



HAL
open science

Propagated Photoconsistency and Convexity in Variational Multiview 3D Reconstruction

Kalin Kolev, Maria Klodt, Thomas Brox, Daniel Cremers

► **To cite this version:**

Kalin Kolev, Maria Klodt, Thomas Brox, Daniel Cremers. Propagated Photoconsistency and Convexity in Variational Multiview 3D Reconstruction. Proceedings of the First International Workshop on Photometric Analysis For Computer Vision - PACV 2007, Oct 2007, Rio de Janeiro, Brazil. 8 p. inria-00265725

HAL Id: inria-00265725

<https://inria.hal.science/inria-00265725>

Submitted on 20 Mar 2008

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.

Propagated Photoconsistency and Convexity in Variational Multiview 3D Reconstruction

Kalin Kolev, Maria Klodt, Thomas Brox, Daniel Cremers

Computer Vision Group, University of Bonn
Römerstr. 164, 53117 Bonn, Germany

{kolev,klodt,brox,dcremers}@cs.uni-bonn.de

Abstract

In this paper, we make two contributions. Firstly, we replace the generic balloon constraints widely used in 3D reconstruction by a more sophisticated data-dependent regional term. The key idea is to propagate classical photoconsistency along visual rays into regional values describing voxel probabilities for being inside or outside the observed object. Secondly, we cast the optimization as one of minimizing a convex functional. Therefore (up to visibility) the reconstruction problem can be solved in a globally optimal manner in a spatially continuous setting. Compared to graph cut methods, this approach does not suffer from discretization artifacts and exhibits considerable reduction in memory requirements. Experimental comparisons clearly show the advantages of the proposed technique.

1. Introduction

Recovering the geometrical structure of a scene from multiple views is one of the most fundamental and extensively studied problems in computer vision with numerous applications. Here we concentrate on algorithms operating on calibrated images, where the projection from 3D world coordinates to image coordinates is assumed to be known. There are two major classes of techniques according to the exploited information: shape from silhouettes and stereo.

In case of sparsely textured objects, silhouettes exhibit the dominant image feature. Most of the silhouette-based methods for reconstruction aim at approximating the *visual hull* [13] of the imaged object. The visual hull of an object can be described as the maximal shape that yields the same silhouettes as the actual object for all views. The earliest attempts use polyhedral [2], volumetric representations [15] or local surface models [4]. In a completely new approach [20] and [17] formulate an energy minimization framework for 3D region segmentation from a collection of images of a

scene using level sets and graph cuts, respectively. The authors couple the segmentations of each image through the evolution of a single 3D surface rather than separate 2D contours. This allows to impose regularization making their method robust to outliers and erroneous camera calibration. In [10] the robustness to noise is increased by incorporating all available image information into a probabilistic framework.

In fact silhouette-based methods cannot recover surface concavities as these do not affect the silhouettes. In order to reconstruct non-convex surface details, stereo-based techniques exploiting texture information are required. Here, one matches points or patches from different images that correspond to the same point in the scene. The earliest algorithms that incorporate a large number of views use carving techniques to obtain a volumetric representation of the scene assuming Lambertian properties of the objects [12]. Others have suggested to guide a deformable surface model by a measure based on photoconsistency between different views toward a steady state utilizing meshes [7], level sets [9], or graph cuts [14].

Some recent approaches use a fusion of silhouette constraints and photoconsistency; [5] combines them into a single cost function; [8] and [18] use silhouette points to constrain the computed surface based on stereo information. However, simultaneous use of silhouettes and photoconsistency may introduce a bias in the reconstruction near the visual hull. Thus, [19] proposed a two-phase approach. In the first phase the visual hull is computed, which is then refined in the second phase considering photoconsistency.

Most previous methods aiming at dense stereo-based 3D reconstruction involve integration of photoconsistency over the sought surface exploiting the similarity of the projections of points lying on it in images, where they are visible. Although local optimization methods can produce high-quality reconstructions [9, 8], their robustness is limited, since they are vulnerable to local minima. On the other

hand, global optimization of this model delivers the empty set as a solution [14]. Currently, this problem is approached by using an artificial ballooning term [19, 18, 14]. It can be seen as a prior that favors surfaces with larger volumes. However, this prior has a decisive limitation: setting the weighting of the expanding force to large values causes filling of concavities, whereas lower weighting leads to cutting of thin protrusive structures. Therefore, in this paper, we propose to replace the generic balloon constraint by a more sophisticated data-dependent regional term, which is obtained by propagating stereo information into the entire 3D space. This avoids the globally optimum solution to be the empty set without introducing a bias towards larger volumes. Upon a closer look, the presented model reveals some similarities with [6], however in a different context.

