

From Photohulls to Photoflux Optimization

Yuri Boykov

*Computer Science Department
University of Western Ontario
London, ON, Canada*

Victor Lempitsky

*Department of Mathematics
Moscow State University
Moscow, Russia*

Abstract

Our work was inspired by recent advances in image segmentation where flux-based functionals significantly improved alignment of object boundaries. We propose a novel *photoflux* functional for multi-view 3D reconstruction that is closely related to properties of photohulls. Our *photohull prior* can be combined with regularization. Thus, this work unifies two major groups of multiview stereo techniques: “space carving” and “deformable models”. Our approach combines benefits of both groups and allows to recover fine shape details without oversmoothing while robustly handling noise. Photoflux provides data-driven ballooning force that helps to segment thin structures or holes. Photoflux maximizing shapes can be also seen as regularized Laplacian zero-crossings [3]. We discuss several versions of photoflux functional based on global, local, or non-deterministic visibility models. Some forms of photoflux can be easily added into standard regularization techniques. For other forms we propose new optimization methods.

1 Introduction

The reconstruction of 3D shape from a set of its views, *multiview reconstruction*, is one of the fundamental problems in computer vision. In this paper, we assume that all views are registered within the global coordinate system, such that for any 3D point $X \in \mathcal{R}^3$ it is possible to find its projection onto each of the views.

The majority of (volumetric) multiview shape-from-photoconsistency methods fall into two categories: greedy space carving approaches aiming at recovering maximal photoconsistent shapes, photohulls, and energy-based methods searching for some regularized surface with good photoconsistency (see the bottom row in the table below).

	Greedy methods	Minimal/deformable surfaces	Flux-based methods
Image segmentation	Thresholding [18] Region growing [1]	Snakes [8] Level-sets [15] Graph cuts [4]	Level-sets [21] Graph cuts [11]
Multiview reconstruction (volumetric)	Voxel coloring [19] Space carving [13]	Mesh-based [6] Level-sets [7, 20, 16] Graph cuts [22, 14]	This work

Normally, energy-based methods achieve better surface quality by adding regularization that removes noise. However, their explicit bias towards minimal surfaces may have an

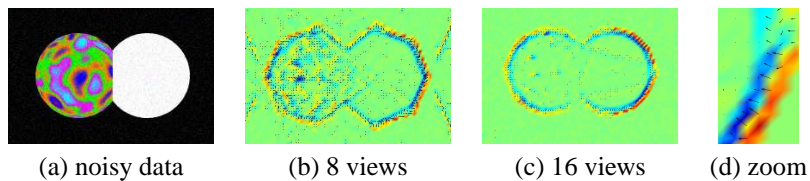


Figure 1: Example of *photoconsistency flow* vector field (Sec.2.3) in a 3D volume. One camera view of a 3D scene is shown in (a). The spectrum of colors in (b-d) indicates divergence of the vector field. Green color corresponds to zero divergence. Pictures (b-d) illustrate vector fields based on “average” photoconsistency gradients (7).

oversmoothing effect. In contrast, greedy methods like space-carving do not have this minimality bias which helps them to outperform energy-based methods on the scenes with thin, protruding, and/or textureless parts.

This work is largely inspired by the recent progress in image segmentation, which is another low-level vision problem that in many aspects is similar to volumetric multi-view reconstruction of shapes (see the upper row in the table above). The field of image segmentation was once dominated by greedy methods (e.g. thresholding or region growing) that were eventually replaced by regularized minimal/deformable surface approaches. However, standard minimum surface methods for image segmentation tend to oversmooth thin structures and undersegment/oversegment blurred boundaries which is very similar to problems of current energy-based multiview reconstruction methods.

Recent results in image segmentation suggest that incorporating flux-based functionals into local and global regularization methods such as level-sets [21, 9] and graph cuts [11] significantly improves segmentation of thin elongated structures and alignment of segments with object boundaries. These approaches advocate energy functionals combining Riemannian area/length with flux of a vector field normally computed from the gradients of a gray-scale image.

