

Continuous Global Optimization in Multiview 3D Reconstruction

Kalin Kolev¹, Maria Klodt¹, Thomas Brox¹,
Selim Esedoglu², and Daniel Cremers¹

¹ Department of Computer Science, University of Bonn, Germany

² Department of Mathematics, University of Michigan, USA

Abstract. In this work, we introduce a robust energy model for multiview 3D reconstruction that fuses silhouette- and stereo-based image information. It allows to cope with significant amounts of noise without manual pre-segmentation of the input images. Moreover, we suggest a method that can globally optimize this energy up to the visibility constraint. While similar global optimization has been presented in the discrete context in form of the maxflow-mincut framework, we suggest the use of a continuous counterpart. In contrast to graph cut methods, discretizations of the continuous optimization technique are consistent and independent of the choice of the grid connectivity. Our experiments demonstrate that this leads to visible improvements. Moreover, memory requirements are reduced, allowing for global reconstructions at higher resolutions.

1 Introduction

We consider the classical problem of inferring a dense 3D structure reconstruction of an object from a collection of views calibrated to a common world coordinate system. Among the multitude of existing methods one can distinguish between two major classes of techniques according to the exploited image information: shape from silhouettes and stereo.

In case of sparsely textured objects, silhouette-based methods exhibit favorable performance. Most of them aim at approximating the *visual hull* [17] of the imaged object. The visual hull is an outer approximation of the observed solid, constructed as the intersection of the visual cones associated with all image silhouettes. The earliest attempts use a volumetric representation of the scene, where each voxel is labeled as opaque or transparent according to each projection onto the images [19]. Latter developments led to the use of surface-based representations, which allow to impose regularization in an energy minimization framework. These methods are able to reconstruct a smooth version of the visual hull from the raw input images without the immediate need for manually outlined silhouettes [23, 27]. This is because the segmentation of each image is obtained through the evolution of a single surface in 3D rather than separate contours in 2D. As a result, such methods exhibit considerable robustness to

outliers and erroneous camera calibration. In [15] the robustness to noise and initialization is further increased by incorporating all available observations into a probabilistic framework.

The main drawback of silhouette-based approaches is their inability to reconstruct concavities, since these do not affect the silhouettes. Stereo-based methods capture such indentations by measuring photoconsistency of surface patches in space. The fundamental idea is that only points on the object’s surface have a consistent appearance in the input images, while all other points project to incompatible image patches. The earliest algorithms use carving techniques to obtain a volumetric representation of the scene by repeatedly eroding inconsistent voxels [16]. They do not enforce the smoothness of the surface, which often results in rather noisy reconstructions. Later, energy minimization techniques based on the integration of the data fidelity criterion on the unknown surface, have become more popular [10, 12, 18]. In these works, one seeks the surface with the smallest weighted area, where the weights reflect the local photoconsistency. Some recent approaches use a fusion of silhouette constraints and stereo information in order to achieve consistency in terms of silhouettes as well as image patches. Generally, there are two types of techniques to combine silhouette information and photoconsistency. The first strategy integrates silhouette constraints into stereo-based optimization [11, 22, 25]. The alternative is to use the visual hull merely as initialization for a stereo-based technique [26].

In this paper, we present an energy model which generalizes [15] by imposing photoconsistency constraints. Since for computing photoconsistency one needs the visibility of surface points, the photoconsistency term is collapsed at the beginning. With the resulting approximate visibility information, we can globally optimize the energy that includes both constraints. Our approach is related to the one introduced in [26]. However, the sought surface is not restricted to lie within some predefined band around the visual hull, which imposes different weighting of silhouette and stereo term. Another closely related work is the one of [18]. However, in this approach the silhouette constraint is replaced by a constant ballooning term that persistently prefers larger surfaces. To this end, visibility estimation is based on local graph edge orientations.