As a second contribution, we cast the optimization as one of minimizing a convex functional. Therefore (up to visibility) the reconstruction problem can be solved in a globally optimal manner while having a spatially continuous setting. The continuous formulation has two important advantages compared to alternative discrete graph cut approaches. Firstly, the solution is free of discretization artifacts. Secondly, the memory requirements are smaller allowing for reconstructions at higher resolutions.

The paper is organized as follows. In the next section the variational formulation is presented. Its conversion to a convex optimization problem and efficient implementation by SOR are explained in Section 3. In Section 4 experimental results demonstrate in particular that the proposed volumetric extension of stereo information compares favorably to the traditional ballooning approach. We conclude the paper with a brief summary in Section 5.

2. Variational Formulation

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. 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 $x \in V$ a photoconsistency value $\rho(x) \in [0, 1]$ describing the probability of a voxel for lying on the surface (figure 1(a)), based on its projections onto the images, where it is visible. Note that in the classical photoconsistency definition small values correspond to photoconsistent points. In a similar manner, we can compute values $\rho_{obj}(x), \rho_{bck}(x) \in [0, 1]$ describing probabilities of a point x to belong to R_{obj}^S and R_{bck}^S , respectively. We will refer to these functions as *interior* and *exterior* photoconsistency (figure 1(c)). Hence, we can for-

mulate the following energy minimization problem:

$$E(S) = \int_{R_{obj}^S} \rho_{obj}(x) dx + \int_{R_{bck}^S} \rho_{bck}(x) dx + \nu \int_S \rho(x) dx \quad (1)$$

$$\hat{S} = \arg \min_{S \subset V} E(S).$$

The first two terms of the functional impose correct subdivision of the volume into interior/exterior according to the respective photoconsistency functions. The last term acts as a constraint both for smoothness and photoconsistency by seeking the minimal surface with respect to a Riemannian metric. Hence, it can be considered as a weighted smoothness term. Note that the photoconsistency functions may also depend on the orientation of the surface estimate S in order to take distortion of the compared image patches into account. This dependency is suppressed here for simplicity.

In classical energy minimization methods using stereo, the regional terms are omitted [9] or replaced by an artificial ballooning term in order to avoid the empty surface as a solution [14] or to reconstruct protrusive details [19]. In the first case only local optimization is possible, since the global minimum is the empty surface. In the second case concavities are oversmoothed, since greater volumes are preferred [18]. Clearly, the proposed model (1) is more appropriate for global optimization, since it naturally overcomes such limitations based on image observations only. This is validated in the experimental section. In the following, we explain how the classical surface-based photoconsistency measure can be used to define interior/exterior photoconsistency.

Based on an initialization S_I , computed by using [10], we can define a signed distance function $\phi : V \rightarrow \mathbb{R}$ as the Euclidean distance to the closest point on S_I with negative values inside and positive values outside S_I . The distance function allows to compute the normal vector at each voxel $x \in R_{obj}^{S_I}$ as $N_x = \frac{\nabla \phi(x)}{|\nabla \phi(x)|}$. Additionally, it allows to perform efficient global visibility computations. In the following global visibility is suppressed for simplicity. At a given location $x \in V$ a matching score of cameras i and j in terms of normalized cross-correlations is defined as

$$C(x) = \frac{1}{N} \sum_i \sum_j NCC(I_i(\pi_i(x)), I_j(\pi_j(x))), \quad (2)$$

where N denotes the number of relevant camera pairs. The summation is performed only over front-facing cameras according to the estimated normal direction N_x . That is, if $V_i(x)$ denotes the normalized viewing direction of camera i , it is considered if $\arccos(V_j(x), N_x) \leq \gamma_{max}$ for some angle γ_{max} . We used $\gamma_{max} = 60^\circ$ in our experiments. The corresponding image patches are distorted using a local planar surface approximation as described in [9]. The normalized cross correlation measure was chosen due