This paper proposes to integrate flux optimization into regularization-based methods for multiview reconstruction. We show that flux of vector fields computed from the gradients of photoconsistency function (*photoflux*) introduces a data-driven ballooning force biasing reconstructed shapes towards maximal photoconsisten surface (photohull). This bias helps to recover textureless (Figure 1) and thin (Figure 5(d)) parts of objects. Photoflux can be easily integrated into the majority of regularization-based methods for multiview reconstruction and this paper discusses a number techniques that can use either local or global optimization methods such as level-sets or graph cuts. We test-prove the concept of photoflux by experimentally comparing results with and without photoflux.

2 Flux for multiview stereo

One motivation for using *flux* in multiview stereo comes from its relationship with a well-known concept of *photohulls* [13]. Photohulls are based on binary “yes” or “no” decisions about photoconsistency of any given point on a current surface. The output of greedy space carving methods, *photohull*, depends on a threshold for making such hard decisions. Like “region growing” is sensitive to “leaking” through weak spots on object boundary, space carving is not robust to noise or specularities on object surface that will often result

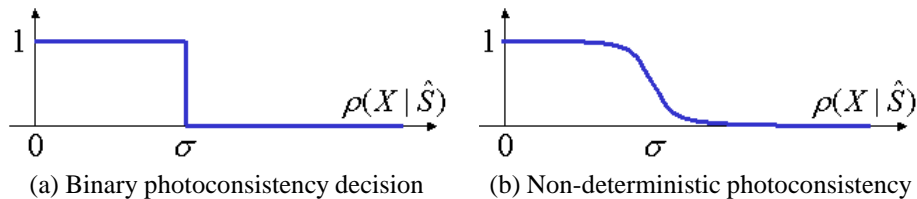


Figure 2: Photoconsistency decision functions used to compute photohulls (a) and photoflux (b). The x -axis shows possible values of inter-camera color variance ρ . Plot (b) represents a non-deterministic strategy that can be seen as a smoothed version of the binary decision function (a). Strategy (b) may better account for noise and outliers in the data. Equation (1) is one example of a differentiable photoconsistency function as in (b).

in “leaking” and eventual cascading erosion of the shape.

However, space carving has some noticeable advantages over existing minimal surface methods for multiview stereo. Space carving can reconstruct thin objects or fine protruding details that previous energy based methods tend to oversmooth. We propose a flux-based surface functional that allows to integrate a bias to photohull into regularization methods for multi-view stereo. This combines the benefits of two standard groups of multiview reconstruction methods (see table above): regularization provides noise robustness while flux counteracts shrinking bias. In the next subsections we formulate our *photoflux* functional motivating it by an argument showing that a surface of a photohull should have large flux of photoconsistency gradients.

2.1 Photoflux and global visibility

For now, assume a fixed surface \hat{S} . Let $\rho(X|\hat{S})$ be some *color variance* function measuring the deviation between the colors I_X^1, \dots, I_X^m of a point X in different views given visibility of X defined by the shape \hat{S} (for example, this could be a current solution in an iterative algorithm). Some examples of color variance measures are mentioned in Section 3.2. Under a Gaussian noise assumption, the likelihood of a surface at point X to be photoconsistent with the images is

$$P(X|\hat{S}) = \Pr(I_X^1, \dots, I_X^m | X \in \hat{S}) \propto \exp\left(-\frac{\rho(X|\hat{S})}{2\sigma^2}\right) \quad (1)$$

where σ is the noise variance. We will use this likelihood function throughout the paper. We will simply refer to (1) as a *photoconsistency function*.

In contrast, space carving algorithm uses a binary photoconsistency function

$$P(X|\hat{S}) = \begin{cases} 1 & \text{if } \rho(X|\hat{S}) \leq \sigma \\ 0 & \text{if } \rho(X|\hat{S}) > \sigma \end{cases}$$

where σ is a photoconsistency threshold (compare plots in Figure 2). The algorithm can be regarded as a successive carving of inconsistent voxels X such that $P(X|\hat{S}) = 0$ and a simultaneous update of \hat{S} . Thus, space carving algorithm converges to shape S (photohull) that is a boundary between points with $P(X|S) = 1$ and points with $P(X|S) = 0$.