All previous methods use either local optimization, which is prone to instabilities and getting stuck in local minima, or *discrete* global optimization based on graph cuts. However, graph cuts can only minimize a certain class of discrete energies that are inconsistent to a corresponding continuous formulation, i.e., the solution does not converge to the continuous solution for finer grids. Thus, graph cuts are not rotationally invariant and favor polyhedral structures. In the scope of multiview reconstruction, an additional practical limitation is the relatively large memory consumption of graph cut methods, which can be decisive when computing reconstructions at a high resolution. The main contribution of the present work is the development of a novel technique for continuous global optimization for multiview reconstruction, which allows to avoid previously mentioned limitations. Similar techniques were recently proposed in the context of image segmentation [8]. In [1] another method for global optimization has been

	continuous local optimization	discrete global optimization	continuous global optimization
image segmentation	Snakes [14] Level Sets [5]	Graph Cuts [3]	TV-L1 [8] CSPs [1]
multiview reconstruction	Mesh-based [11] Level Sets [12]	Graph Cuts [18, 22, 25, 26]	this work

Table 1. Optimization techniques used in image segmentation and multiview 3D reconstruction.

proposed, which has been extended in [2] to 3D segmentation. However, this approach does not allow to incorporate regional information, which makes it inappropriate for our model. Both techniques were inspired by the original works of [13], [24]. Table 1 provides a number of representative works on local optimization, discrete and continuous global optimization in the context of image segmentation and multiview reconstruction, respectively.

The paper is laid out as follows. The next section contains a brief reviewing of related continuous global optimization techniques in the context of image segmentation. In section 3 we present and discuss the energy model. Section 4 is devoted to the optimization technique including implementation details. We show experimental results in section 5 and conclude the paper with a brief summary in section 6.

2 Continuous Convex Formulation for Image Segmentation

In a series of works [8, 4, 7] image segmentation functionals, namely the piecewise constant Mumford-Shah model and the snakes were addressed by means of convex formulations. For a given intensity image $I : \Omega \rightarrow \mathbb{R}$ the Mumford-Shah energy is defined as

$$E(C, \mu_1, \mu_2) = \int_{C_{in}} (\mu_1 - I(x))^2 dx + \int_{C_{out}} (\mu_2 - I(x))^2 dx + \nu \int_C ds, \quad (1)$$

where $C_{in}, C_{out} \subset \Omega$ denote the interior and exterior regions, separated by a contour C (see [20]). This optimization problem can be interpreted to be looking for the best approximation in the L^2 sense to the given image $I(x)$ among all functions that take only two values. In addition, there is a penalty on the length of the interface C that separates the regions where the two values μ_1 and μ_2 are taken. This model can be extended by replacing the classical Euclidean metric by a more sophisticated one defined in terms of the image content (see [6]):

$$E(C, \mu_1, \mu_2) = \int_{C_{in}} (\mu_1 - I(x))^2 dx + \int_{C_{out}} (\mu_2 - I(x))^2 dx + \nu \int_C g ds, \quad (2)$$

where $g = \frac{1}{1 + |\nabla I|^p}$ for some $p \geq 1$. Hence, the contour C is additionally pulled towards image edges. Functional (2) can be further expressed in terms of a set of binary functions (see [9]):

$$E(u, \mu_1, \mu_2) = \int_{\Omega} (\mu_1 - I)^2(1 - u) + (\mu_2 - I)^2u \, dx + \nu \int_{\Omega} g|\nabla u| \, dx \quad (3)$$

s. t. $u \in \{0, 1\}$.

The weighted total variation (TV) norm $TV_g(u)$ with the weight function g was introduced in [4]. It builds upon the TV regularizer introduced in the context of image denoising [21] and is defined as follows

$$TV_g(u) = \int_{\Omega} g|\nabla u| \, dx := \sup_{\phi \in \Phi_g} \left\{ \int_{\Omega} u(x) \operatorname{div} \phi(x) \, dx \right\}, \quad (4)$$

where

$$\Phi_g := \left\{ \phi \in C^1(\Omega, \mathbb{R}) \mid |\phi(x)| \leq g, \text{ for all } x \in \Omega \right\} \quad (5)$$

and g is a positive valued continuous and bounded function in Ω . Intuitively, this norm measures the area of the surface separating homogeneous regions locally weighted by g , in this case the edge-based measure.

