy. A volumetric method for building complex models from range images. *SIGGRAPH '96*.
5. O. Faugeras and R. Keriven. Complete dense stereovision using level set methods. *ECCV'98*.

6. O. Faugeras, Q. Luong, and T. Papadopoulos. Geometry of Multiple Images. *MIT Press*, 2001.
7. L. Ford and D. Fulkerson. Flows in Networks. *Princeton University Press*, 1962.
8. P. Fua. From multiple stereo views to multiple 3d surfaces. *IJCV'97*.
9. R.I. Hartley and A. Zisserman. Multiple View Geometry in Computer Vision. *CU Press*, 2000.
10. W. Hoff, N. Ahuja. Surfaces from Stereo: Integrating Feature Matching, Disparity Estimation, and Contour Detection. *TPAMI'89*.
11. H. Hoppe, T. Deroose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. *Computer Graphics*, 1992.
12. H. Ishikawa and D. Geiger. Occlusions, discontinuities, and epipolar lines in stereo. *ECCV'98*.
13. R. Kimmel. 3D Shape Reconstruction from Autostereograms and Stereo. *Journal of Visual Communication and Image Representation*, 2002.
14. V. Kolmogorov and R. Zabih. Multi-camera scene reconstruction via graph cuts. *ECCV'02*.
15. K.N. Kutulakos and S.M. Seitz. A theory of shape by space carving. *ICCV'99*.
16. M.S. Lee, G. Medioni and P. Mordohai. Inference of Segmented Overlapping Surfaces from Binocular Stereo. *TPAMI'02*.
17. M. Lhuillier and L. Quan. Quasi-dense reconstruction from image sequence. *ECCV'02*.
18. M. LHUILLIER and L. QUAN. Surface Reconstruction by Integrating 3D and 2D Data of Multiple Views. *ICCV'03*.
19. W.E. Lorensen and H.E. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. *SIGGRAPH'87*.
20. P.J. Narayanan, P.W. Rander and T. Kanade. Constructing Virtual Worlds using Dense Stereo. *ECCV'98*.
21. M. Okutomi and T. Kanade. A multiple-baseline stereo. *TPAMI'93*.
22. S. Paris and F. Sillion and L. Quan. A Volumetric Reconstruction Method from Multiple Calibrated Views Using Global Graph Cut Optimization. *ACCV'04*.
23. S. Roy. Stereo without epipolar lines : A maximum-flow formulation. *IJCV'99*.
24. H.S. Sawhney, H. Tao and R. Kumar. A global matching framework for stereo computation. *ICCV'01*.
25. S.M. Seitz and C.R. Dyer. Photorealistic scene reconstruction by voxel coloring. *CVPR'97*.
26. J.A. Sethian. Level-set methods and fast marching methods. *Cambridge University Press*, 1999.
27. R. Szeliski, D. Tonnesen, and D. Terzopoulos. Modelling surfaces of arbitrary topology with dynamic particles. *CVPR'93*.
28. R. Szeliski and P. Golland. Stereo matching with transparency and matting. *IJCV'99*.
29. C.K. Tang and G. Medioni. Curvature-augmented tensor voting for shape inference from noisy 3d data. *TPAMI'02*.
30. R. Whitaker. A level-set approach to 3d reconstruction from range data. *IJCV'98*.
31. Z. Zhang and Y. Shan. A progressive scheme for stereo matching. *SMILE'2, ECCV'00 Workshop*.
32. H.K. Zhao, S. Osher, B. Merriman, and M. Kang. Implicit and non-parametric shape reconstruction from unorganized data using a variational level set method. *CVIU'00*.