

Fast and Globally Optimal Single View Reconstruction of Curved Objects

Martin R. Oswald, Eno Töppe, Daniel Cremers
TU Munich, Germany

Abstract

We propose a novel algorithmic solution for estimating a three-dimensional model of an object observed in a single image. Based on a minimal user input, the algorithm interactively determines the objects' silhouette and subsequently computes a silhouette-consistent 3D model which is precisely the globally minimal surface with user-specified volume. In contrast to a recently published approach to single view reconstruction, the proposed algorithm does not constrain the resolution in the depth-direction, it assures the global optimum and is faster by about an order of magnitude. Experiments demonstrate that plausible high-resolution 3D models can be generated in fractions of a second and compare favorably with other methods.

1. Introduction

1.1. 3D Reconstruction from a Single Image

The reconstruction of 3D shape from a single image is a very attractive challenge as it serves as a basis for generating novel views of an object observed in one image – under different viewpoint, different texture or illumination. More than most other computer vision problems single-view reconstruction is a highly ill-posed problem. As a consequence, one needs to make additional assumptions on the object's geometry (such as piecewise planarity [7],[10],[6]), its albedo (shape from texture [11]), its reflectance properties (shape from shading [8]), or the image formation process (shape from defocus [5]). Although such approaches were demonstrated to generate plausible 3D models, they all impose more or less strong limitations on the applicability to objects in real-world images. Moreover, many of these approaches give rise to hard computational challenges and are therefore unlikely to work well in an interactive real-time application.

A class of reconstruction algorithms that was shown to work well on a fairly large set of objects and images are silhouette-based reconstruction methods. Pioneering work for the reconstruction of curved surfaces was done by Terzopoulos et al. [16] although their reconstructions are restricted to tube-like shaped objects with genus 0.




Figure 1. The proposed algorithm computes optimal silhouette-consistent minimal surfaces of given volume in computation times below 1s. From left to right: input image, reconstructed geometry and textured geometry.

The work of Prasad et al. [14] made further significant advances towards the reconstruction of arbitrary curved surfaces and generalized the class of reconstructable objects to those of higher genus. In practice, however, only objects with rather simple topology can be reconstructed due to the use of a parametrized surface representation. Moreover, the required amount of user input is considerably high.

Similar to this work, Zhang et al. [21] make use of a depth map in order to calculate a smooth minimal surface which is subject to user defined constraints. However, for reasonable reconstructions the amount of necessary user input is much larger compared to the proposed method.

Somewhat related to this work is the paper of Joshi et al. [9]. There, a 2D surface generated from an input silhouette is inflated which can then be used e.g. in the generation of font effects. Inflation in their formulation is a non-intuitive local parameter of the silhouette border and differs significantly from the intuitive notion of specifying a global volume. Further, there are learning-based approaches like the one of Hassner and Basri [17] who learn depth values of image patches from a database.

1.2. The Cheeger Set Approach

Recently, Töppe et al. [18] introduced a single view reconstruction method based on Cheeger sets: It is aimed at computing silhouette-consistent minimal surfaces of a user-specified volume and provides plausible reconstructions for a large number of real-world objects. While we agree that the concept of fixed-volume minimal surfaces is an intuitive and useful paradigm for computing plausible single-view reconstructions, the formulation in [18] has a number of shortcomings:

- The volumetric representation imposes strong constraints on memory and runtime. Even with an efficient GPU-accelerated primal-dual algorithm the method requires around a second of computation time for moderate resolution reconstructions. As a consequence, higher-resolution 3D models cannot be generated at interactive speeds.
- Although the method in [18] was shown to provide exactly volume-consistent solutions, the algorithm is based on a convex relaxation and thresholding. In the absence of a threshold theorem, the method is not guaranteed to provide the globally minimal surface of specified volume. Furthermore, it is not clear whether subsequent thresholding of the relaxed solution actually leads to a spatially coherent structure (rather than a scattered set of voxels).
- Although the method is essentially computing a depth map, the authors in [18] make use of a fully volumetric representation. The required discretization of possible depth values imposes a limitation on the possible resolution in the z -direction.

1.3. Contribution of this Work

In this paper, we revisit the Cheeger set approach to single view reconstruction. In particular, we propose a novel algorithm for computing single view reconstructions which remedies the above shortcomings. More precisely:

- We propose to solve the above problem by means of a height-field representation. As a consequence, we can allow for a spatially continuous set of depth values.
- Due to the 2D representation we have substantially reduced computation time and memory requirement (quadratic rather than cubic). Experiments confirm that the proposed method allows to compute solutions about an order of magnitude faster, even for higher resolutions.
- In contrast to the algorithm in [18], the proposed method does not require convex relaxation and thresholding. As a consequence, the algorithm provably computes silhouette-consistent minimal surfaces of a specified volume.

2. Fixed Volume Minimal Surfaces: A Two-Dimensional Formulation

For the following, we will assume that the user has interactively extracted the object’s silhouette using an interactive segmentation tool like the ones in [2, 19]. Let $S \subset \mathbb{R}^2$ be the object’s silhouette, i.e. those points of the image plane corresponding to the interior of the object. In the following, we will recover a plausible estimate of the object’s




Figure 2. Area of an infinitesimal surface element dA and partial derivatives of u .