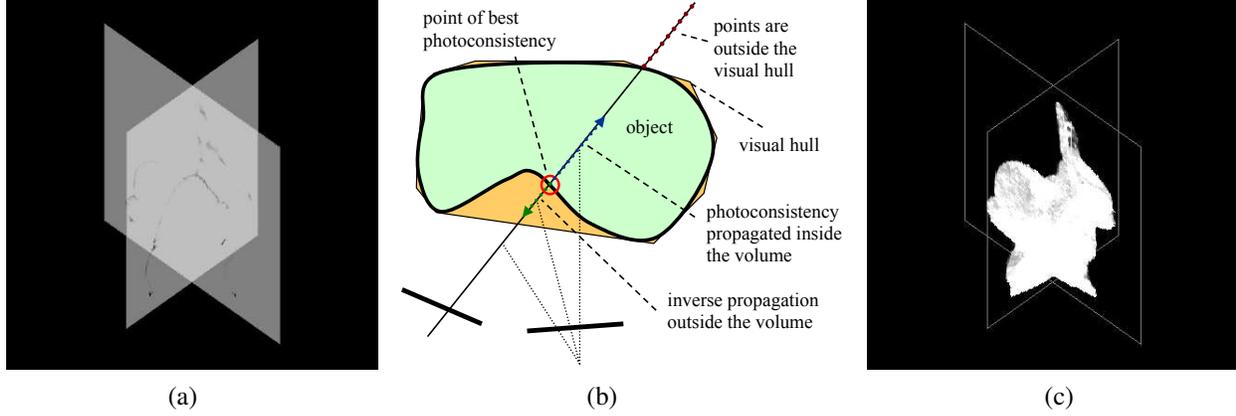


Figure 1. Volumetric propagation of photoconsistency. (a) Classical photoconsistency ρ for the "bunny" data set in Figure 4. Dark values correspond to consistent points. Photoconsistency is usually best for surface points. (b) Illustration of the proposed approach to spread photoconsistency inside a volume. (c) Resulting regional photoconsistency ρ_{bck} .

to its invariance to additive and multiplicative illumination changes. The resulting value between -1 and 1 is mapped to the interval $[0, 1]$ by applying

$$f(s) = 1 - \exp\left(-\frac{\tan\left(\frac{\pi}{4}(s-1)\right)^2}{\sigma^2}\right), \quad (3)$$

as proposed in [19]. We used $\sigma = 0.25$ in our experiments. Hence, the photoconsistency function can be written as $\rho(x) = f(C(x))$.

The main difficulty in defining spatial likelihoods is the fact that the state of each voxel in space (inside/outside the object) is affected by potentially distant points. We solve this problem by measuring photoconsistency along visual rays exploiting the following property of silhouette-consistent shapes:

Property: *Let S be an arbitrary surface, which is consistent with the silhouettes of a set of input images I_1, \dots, I_n . Then, each visual ray passing through a point x in the interior of S intersects the real observed surface \hat{S} at least once.*

If there exists a visual ray through a point x , which does not intersect the real surface \hat{S} , x does not project within the silhouette of the respective image. Hence, this point cannot lie in the interior of a silhouette-consistent shape. Note that the above property is fulfilled for the maximal consistent shape as well as for any subset of it. This leads to the following idea. We can compute photoconsistency along each visual ray and take the position, where its maximum is reached, as a potential intersection with the real surface \hat{S} ; see figure 1(b).

Based on this observation, we can convert classical photoconsistency into interior/exterior photoconsistency. More precisely, we consider all voxels x lying in the interior $R_{obj}^{S_I}$

of the surface S_I , used as initialization. Let $r_j(x, t)$ be the visual ray of camera j through voxel x , parametrized by t starting at the camera position. Let t_{cur} be the position of x along the ray. We measure photoconsistency along the ray according to another camera i :

$$C_i^j(x, t) = NCC(I_i(\pi_i(r_j(x, t))), I_j(\pi_j(r_j(x, t)))). \quad (4)$$

Here, the computation of the NCC score is based on the normal vector at x only, used to estimate image patch distortion. Note that the second term in (4) stays constant for varying t , since points on the ray r_j always project onto the same location in image I_j . This formulation can be extended to multiple cameras:

$$C^j(x, t) = \sum_{i=1}^m w_i^j(x) C_i^j(x, t). \quad (5)$$

We sum only over neighboring cameras according to the normalized viewing direction $V_j(x)$ from x to camera j . That is, camera i is excluded if $\alpha_i^j(x) := \arccos(V_i(x), V_j(x)) > \alpha_{max}$ for some bounding angle α_{max} . The weights w_i^j are computed as

$$w_i^j(x) = \frac{\alpha_{max} - \alpha_i^j(x)}{\sum_{k=1}^m \alpha_{max} - \alpha_k^j(x)} \quad (6)$$

in order to compensate for non-linear projective warping and violations of the occlusion approximation. We set $\alpha_{max} = 45^\circ$ in our experiments. As mentioned above, we determine the maximal photoconsistency along r_j together with the location, where it is reached:

$$\begin{aligned} C_{max}^j(x) &= \max_t C^j(x, t) \\ t_{max} &= \arg \max_t C^j(x, t). \end{aligned} \quad (7)$$

The sampling rate along the ray was set to the current volume resolution. Finally, we can define interior/exterior photoconsistency according to ray r_j as