Due to the constraint that u is a binary function, the minimization problem (3) is non-convex (because the space of binary functions is non-convex). However, in [8] it was found that when minimizing the total variation norm over all functions $u : \Omega \rightarrow \mathbb{R}$ where also intermediate values can be taken, the values of $u(x)$ converge to $\pm\infty$ almost everywhere. By enforcing $0 \leq u(x) \leq 1$ via a convex penalizer $\theta(u) := \max\{0, 2|u - \frac{1}{2}| - 1\}$ one obtains the energy

$$E(u, \mu_1, \mu_2) = \int_{\Omega} ((\mu_2 - I)^2 - (\mu_1 - I)^2) u \, dx + \nu \int_{\Omega} g|\nabla u| \, dx + \alpha \int_{\Omega} \theta(u(x)) \, dx. \quad (6)$$

Since this unconstrained version of (3) is convex, it can be globally minimized via standard variational methods. Moreover, minimization yields a binary function u almost everywhere. It was further shown in [8] that, by applying almost any threshold $\mu \in (0, 1)$, the resulting u will be a global minimizer of the non-convex functional (3).

The present work revisits this technique by formulating multiview reconstruction as a continuous convex optimization problem.

3 A Continuous Energy Model for Multiview Reconstruction

Let $V \subset \mathbb{R}^3$ be a volume, which contains the scene of interest, and $I_1, \dots, I_n : \Omega \rightarrow \mathbb{R}^3$ a collection of calibrated color images with perspective projections

π_1, \dots, π_n . We are looking for some surface $\hat{S} \subset V$ that gives rise to these images. This can be formulated as the energy minimization problem

$$E(S) = - \int_{R_{obj}^S} \log P_{obj}(x) dx - \int_{R_{bck}^S} \log P_{bck}(x) dx + \nu \int_S \rho(x) dA \rightarrow \min. \quad (7)$$

The energy consists of two parts. The first two terms impose the silhouette constraint via a probabilistic segmentation of the volume into object and background. The third term acts as a constraint both for smoothness and photoconsistency by seeking the minimal surface with respect to a Riemannian metric. The parameter ν controls the weighting of both parts of the energy.

The definition of the probability terms follows [15]. Regarding the silhouette constraint, according to a certain surface estimate S , all points in V can be divided into two classes: lying inside S or belonging to the background, i.e. $V = R_{obj}^S \cup R_{bck}^S$, where R_{obj}^S denotes the interior and R_{bck}^S the exterior. Considering the given image content we can assign each point $x \in V$ two probabilities $P_{obj}(x)$ and $P_{bck}(x)$ associated with R_{obj}^S and R_{bck}^S , respectively. More precisely

$$\begin{aligned} P_{obj}(x) &:= P(\{I_l(\pi_l(x))\}_{l=1, \dots, n} \mid x \in R_{obj}^S) \\ P_{bck}(x) &:= P(\{I_l(\pi_l(x))\}_{l=1, \dots, n} \mid x \in R_{bck}^S). \end{aligned} \quad (8)$$

Note that in this formulation $P_{obj}(x)$ and $P_{bck}(x)$ will generally not sum to 1. Considering dependence of the image observations we can write

$$\begin{aligned} P_{obj}(x) &= \sqrt[n]{\prod_{i=1}^n P(I_i(\pi_i(x)) \mid x \in R_{obj}^S)} \\ P_{bck}(x) &= 1 - \sqrt[n]{\prod_{i=1}^n [1 - P(I_i(\pi_i(x)) \mid x \in R_{bck}^S)]}. \end{aligned} \quad (9)$$

The probability of a voxel being part of the foreground is equal to the probability that *all* cameras observe this voxel as foreground, whereas the probability of background membership describes the probability of *at least one* camera seeing background. The root is for normalization with respect to the number of camera views, since both products will converge to 0 for $n \rightarrow \infty$. Thus, dependency between single image observations is expressed in terms of their geometric mean. Note that the fusion of all available image observations allows for quite robust silhouette-based surface estimation.