In the context of our non-deterministic photoconsistency function (1) as in Figure 2(b) photohull’s boundary can be characterized by a large gradient of $P(X|S)$ directed from

exterior of photohull to its interior. Based on this observation, we propose the following functional for energy-based multiview reconstruction

$$\text{PHOTOFLUX I : } \quad F(S) = - \int_S \langle \nabla P(X|S), N_X \rangle dS \quad (2)$$

where N_X is an outward looking unit normal of surface S at point X , and $\langle \cdot, \cdot \rangle$ is a scalar product that does not have to be necessarily Euclidean. Functional (2) corresponds to flux of photoconsistency gradients through the surface. To be short, we call this functional a *photoflux* through an oriented surface S ¹.

Remember that reconstruction of photohulls using space carving algorithms does not include any regularization. Analogously, optimization of flux alone does not provide any regularization either [11]. Divergence theorem suggest that flux can be seen as an “intelligent” regional balloon force. To add regularization, we augment the functional (2) with an area surface term. This term can correspond either to Euclidean area $\int_S dS$ or to a photoconsistency-based area, used in other energy-based methods $\int_S \rho(X|S) dS$.

2.2 Photoflux and local visibility

In case of convex shapes we have $P(X|S) = P(X|N_{dS})$ where N_{dS} is an outward normal of a surface patch containing point X and

$$P(X|N) = \Pr(I_X^1, \dots, I_X^m | X, N) \propto \exp\left(-\frac{\rho(X|N)}{2\sigma^2}\right) \quad (3)$$

is a photoconsistency of point X based on its *local visibility* [14] defined by a given surface normal N . Clearly, photoconsistency $P(X|N)$ should be computed from cameras located in the half-space defined by a tangent plane at point X and its outward normal N .

Thus, convexity of shape S implies $\nabla P(X|S) = \nabla P(X|N_X)$ and photoflux of such shape can be written as

$$\text{PHOTOFLUX II : } \quad F(S) = - \int_S \langle \nabla P(X|N_X), N_X \rangle dS. \quad (4)$$

For general non-convex shapes S equation (4) is an approximation of the definition of photoflux in (2). In case of occlusions due to non-convexity of shape (e.g. in case of multiple objects) we have $P(X|S) \geq P(X|N_X)$ since some cameras used to compute photoconsistency based on local visibility contribute an error. Our experiments show even local visibility approximations like (4) can still handle scenes with significant amount of such occlusions.

2.3 Photoflux and non-deterministic visibility

Note that the earlier two definitions of *photoflux* are based on vector fields of gradients $\nabla P(X|S)$ or $\nabla P(X|N)$. These vectors are defined only at points X on some given surface S or on a given local patch with normal N . The surface or a patch must be fixed in order to estimate visibility of point X before photoconsistency $P(X|S)$ or $P(X|N)$ can be computed. We also assume that small perturbations of object shape near X do not change visibility so that the gradients of photoconsistency can be evaluated at X .

¹The orientation of S is defined by a field of its “outward” normals.

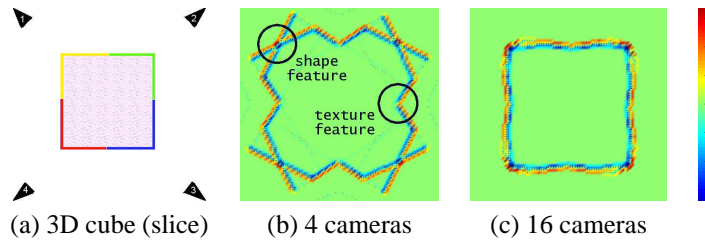


Figure 3: Reconstruction of a cube. Divergence of estimated photoconsistency gradients is shown in (b) and (c). The images of divergence demonstrate “apparent contours” of shape features (sharp corners) or surface texture features.

This section provides a more general definition of photoflux. Instead of computing visibility of point X deterministically for a given surface or a patch, we assume some probability distribution $\Pr(W|X)$ for possible values of visibility W at any given point X . Note that visibility $W = (w_1, \dots, w_m)$ is a vector of non-negative weights $w_k \geq 0$ describing correlation between intensity I_X^k of a pixel in camera k where point $X \in \mathcal{R}^3$ projects and a color of point X assuming it belongs to some surface.