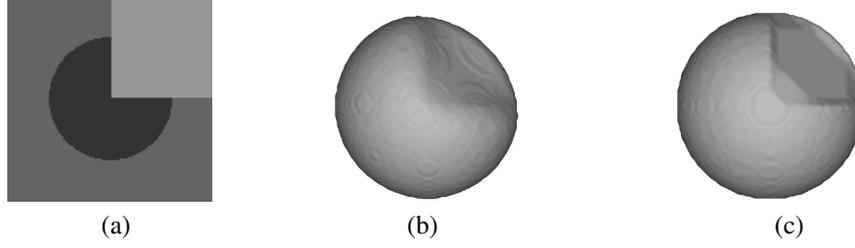


Figure 2. Continuous vs. discrete optimization. (a) A slice through the data volume. Increasing intensities denote regions with $\rho_{obj}(x) < \rho_{bck}(x)$, $\rho_{obj}(x) > \rho_{bck}(x)$ and $\rho_{obj}(x) = \rho_{bck}(x)$ respectively. Surface photoconsistency function ρ is constant throughout the volume. (b) Reconstruction obtained with the continuous optimization technique. (c) Reconstruction computed by graph cuts. In contrast to the graph cut solution, the proposed continuous optimization does not suffer from discretization artefacts.

$$\begin{aligned}\rho_{obj}^j(x) &= H(t_{max} - t_{cur}) \cdot (1 - f(C_{max}^j)) \\ &\quad + (1 - H(t_{max} - t_{cur})) \cdot f(C_{max}^j) \\ \rho_{bck}^j(x) &= H(t_{max} - t_{cur}) \cdot f(C_{max}^j) \\ &\quad + (1 - H(t_{max} - t_{cur})) \cdot (1 - f(C_{max}^j)),\end{aligned}\quad (8)$$

where H is the Heaviside function

$$H(z) = \begin{cases} 1, & \text{if } z \geq 0 \\ 0, & \text{otherwise} \end{cases}\quad (9)$$

and f is defined in (3). Note that the computed values depend on whether the maximal photoconsistency location t_{max} lies before or behind the current voxel t_{cur} . If for example $t_{max} < t_{cur}$, ρ_{obj}^j decreases and ρ_{bck}^j increases with the maximal photoconsistency C_{max}^j accounting for uncertainties because of mismatches. Averaging photoconsistency values according to single rays r_j yields

$$\begin{aligned}\rho_{obj}(x) &= \frac{1}{l} \sum_{j=1}^l \rho_{obj}^j(x) \\ \rho_{bck}(x) &= \frac{1}{l} \sum_{j=1}^l \rho_{bck}^j(x).\end{aligned}\quad (10)$$

In practice, only visual rays of front-facing cameras due to the normal N_x are considered, as previously described. Note that $\rho_{obj}(x) + \rho_{bck}(x) = 1$ for all $x \in V$. In case of photometrically difficult scenes containing highlights and shading effects a more sophisticated weighting could be used. For example, the weight of each visual ray can be chosen to be inversely proportional to the variance of the photoconsistency values.

3. Continuous Global Optimization

A discrete version of energy (1) can be globally optimized via graph cuts [11], provided the visibility of spatial points is given. Instead of using graph cuts, however, our optimization is based on a novel technique for finding the global minimum of certain *continuous* energy functionals. In [3], the authors show how certain binary non-convex minimization problems can be transformed into convex minimization problems and therefore be globally solved

by means of variational techniques. The method has been applied so far in the context of image segmentation and denoising. In this work we adopt it for the minimization problem stated in (1).

Compared to graph cuts, the proposed technique does not suffer from metrication errors while computing a globally optimal solution. This is emphasized by a synthetic experiment shown in Figure 2. Both optimization methods were applied on a sphere with a missing piece of data. At such locations the difference between both models becomes obvious. Some discretization artefacts in terms of blocky structures are available in the graph cut reconstruction even with a 26-neighborhood system. In contrast, the continuous technique yields a nice, smooth continuation of the missing part of the surface. In addition, the proposed optimization method exhibits considerable reduction in memory requirements compared to graph cuts (in our implementation about factor 20), which can be decisive when processing volumes of high resolution.

3.1. Convex Formulation

In order to cast (1) as a convex optimization problem, the surface S is represented implicitly by the characteristic function $u : V \rightarrow \{0, 1\}$ of R_{bck}^S , i. e. $u = \mathbf{1}_{R_{obj}^S}$ and $1 - u = \mathbf{1}_{R_{bck}^S}$. 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 functional equivalent to (1):

$$\begin{aligned}E(u) &= \int_V (\rho_{bck}(x) - \rho_{obj}(x))u(x)dx + \nu \int_V \rho(x)|\nabla u_\epsilon|dx, \\ \text{s. t. } u &\in \{0, 1\},\end{aligned}\quad (11)$$

where u_ϵ is a smoothed version of u .