The foreground/background probabilities for the single image observations

$$\begin{aligned} P(I_i(\pi_i(x)) \mid x \in R_{obj}^S) &\sim \mathcal{N}(\mu_{obj}, \Sigma_{obj}) \\ P(I_i(\pi_i(x)) \mid x \in R_{bck}^S) &\sim \mathcal{N}(\mu_{bck}, \Sigma_{bck}). \end{aligned} \quad (10)$$

are modeled to be Gaussian distributed. The parameters of both models, i.e. mean vectors and covariance matrices, can be updated during optimization by projecting the current surface estimate onto the images in order to collect pixels, which belong to the respective regions. However, in our implementation we




Fig. 1. Orthogonal image features used for multiview 3D reconstruction: shape from silhouettes vs. shape from stereo. (a) Region-based color information. Surface concavities are not presented in the reconstruction. (b) Stereo-based point matching. Surface indentations can also be captured.

replace this iterative scheme by estimating the parameters interactively marking a small object and background region in one of the images. This is a requirement for the energy to be globally minimizable. Minimization of the first two terms in (7) results in the most probable surface with respect to the probability distributions P_{obj} and P_{bck} .

The last term in (7)

$$E_{stereo}(S) = \int_S \rho(x) dA \quad (11)$$

accounts for photoconsistency and smoothness of the sought surface. It is particularly important in order to reconstruct concavities that are not visible from the silhouettes; see Figure 1 for an illustration of the conceptual difference between the silhouette- and stereo-based constraints. Computation of ρ requires visibility estimation. To this end, we minimize the energy with Euclidean regularizer $\rho(x) = 1$. From the resulting surface, one can compute a signed distance function $\phi : V \rightarrow \mathbb{R}$, which in turn allows for normal estimation $N_x = \frac{\nabla \phi}{|\nabla \phi|}$ to each voxel $x \in V$. Hence, visibility is determined by front-facing cameras according to the estimated normal direction. The term (11) is equivalent to the one suggested in [12]. In particular, photoconsistency is computed in terms of the normalized cross-correlations by averaging over front-facing cameras

$$c(x) = \frac{1}{N} \sum_i \sum_j NCC(I_i(\pi_i(x)), I_j(\pi_j(x))), \quad (12)$$

where N denotes the number of relevant camera pairs. In order to take patch distortion into account, the surface is locally approximated by its tangent plane [12]. For each point $x \in V$ this yields some measure $c(x)$ between -1 and 1, where 1 means perfect correlation. This value is then mapped to the unit interval $[0, 1]$

using the following function proposed in [26]:

$$\rho(x) = 1 - \exp\left(-\tan\left(\frac{\pi}{4}(c(x) - 1)\right)^2 / \sigma^2\right). \quad (13)$$

Smoothness is implicitly enforced since minimizing (11) corresponds to finding the minimal surface with respect to a Riemannian metric [5]. Note that global optimization of this energy alone yields the empty surface. In our energy (7), the silhouette-based terms naturally prevent the empty surface without requiring any user interaction or additional knowledge about the scene.

4 Continuous Global Optimization

4.1 An Equivalent Convex Formulation

Energy (7) can be globally optimized, provided the object and background parameters of the Gaussian distribution and the visibility of points are given. In this paper we build upon the optimization technique described in section 2 by formulating (7) as a continuous convex optimization problem.

To this end, the surface S is represented implicitly by the characteristic function $u : V \rightarrow \{0, 1\}$ of R_{bck}^S , i. e. $u = \mathbf{1}_{R_{bck}^S}$ and $1 - u = \mathbf{1}_{R_{obj}^S}$. Hence, changes in the topology of S are handled automatically without reparametrization. With the implicit surface representation we have the following constrained, non-convex energy minimization problem corresponding to (7):

$$\begin{aligned} E(u) &= \int_V (\log P_{obj}(x) - \log P_{bck}(x))u(x) dx + \nu \int_V \rho(x)|\nabla u| dx \rightarrow \min, \\ \text{s. t. } u &\in \{0, 1\}. \end{aligned} \quad (14)$$