The most general definition of photoflux in this paper is motivated by flux-based shape priors for image segmentation [11]. Our “photohull shape prior” is based on

$$\text{PHOTOFLUX III : } F(S) = - \int_S \langle v_X, N_X \rangle dS. \quad (5)$$

where a field of vectors $\{v_X\}$ describes *photoconsistency flow*, that is, direction and rate of increase in photoconsistency at every point X . One way to obtain such vectors is to compute expected gradients of photoconsistency by averaging unknown visibility out

$$v_X = \sum_W \nabla P(X|W) \cdot \Pr(W|X). \quad (6)$$

Functionals (2) and (4) are special cases of (5) when visibility of X is deterministic.

We also suggest a simpler but more practical version of photoconsistency flow

$$v_X = \sum_N \nabla P(X|N) \cdot \Pr(N|X) \quad (7)$$

where visibility is represented by a possible outward normal direction N . A possible heuristic for estimating probability $\Pr(N|X)$ in practice could be based on gradients of photoconsistency

$$\Pr(N|X) \propto \langle \nabla P(X|N), N \rangle^+ \quad (8)$$

where $\langle \cdot, \cdot \rangle^+$ is a positive part of the dot product. Figures 1(b-d) and 3 demonstrate vector fields computed according to equations (7) and (8).

Equation (8) requires a normalization multiplier to make it a true probability distribution. In practice, however, this may cause “division by zero” problems which can be avoided by truncating a vector field where the multiplier gets too small. This “division by zero” corresponds to points where surface is unlikely to pass and “visibility” is ill-defined. Practically speaking, equation (8) without any normalization generates very similar photoconsistency flows $\{v_X\}$. This heuristic is also easier to implement since no “division by zero” numerical problems arise.

In general, there are other options for computing vectors v_X for photohull prior (5). For example, we also experimented with vector fields

$$v_X = \sum_N N \cdot \Pr(N|X). \quad (9)$$

shown in examples of Figure 1)(e-h). Vector fields could also be obtained from approximate and/or sparse disparity maps. Such maps can be computed at each camera by evaluating correspondences with nearby cameras using narrow-base stereo methods or robust feature matching.

3 Properties of photoflux optimization

The definitions of photoflux in (2) and (4) use deterministic visibility and the corresponding vector fields are defined only on a surface of any fixed shape. More general form of photoflux (5) allows non-deterministic visibility and the corresponding vector field can be estimated at any point $X \in \mathcal{R}^3$. Despite these obvious distinctions, there are many important common properties of different forms of photoflux which we describe in this section. For simplicity, we discuss these properties in the context of equations (5) and (7).

3.1 Combining *photoflux* and regularization

Standard *space carving* methods for photohull reconstruction are analogous to adaptive region-growing techniques in image segmentation and they do not include any surface regularization. Section 2.1 “derived” photoflux functional from properties of photohulls. Thus, it should not be surprising that photoflux maximization itself does not imply any regularization. As discussed in [11], flux optimization is equivalent to thresholding a divergence $\text{div}(v)$ of the corresponding vector field v .

In the context of Figure 3, “thresholding” implies that all voxels of blue color (negative divergence values) should be selected as surface interior. It is clear, however, that this approach would generate noisy and/or ambiguous results. Therefore, we propose to augment the photoflux functional (2) or (5) with regularization term. For example, we can combine photoflux with an integral of photoconsistency over the whole surface obtaining a regularized photoflux functional

$$E(S) = - \int_S \langle \nabla P(X|N_X), N_X \rangle dS + \lambda \cdot \int_S \rho(X|N_X) dS. \quad (10)$$

Alternatively, it might also be possible to regularize photoflux via Euclidean area of the surface S . It is also possible to use non-Euclidean dot products in the definition of photoflux that may be based on photoconsistency.

Optimization of an energy combining photoflux and regularization as in (10) should reconstruct fine object details while handling noise. Photoflux can be seen as an intelligent balloon force making minimal surfaces better align with thin protrusions or concavities. Combining flux and regularization has similar effects in image segmentation [21, 9, 11].

Integrating flux of image gradients into image segmentation energy results in better alignment with object boundaries (edges). In particular, [9] demonstrated that regularized flux optimization can be regarded as a regularized Laplacian zero-crossing of image intensity. Analogously, optimization of photoflux can be regarded as a regularized Laplacian zero-crossing of a “photohull likelyhood” function (for details, see [3]).