Due to the constraint that u is a binary function, the minimization problem (11) is non-convex (because the space of binary functions is non-convex). However, when minimizing the total variation norm over all functions $u : V \rightarrow \mathbb{R}$ 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

$$E(u) = \int_V (\rho_{bck}(x) - \rho_{obj}(x)) u(x) + \nu \rho(x) |\nabla u| + \alpha \theta(u(x)) dx, \quad (12)$$

where α has to be chosen sufficiently large in order to ensure that u does not leave the interval $[0, 1]$. Since this unconstrained version of (11) is convex, it can be globally minimized via standard variational methods. This minimization already yields a binary function u almost everywhere. Moreover, as proven in [3], by applying almost any threshold $\mu \in (0, 1)$, the resulting u will be a global minimizer of the non-convex functional (11). 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 (12).

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 (12) is stated by the associated Euler-Lagrange equation

$$0 = (\rho_{bck} - \rho_{obj}) - \nu \operatorname{div} \left(\rho \frac{\nabla u}{|\nabla u|} \right) + \alpha \theta'_\epsilon(u), \quad (13)$$

where θ_ϵ is a regularized version of the derivative of θ with respect to its argument. Since (12) is convex, all minimizers are global minimizers of (12). For the application at hand, this is an important property, as it avoids suboptimal solutions of local optimization techniques, due to the spatial jaggedness of the normalized cross correlation measure.

3.2. Implementation by SOR

Discretization of the Euler-Lagrange equation leads to a sparse nonlinear system of equations. We suggest to solve this nonlinear system by a fixed point iteration scheme that transforms the nonlinear system into a sequence of linear systems. Those 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 (13) 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 , (13) yields a linear system of equations, which we solve with SOR. This means,

we iteratively compute an update of u at voxel i by

$$u_i^{l,k+1} = (1 - \omega) u_i^{l,k} + \omega \frac{\nu \sum_{j \in \mathcal{N}(i), j < i} \rho_j g_{i \sim j}^l u_j^{l,k+1} + \nu \sum_{j \in \mathcal{N}(i), j > i} \rho_j g_{i \sim j}^l u_j^{l,k} - b_i}{\nu \sum_{j \in \mathcal{N}(i)} \rho_j g_{i \sim j}^l} \quad (14)$$

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 (13) that does not depend on u , i.e. the fidelity term $b_i = \rho_{bck,i} - \rho_{obj,i}$. The over-relaxation parameter ω has to be chosen in the interval $(0, 2)$ for SOR 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 $k = 1, \dots, 10$), we 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.

4. Experiments

We applied our method on volumetric grids of about 15-20 million voxels. Computation of the photoconsistency functions and optimization of the surface were iterated for taking the more accurate visibility estimate of the most recent surface into account. In order to accelerate the approach, we used a multi-resolution scheme of 2-3 levels and performed 1-2 iterations at each level. In addition, the values of the photoconsistency functions at successive levels were updated only locally in a small vicinity around the current surface estimate. However, in contrast to some banded graph cut techniques, the surface evolution is not locally bounded, since the local band is adapted if the surface touches its boundary.

In figure 3 we show results of our approach, when applied to two publicly available data sets used for benchmarking purposes [16]. The data sets contain 48 and 47 input images, respectively, with resolutions of 640×480 . Deep indentations such as the area around the legs of dino or the back of temple, which exhibit a difficulty for some previous globally optimal approaches, are reconstructed accurately. A direct quantitative comparison in [1] reveals that the proposed method is among the most accurate approaches.

Figure 4 depicts a comparison of the proposed approach to the commonly used technique based on a ballooning term [19, 18, 14]. To this end, the first two terms in (1) were replaced by a constant expanding term preferring larger volumes. The range of expansion was restricted by the initial surface estimate based on image silhouettes. Whereas the proposed approach based on propagating photoconsistency achieves accurate reconstruction of protrusive parts