The minimization problem stated in (14) is non-convex, since the optimization is carried out over a non-convex set of binary functions. However, omitting the binary condition and extending the optimization to all functions $u : V \rightarrow \mathbb{R}$ where also intermediate values can be taken, will cause the values of $u(x)$ to converge to $\pm\infty$ almost everywhere. In order to circumvent this difficulty, one can restrict the domain by enforcing $0 \leq u(x) \leq 1$ via a convex penalizer $\theta(u)$:

$$E(u) = \int_V \left(\log P_{obj}(x) - \log P_{bck}(x) \right) u(x) + \nu \rho(x) |\nabla u| + \alpha \theta(u(x)) dx, \quad (15)$$

where α has to be chosen sufficiently large in order to ensure that u does not leave the interval $[0, 1]$. This leads to a convex formulation, which allows for global optimization by using standard techniques like gradient descent. Finally, we come up with a global minimizer of the original non-convex functional (14) by thresholding the result at any $\mu \in (0, 1)$. In our experiments, we chose $\mu = 0.5$, but we obtained virtually the same results with $\mu \in [0.1, 0.9]$.

In summary, the optimization can be split into two steps:

1. Find a minimizer u of (15).
2. Threshold the result: $R_{obj}^S = \{x \in V \mid u(x) < \mu \text{ for some } \mu \in (0, 1)\}$.

A necessary condition for a minimum of (15) is stated by the associated Euler-Lagrange equation

$$\begin{aligned} 0 &= (\log P_{obj} - \log P_{bck}) - \nu \rho \operatorname{div} \left(\frac{\nabla u}{|\nabla u|} \right) - \langle \nabla \rho, \frac{\nabla u}{|\nabla u|} \rangle + \alpha \theta'_\epsilon(u) \\ &= (\log P_{obj} - \log P_{bck}) - \nu \operatorname{div} \left(\rho \frac{\nabla u}{|\nabla u|} \right) + \alpha \theta'_\epsilon(u), \end{aligned} \quad (16)$$

where θ_ϵ is a regularized version of the derivative of θ with respect to its argument.

4.2 Fast Minimization by Successive Overrelaxation

Discretization of the Euler-Lagrange equation (16) leads to a sparse nonlinear system of equations, which can be solved via gradient descent. However, gradient descent converges very slowly. Thus, we suggest to use a fixed point iteration scheme that transforms the nonlinear system into a sequence of linear systems. These can be efficiently solved with iterative solvers, such as Gauss-Seidel, successive over-relaxation (SOR), or even multi-grid methods.

Neglecting the term $\alpha \theta'_\epsilon(u)$, which can in practice be replaced by simply clipping values of u that fall out of the interval $[0, 1]$, the only source of nonlinearity in (16) is the diffusivity $g := \frac{1}{|\nabla u|}$. Starting with an initialization $u^0 = 0.5$, we can compute g and keep it constant. For constant g , (16) yields a linear system of equations, which we solve with the SOR method. This means, we iteratively compute an update of u at voxel i by

$$u_i^{l,k+1} = (1 - \omega) u_i^{l,k} + \omega \frac{\nu \sum_{j \in \mathcal{N}(i), j < i} \rho_j g_{i \sim j}^l u_j^{l,k+1} + \nu \sum_{j \in \mathcal{N}(i), j > i} \rho_j g_{i \sim j}^l u_j^{l,k} - b_i}{\nu \sum_{j \in \mathcal{N}(i)} \rho_j g_{i \sim j}^l} \quad (17)$$

where $\mathcal{N}(i)$ denotes the neighborhood of i , $g_{i \sim j}$ denotes the diffusivity between voxel i and its neighbor j , and the vector b_i contains the constant part of (16) that does not depend on u , i.e. the fidelity term $b_i = \log P_{obj,i} - \log P_{bck,i}$. The over-relaxation parameter ω has to be chosen in the interval $(0, 2)$ for the method to converge. The optimal value depends on the linear system to be solved. Empirically we obtained the fastest convergence rate for $\omega = 1.85$. After being sufficiently close to a fixed point u^l (we iterated for $k = 1, \dots, 10$), one can update the diffusivities and solve the next linear system. Iterations are stopped as soon as the energy decay in one iteration is in the area of number precision.