3.2 Integrating *apparent contours* and texture features

This subsection demonstrates that photoflux based multi-view reconstruction implicitly relies on *apparent contours* and on detecting object textural features. To be specific, assume that color variance is measured by function

$$\rho(X|N_X) = \sum_k w_k \cdot (I_X^k - \bar{I}_X)^2, \quad \bar{I}_X = \sum_k w_k \cdot I_X^k \quad (11)$$

or by

$$\rho(X|N_X) = \sum_{i,j} w_i \cdot w_j \cdot (I_X^i - I_X^j)^2 \quad (12)$$

where I_X^k is intensity of a pixel in camera k where point $X \in \mathcal{R}^3$ projects, \bar{I}_X is a weighted average of such intensities, and w_k is a non-negative weight describing visibility of point X in camera k . Typically, visibility w_k is large for a camera observing surface point X with normal N_X under a small angle and w_k approaches zero for angles near 90 degrees. Note that color consistency penalty as in (11) was proposed for 3D reconstruction in [20] and functionals like (12) were proposed in [16] (image-based regularization).

For both (11) and (12), differentiation of photoconsistency $P(X|N_X)$ in (1) gives

$$\nabla P(X|N_X) \propto P(X|N_X) \cdot \sum_k w_k \cdot (\bar{I}_X - I_X^k) \cdot \nabla I_X^k \quad (13)$$

where ∇I_X^k is a gradient of an image in camera k at a pixel where 3D point X projects. We treat ∇I_X^k as a 3D vector since k -th camera image plane orientation in \mathcal{R}^3 is known. Expressions analogous to (13) can be derived for other color variance models.

Equation (13) breaks down the gradient of photoconsistency into a linear combination of vectors contributed by individual cameras where point X is visible, that is $w_k > 0$. It is not surprising that each camera contribution is based on its image gradient ∇I_X^k . For example, a small number of cameras in Figure 3(b) allows to see straight “rays” of points with large gradients of photoconsistency. Each “ray” in Figure 3(b) is formed by points $X \in \mathcal{R}^3$ projecting onto the same pixel in one of the cameras where image gradient is large. Note that multiplication by a factor $P(X|N_X)$ in (13) discards gradients in all points $X \in \mathcal{R}^3$ that are not photoconsistent.

Computing a gradient of photoconsistency automatically obtains a “good” linear combination of image gradients (13) appropriate for estimating surface normals. Interestingly, each camera k individually does not know how to interpret its image gradients. In general, there are two possible reasons for a large image gradient at any given pixel; the camera could observe either a boundary between an object and a background (*apparent contour*) or a textural detail on a surface. Figure 3(b) demonstrates that rays of gradients from different cameras intersect at points that align with an object surface either at its *shape features* or at its *texture features*. Note that by computing gradients of photoconsistency we automatically detect both types of features and use them as surface “anchors”. Larger number of cameras, e.g. in Figure 3(c), gives tighter alignment of large photoconsistency gradients with such anchors producing a close outline of the true object surface.

4 Efficient algorithms

The photoflux optimization can be plugged into virtually any energy optimization framework. In particular, level-sets [7, 16] and banded graph cuts [22] reconstructions can

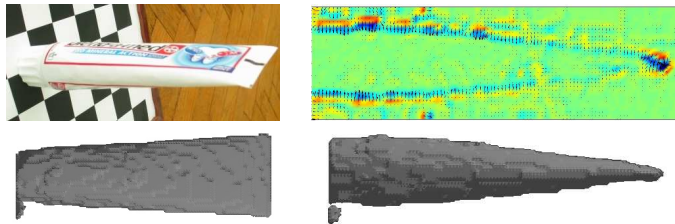


Figure 4: Results of our method for the *toothpaste* sequence. Note the sharp top of the tube. The surface also exhibits significant non-lambertian reflectivity effects. Top-right is the profile slice of a vector field. A strong photoflux vector field appears in the most difficult part (top) of the tube. Bottom row - two views of the reconstruction.

include the photoflux optimization either in the deterministic form (2) or in the probabilistic form (5). In the former case, the current solution defines the shape estimate \hat{S} determining global visibility.

Local visibility photoflux in (4) allows to use global optimization methods using explicit graph cuts on a complex [10, 14]. However, the corresponding energy is not guaranteed to be submodular. Thus, it may be necessary to truncate supermodular terms as in [17] (see our results in [3]), or to use QPBO approach to non-submodular function [2] that was first proposed for applications in computer vision by Kolmogorov et al. [12].