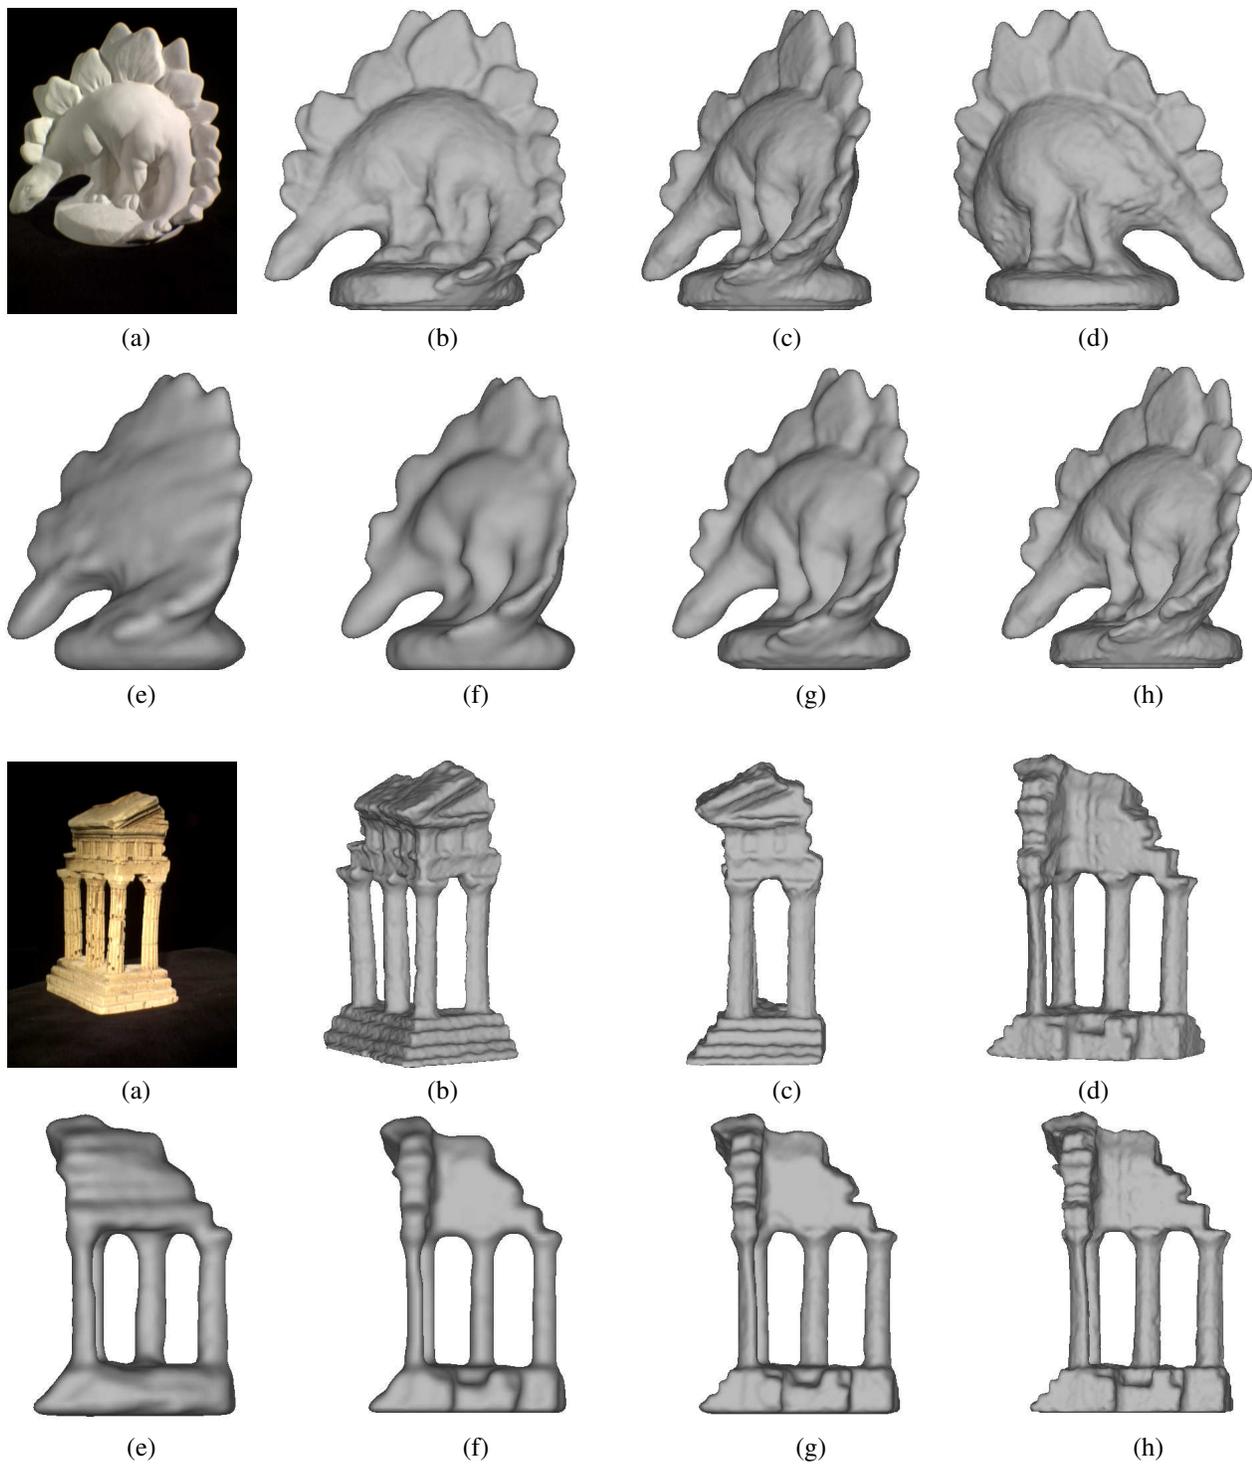


Figure 3. Dino and temple sequences. **First row, from left to right:** (a) One of the input images. (b)(c)(d) Different views of the reconstructed surface. **Second row:** (e) Visual hull reconstruction obtained by employing the technique in [10]. (f)(g)(h) Successively refined reconstructions with increasing resolutions of the volume (64^3 , 128^3 , and 256^3 for dino, $64 \times 96 \times 48$, $128 \times 192 \times 96$, and $256 \times 384 \times 192$ for temple). Note the accurate reconstruction of the concavities between the legs of dino and on the back of temple. For a qualitative and quantitative comparison with other approaches see [1].