5 Experiments

Figure 2 depicts 3 of 33 input images of resolution 640×480 used for reconstruction. The input images are pretty challenging because of the presence of



Fig. 2. Some of the input images used for 3D reconstruction. The image sequence is pretty challenging because of the presence of reflections and illumination artefacts.

illumination artefacts and specular reflections. Note that automatic color-based segmentation of the single images is infeasible due to the similarity in color of the bunny figure and the illumination effects in the background.

Figure 3 shows reconstructions from the above image sequence by using discrete graph cuts and the proposed continuous optimization technique applied on the model described in section 3. Both reconstructions look accurate. However, the graph cut reconstruction looks generally slightly oversmoothed because of the discrete approximation of the smoothness term. In addition, the proposed minimization exhibits considerable reductions in memory compared to graph cuts (in our implementation about a factor of 20), which allows to perform global optimization at higher volume resolutions. We ran the proposed optimization on an architecture with 2 GB of main memory and volume of more than 20 million voxels (see Figure 4). The corresponding graph cut computation is infeasible for this resolution.

The evolution of an initial surface towards the final result is depicted in Figure 5. Note that the final reconstruction does not depend on the initialization, since the used cost function is minimized globally. A closer look at the evolution process reveals the difference to local optimization techniques like level sets, where the surface evolves coherently, i.e. there are no unnecessary topological changes.

Figure 6 additionally emphasizes a comparison between graph cuts and the proposed continuous optimization when applied on a synthetic sphere with a missing piece of data. At such locations the difference between both models becomes obvious. Note that some discretization artefacts in terms of blocky structures are available in the graph cut reconstruction because of metrication errors, even with 26-neighborhood system. In addition, sharp corners occur, since the discrete model does not take the curvature of the surface into account. In contrast, the continuous optimization achieves nice and smooth continuation of the missing part of the surface.

6 Summary

In this paper an energy model for multiview 3D reconstruction allowing global optimization is proposed. To the best of our knowledge this is the first work to




Fig. 3. Multiview reconstruction of the sequence in figure 2. **First row:** reconstruction with the proposed method. **second row:** reconstruction obtained by minimizing the same energy functional (7) via graph cuts. Volume resolution was set to $108 \times 144 \times 162$. At this resolution both reconstructions look quite similar.




Fig. 4. Reconstruction obtained by the proposed approach at a volume resolution of $216 \times 288 \times 324$. Increasing the resolution by a factor of 2 in each dimension allows for the emergence of fine-scale details (compare to Figure 3). Graph cut reconstruction at such a resolution was infeasible on our machines due to memory overflow.




Fig. 5. Surface evolution towards the final result. Intermediate surfaces were generated by thresholding the evolving function u at 0.5 (see section 4). In contrast to level set schemes the evolution process is not coherent.




Fig. 6. Continuous vs. discrete optimization. (a) A slice through the data volume. Increasing intensities denote regions with $P_{obj}(x) > P_{bck}(x)$, $P_{obj}(x) < P_{bck}(x)$ and $P_{obj}(x) = P_{bck}(x)$ respectively. Photoconsistency function ρ is constant throughout the volume. (b) Reconstruction obtained with the optimization technique described in section 4. (c) Reconstruction computed by graph cuts. In contrast to the graph cut solution, the proposed continuous optimization does not suffer from discretization artefacts.

cast multiview 3D reconstruction as a continuous convex optimization problem (up to visibility). As for graph cuts this allows to compute globally optimal shapes. However, in contrast to discrete techniques, the proposed continuous formulation does not suffer from metrication errors. Moreover, it requires considerably less memory, thereby allowing for optimal reconstructions at higher resolutions. All these properties are demonstrated experimentally.

Acknowledgments

We thank Reinhard Klein and his group for helping us with the data acquisition.