Non-deterministic approach to photoflux (5) can be implemented by global optimization methods based on either implicit surface representation [11] or complex-based graph cuts [10]. It is also interesting to investigate global methods that update estimate of visibility once parts of a surface are recovered in optimization methods that generate partial solutions.

Another possibility of visibility handling is the iterative estimation of visibility for photoflux which can be achieved via local improvement methods (e.g. level-sets [16], banded cuts [22]), or via discrete algorithms demonstrating intermediate solutions. It would be also interesting to compare the results of explicit graph cuts on complexes [14] with implementing regularized photoflux via implicit graph cuts techniques [4, 11].

5 Experimental Results

In order to validate the basic concept of photoflux, this paper presents our results for *complex*-based multiview reconstruction [14] after integrating photoflux into it (more results and implementational details can be found in [3]). In the absence of any scene approximation, [14] treats visibility and photoconsistency $P(X|N)$ of a patch only as functions of its position and orientation. This technique allows to obtain the global minimum of a discretized version of regularized photoflux functionals as in (10). In particular, we tested two forms of photoflux (4) and (5).

To construct a complex, we subdivide the approximate bounding box with 9 families of parallel planes. As a result, we get a *CW-complex* comprising polyhedral cells and the facets separating cells from each other. For each facet, we distinguish between two faces of this facet having opposite orientations. For every oriented facet F a regularized cost (10) is computed based on the facet' area, position, and normal orientation. The minimum

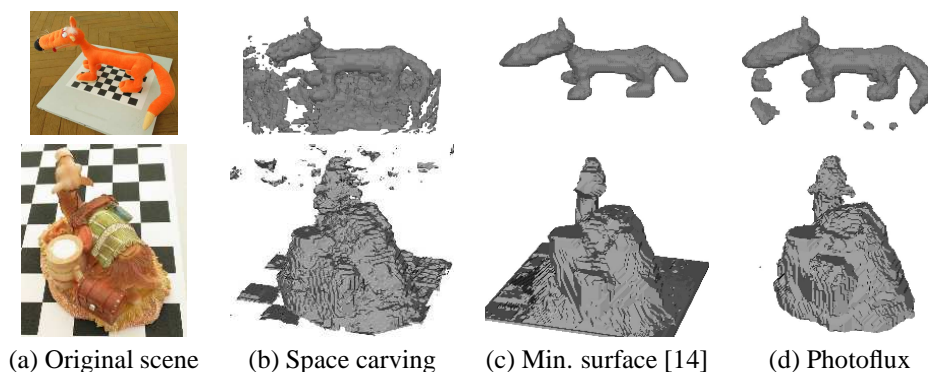


Figure 5: Comparison of different methods for multiview reconstruction on the fox and camel sequence (similar photoconsistency measure is used for all three algorithms). Typically, space carving generates noisy results while minimum surface methods oversmooth the shape. Adding maximization of photoflux into energy-based methods for 3D reconstruction allows to accurately restore thin protrusions and other fine details of the shape (ears of camel and a long tail of the fox).




Figure 6: Results of our method on standard *gargoyle* and *hand* sequences (courtesy of Prof. Kutulakos). Note the complex topology of the gargoyle and the weak texture of the hand. Please, refer to other publications (e.g. [13, 16]) using these sequences for comparison.

of a regularized cost (10) over all shapes comprised from the cells of the CW-complex can be efficiently found using a variant of a min-cut algorithm [5] applied to the graph, which is dual to the complex.

Non-deterministic photoflux: Integrating functional (5) into complex-based regularization algorithm above is straightforward. One can estimate vector v_X at each facet and then introduce directed costs for an n-link corresponding to this facet as described in [11]. This work guarantees that the corresponding energy is submodular and that the global minima can be computed via max-flow/min-cut algorithms, e.g. [5]. The results are presented on the figures 4, 5, 6.

Local-visibility deterministic photoflux: Photoflux in (4) can also be integrated into complex based regularization framework since each facet on the complex has normal N and $\nabla P(X|N)$ can be evaluated when the complex is constructed. Please, refer to [3] for the details and experimental results in this case.