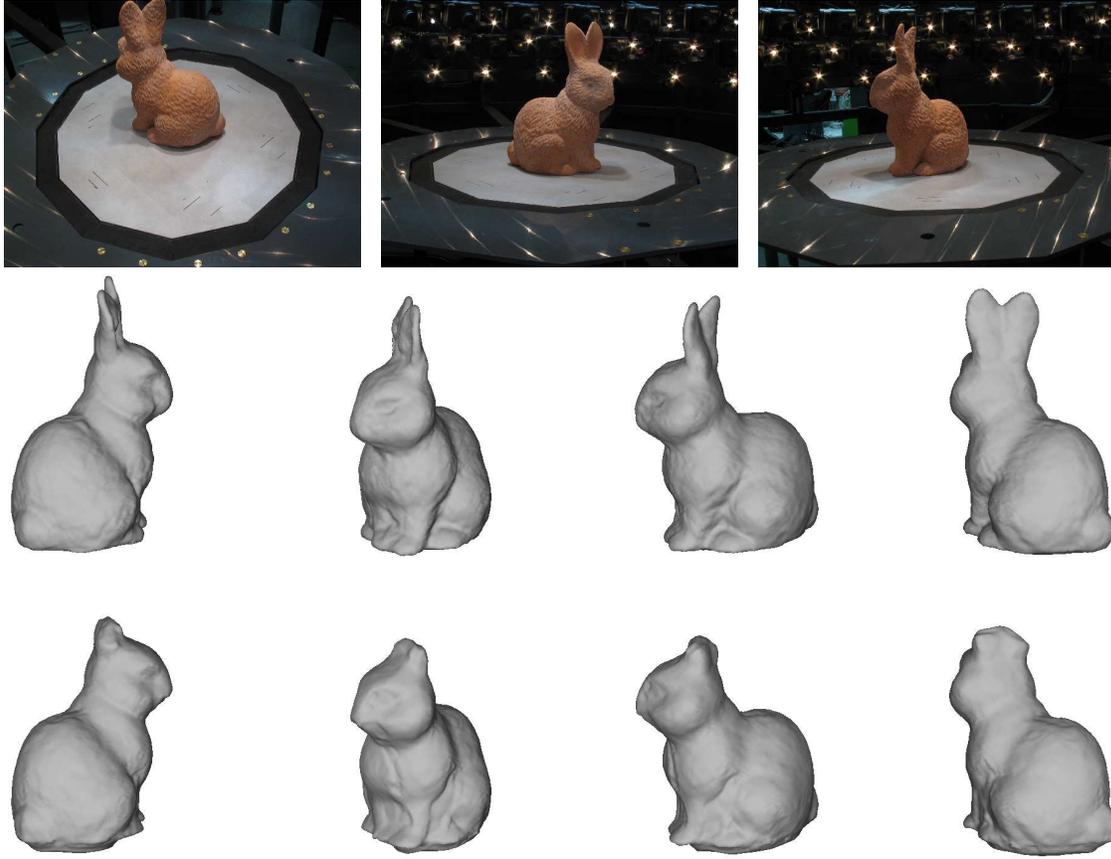


Figure 4. Bunny sequence. **First row:** Some of the input images (3 of 33) with a resolution of 640×480 . **Second row:** Reconstruction using the proposed approach. **Third row:** Reconstruction obtained when the first two terms in (1) are replaced by a ballooning force. Clearly, propagating photoconsistency produces favorable results, while the ballooning force tends to oversmooth the surface.

(the ears) as well as indentations (at the legs), the same method based on the ballooning force oversmooths both of them, although the surface photoconsistency was strengthened by a large value for σ in (3) ($\sigma = 0.5$). Increasing the weighting constant ν in (1) relative to the ballooning force leads to more accurate reconstruction of surface concavities but further cutting of protrusions. On the other hand, decreasing it preserves the protrusions but closes the concavities. Similar observations were reported in [19] and [18].

Finally, Figure 5 shows a reconstruction on an image sequence of 33 images with resolution of 640×480 . The images are challenging due to the absence of texture and the highly detailed surface. Despite these difficulties the reconstruction is accurate including deep concavities under the chin and at the ears as well as small protrusive structures like the nose.