References

1. B. Appleton and H. Talbot. Globally optimal geodesic active contours. *J. Math. Imaging Vis.*, 23(1):67–86, 2005.
2. B. Appleton and H. Talbot. Globally minimal surfaces by continuous maximal flows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(1):106–118, 2006.
3. Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001.
4. X. Bresson, S. Esedoğlu, P. Vanderghenst, J. P. Thiran, and S. Osher. Global minimizers of the active contour/snake model. Technical Report CAM-05-04, Department of Mathematics, University of California at Los Angeles, CA, U.S.A., Jan. 2005.
5. V. Caselles, F. Catté, T. Coll, and F. Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66:1–31, 1993.
6. V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *Proc. Fifth International Conference on Computer Vision*, pages 694–699, Cambridge, MA, June 1995. IEEE Computer Society Press.
7. A. Chambolle. Total variation minimization and a class of binary MRF models. In *EMMCVPR*, pages 136–152, 2005.
8. T. Chan, S. Esedoğlu, and M. Nikolova. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Journal on Applied Mathematics*, 66(5):1632–1648, 2006.
9. T. Chan and L. Vese. An active contour model without edges. In M. Nielsen, P. Johansen, O. F. Olsen, and J. Weickert, editors, *Scale-Space Theories in Computer Vision*, volume 1682 of *LNCS*, pages 141–151. Springer, 1999.
10. Y. Duan, L. Yang, H. Qin, and D. Samaras. Shape reconstruction from 3D and 2D data using PDE-based deformable surfaces. In *Proc. European Conference on Computer Vision*, pages 238–251, 2004.
11. C. H. Esteban and F. Schmitt. Silhouette and stereo fusion for 3D object modeling. *Computer Vision and Image Understanding*, 96(3):367–392, 2004.
12. 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, Mar. 1998.
13. T. C. Hu. *Integer Programming and Network Flows*. Addison-Wesley, Reading, MA, 1969.
14. M. Kass and A. Witkin. Analyzing oriented patterns. *Computer Vision, Graphics and Image Processing*, 37:362–385, 1987.
15. K. Kolev, T. Brox, and D. Cremers. Robust variational segmentation of 3D objects from multiple views. In K. F. et al., editor, *Pattern Recognition (Proc. DAGM)*, volume 4174 of *LNCS*, pages 688–697, Berlin, Germany, September 2006. Springer.
16. K. N. Kutulakos and S. M. Seitz. A theory of shape by space carving. *International Journal of Computer Vision*, 38(3):199–218, 2000.
17. A. Laurentini. The visual hull concept for visual-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):150–162, 1994.
18. V. Lempitsky, Y. Boykov, and D. Ivanov. Oriented visibility for multiview reconstruction. In *Proc. European Conference on Computer Vision*, volume 3953 of *LNCS*, pages 226–238, 2006.
19. W. N. Martin and J. K. Aggarwal. Volumetric descriptions of objects from multiple views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5(2):150–158, 1983.

20. D. Mumford and J. Shah. Boundary detection by minimizing functionals, I. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 22–26, San Francisco, CA, June 1985. IEEE Computer Society Press.
21. L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
22. S. Sinha and M. Pollefeys. Multi-view reconstruction using photo-consistency and exact silhouette constraints: A maximum-flow formulation. In *Proc. International Conference on Computer Vision*, pages 349–356, Washington, DC, USA, 2005. IEEE Computer Society.
23. D. Snow, P. Viola, and R. Zabih. Exact voxel occupancy with graph cuts. In *Proc. International Conference on Computer Vision and Pattern Recognition*, volume 1, pages 345–353, 2000.
24. G. Strang. Maximal flow through a domain. *Mathematical Programming*, 26:123–243, 1983.
25. S. Tran and L. Davis. 3D surface reconstruction using graph cuts with surface constraints. In *Proc. European Conference on Computer Vision*, volume 3952 of *LNCS*, pages 219–231, 2006.
26. G. Vogiatzis, P. Torr, and R. Cippola. Multi-view stereo via volumetric graph-cuts. In *Proc. International Conference on Computer Vision and Pattern Recognition*, pages 391–399, 2005.
27. A. Yezzi and S. Soatto. Stereoscopic segmentation. *International Journal of Computer Vision*, 53(1):31–43, 2003.