In our experiments, we use complexes with 5–10 million of cells, with a typical running time of few minutes on a P4-2.8 architecture.

We are working on iterative estimation of visibility for photoflux which can be achieved

via local improvement methods (e.g. level-sets [16], banded cuts [22]), or via discrete algorithms demonstrating intermediate solutions. It is also interesting to compare our results with explicit graph cuts on complexes [14] with implementing regularized photoflux via implicit graph cuts techniques [4, 11]. In general, there are also other methods for computing vector fields that can estimate photoconsistency flow or surface normals defining some flux-based shape bias that can help to overcome “shrinking” problem of standard regularization approach to multiview reconstruction.

References

- [1] R. Adams and L. Bischof. Seeded region growing. *IEEE Trans. PAMI*, 16(6):641–647, 1994.
- [2] E. Boros and P. L. Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1-3):155–225, 2002.
- [3] Y. Boykov and V. Lempitsky. Photoflux maximizing shapes. In *University of Western Ontario, Comp. Sci. Tech. Rep. 670, ISBN-13: 978-0-7714-2580-5*, May 3, 2006.
- [4] Yuri Boykov and Vladimir Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. In *ICCV*, volume I, pages 26–33, 2003.
- [5] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. PAMI*, 26(9):1124–1137, 2004.
- [6] Carlos Hernandez Esteban and Francis Schmitt. Silhouette and stereo fusion for 3d object modeling. *CVIU*, 96(3):367–392, 2004.
- [7] Olivier D. Faugeras and Renaud Keriven. Complete dense stereovision using level set methods. In *ECCV '98(1)*, pages 379–393, 1998.
- [8] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *IJCV*, 1(4):321–331, 1988.
- [9] R. Kimmel and A. M. Bruckstein. Regularized laplacian zero crossings as optimal edge integrators. *IJCV*, 53(3):225–243, 2003.
- [10] D. Kirasanov and S. J. Gortler. A discrete global minimization algorithm for continuous variational problems, 2004.
- [11] Vladimir Kolmogorov and Yuri Boykov. What metrics can be approximated by geo-cuts, or global optimization of length/area and flux. In *ICCV*, volume I, pages 564–571, 2005.
- [12] Vladimir Kolmogorov and Carsten Rother. Minimizing non-submodular functions with graph cuts - a review. *Microsoft Research, Cambridge, UK, tech. rep. MSR-TR-2006-100*, 2006 (under review in PAMI).
- [13] K. N. Kutulakos and S. M. Seitz. A theory of shape by space carving. In *IJCV*, volume 38(3), pages 199–218, 2002.
- [14] V. Lempitsky, Y. Boykov, and D. Ivanov. Oriented visibility for multiview reconstruction. In *ECCV*, volume III, pages pp.226–238, 2006.
- [15] R. Malladi, J. A. Sethian, and B. C. Vemuri. Evolutionary fronts for topology-independent shape modeling and recovery. In *ECCV*, pages 3–13, 1994.
- [16] Jean-Philippe Pons, Renaud Keriven, and Olivier Faugeras. Modelling dynamic scenes by registering multi-view image sequences. In *CVPR '05 (2)*, pages 822–827, 2005.
- [17] Carsten Rother, Sanjiv Kumar, Vladimir Kolmogorov, and Andrew Blake. Digital tapestry. In *CVPR*, volume I, 2005.
- [18] P. K. Sahoo, S. Soltani, A. K.C. Wong, and Y. C. Chen. A survey of thresholding techniques. *CVGIP*, 41(2):233–260, 1988.
- [19] Steven M. Seitz and Charles R. Dyer. Photorealistic scene reconstruction by voxel coloring. In *CVPR '97*, page 1067, 1997.
- [20] Stefano Soatto, Anthony J. Yezzi, and Hailin Jin. Tales of shape and radiance in multi-view stereo. In *ICCV '03*, 2003.
- [21] Alexander Vasilevskiy and Kaleem Siddiqi. Flux maximizing geometric flows. *IEEE Trans. PAMI*, 24(12):1565–1578, 2002.
- [22] G. Vogiatzis, P. H. S. Torr, and R. Cipolla. Multi-view stereo via volumetric graph-cuts. In *CVPR '05(2)*, 2005.