5. Conclusion

In this paper a new variational model for 3D reconstruction from multiple views is presented. The formulation is

based on propagating surface photoconsistency in space, which allows to replace currently used artificial ballooning terms by regional data terms. This results in more accurate reconstructions, especially at surface protrusions and concavities, as demonstrated by experiments on real data sets. Moreover, up to the visibility constraint, we find a continuous, globally optimal solution according to our energy model.

References

- [1] <http://vision.middlebury.edu/mview/>. 5, 6
- [2] B. Baumgart. *Geometric modeling for computer vision*. PhD thesis, Department of Computer Science, Stanford University, USA, 1974. 1
- [3] 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. 4, 5
- [4] R. Cipolla and A. Blake. Surface shape from the deformation of apparent contours. *International Journal of Computer Vision*, 9(2):83–112, 1992. 1

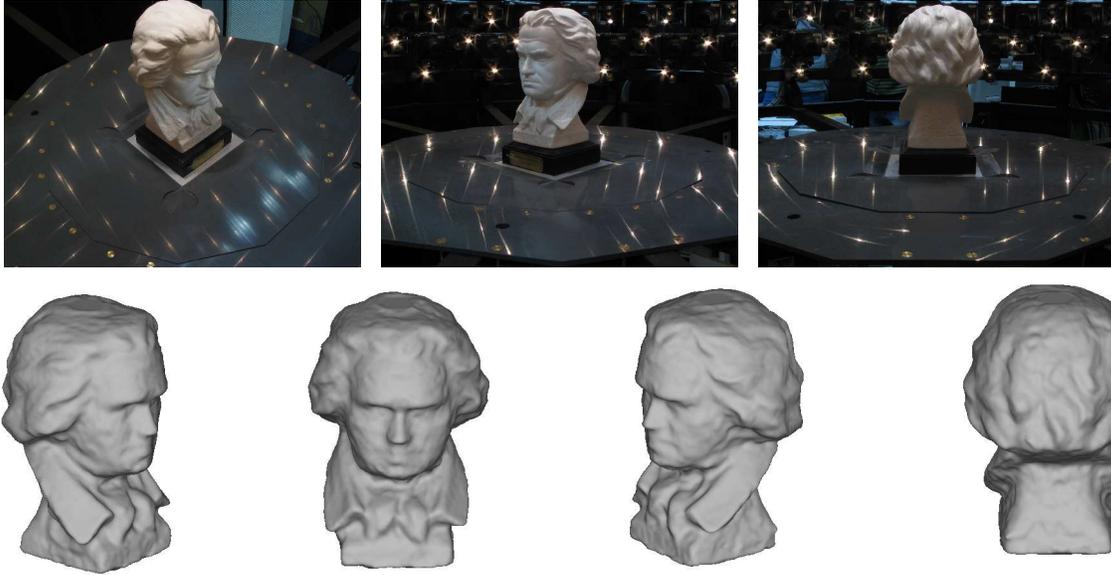


Figure 5. Beethoven sequence. **First row:** Three out of 33 input images with a resolution of 640×480 . **Second row:** Multiple views of the reconstructed surface. The highly detailed surface is reconstructed accurately despite the absence of texture.

- [5] G. Cross and A. Zisserman. Surface reconstruction from multiple views using apparent contours and surface texture. In *Confluence of Computer Vision and Computer Graphics*, pages 25–47, Norwell, MA, USA, 2000. Kluwer Academic Publishers. **1**
- [6] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, New York, NY, USA, 1996. ACM Press. **2**
- [7] 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. **1**
- [8] C. H. Esteban and F. Schmitt. Silhouette and stereo fusion for 3D object modeling. *Computer Vision and Image Understanding*, 96(3):367–392, 2004. **1**
- [9] 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. **1, 2**
- [10] 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. **1, 2, 6**
- [11] V. Kolmogorov and Y. Boykov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(9):1124–1137, 2004. **4**
- [12] K. N. Kutulakos and S. M. Seitz. A theory of shape by space carving. *International Journal of Computer Vision*, 38(3):199–218, 2000. **1**
- [13] A. Laurentini. The visual hull concept for visual-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):150–162, 1994. **1**
- [14] 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. **1, 2, 5**
- [15] 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. **1**
- [16] S. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *Proc. International Conference on Computer Vision and Pattern Recognition*, pages 519–528, Washington, DC, USA, 2006. IEEE Computer Society. **5**
- [17] 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. **1**
- [18] 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. **1, 2, 5, 7**
- [19] 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. **1, 2, 3, 5, 7**
- [20] A. Yezzi and S. Soatto. Stereoscopic segmentation. *International Journal of Computer Vision*, 53(1):31–43, 2003. **1**