3D geometry by determining silhouette consistent surfaces of minimal area and a user-specified volume. Of course, we do not know the true depth of the object at any given point. Rather than specifying individual depth values, the user simply provides the volume which is the product of the silhouette area and the object’s average depth value. Intuitively, this approach to single-view reconstruction corresponds to a balloon being placed inside the silhouette-constrained domain and being inflated to a given volume.

The object’s surface will be represented by means of a height map

$$u : S \rightarrow \mathbb{R}, \quad S \subset \Omega \quad (1)$$

assigning a depth value $u(x, y)$ to each point $(x, y) \in S$ of the silhouette which is embedded in the image plane Ω . As shown in the schematic plot in Fig. 2, an infinitesimal surface area element dA of the surface represented by the function u is given by

$$dA = \left| \begin{pmatrix} dx \\ 0 \\ u_x dx \end{pmatrix} \times \begin{pmatrix} 0 \\ dy \\ u_y dy \end{pmatrix} \right| = \sqrt{1 + |\nabla u|^2} dx dy \quad (2)$$

The overall area of the surface denoted by u is given by

$$E(u) = \int dA = \int_S \sqrt{1 + |\nabla u|^2} dx dy \quad (3)$$

For brevity dx will denote a 2 dimensional integrand for the rest of the paper. Reconstructing a minimal surface of volume V can therefore be done by solving the minimization problem

$$\min_{u \in \mathcal{C}} E(u), \quad \text{with } \mathcal{C} = \left\{ u \mid \int_S u dx = V \right\}. \quad (4)$$

Proposition 1. *The two-dimensional fixed volume minimal surface problem defined in (4) is convex.*

Proof. The volume constraint on u is obviously a linear constraint, thus the domain of optimization is convex. Moreover, the functional E is convex because for any func-

tions u_1 and u_2 and any $\alpha \in (0, 1)$ we have:

$$\begin{aligned}
& E(\alpha u_1 + (1 - \alpha)u_2) \\
&= \int \sqrt{1 + |\nabla(\alpha u_1 + (1 - \alpha)u_2)|^2} dx \\
&= \int \sqrt{1 + |(\alpha \nabla u_1 + (1 - \alpha)\nabla u_2)|^2} dx \quad (5) \\
&\leq \int \alpha \sqrt{1 + |\nabla u_1|^2} + (1 - \alpha) \sqrt{1 + |\nabla u_2|^2} dx \\
&= \alpha E(u_1) + (1 - \alpha)E(u_2).
\end{aligned}$$

□

In contrast to the volumetric formulation proposed in [18], the two-dimensional formulation proposed here is convex. As a consequence, we do not need to revert to the generally suboptimal strategy of convex relaxation and thresholding. Instead we can directly compute globally optimal solutions by solving (4).

3. Minimization of the Proposed Energy

Minimization of the convex problem (4) can be achieved by solving the Euler-Lagrange extremality condition given by the partial differential equation

$$\frac{dE}{du} = -\operatorname{div} \left(\frac{1}{\sqrt{1 + |\nabla u|^2}} \nabla u \right) = 0. \quad (6)$$

This is a nonlinear diffusion equation which is similar to the well-known model by Perona and Malik [12] for edge-preserving image smoothing, but with a different diffusivity $g(x) = 1/\sqrt{1 + |\nabla u|^2}$ which was proposed by Charbonnier et al. [3].

Our derivation of Eq. (6) via Eq. (3) therefore provides a geometric interpretation of the Perona-Malik diffusion with the Charbonnier-diffusivity: In image diffusion the image gray values can be interpreted as a height map whose surface area is minimized as the diffusion process minimizes energy (3) (see also [15] for more details).

However, we use Eq. (6) in a completely different setting. Instead of using a data term we impose a global volume constraint and special boundary conditions which depend on the input silhouette. In the following we describe how these constraints are chosen and incorporated into the numerical optimization of Eq. (6).

3.1. Numerical Optimization

We employed three optimization schemes and compared their performance. We briefly sketch all three methods in the following. In [4] Paul Concus proposed a numerical scheme for solving the minimal surface problem (4) except that he did not consider a volume constraint.

Gradient Descent. An iterative method for solving energy minimization problems is gradient descent. In each iteration we advance in the direction of the negative gradient of the energy, finding a better solution with each step:

$$u_{t+1} = u_t - \tau \cdot \frac{dE}{du} \quad (7)$$

where τ is the step size. Since the minimization problem (4) is convex, the gradient descent method will converge to the global optimum of the energy.

FISTA. The fast iterated shrinkage and thresholding algorithm [1] can be considered as a generalized gradient descent scheme for a certain class of functions. Applied to our case it amounts to an ordinary gradient descent with adaptive over-relaxation of the current solution every few iteration steps:

$$\tau_{k+1} = \frac{1}{2}(1 + \sqrt{1 + 4\tau_k^2}) \quad (8)$$

$$u_{k+1} = u_k + \frac{\tau_k - 1}{\tau_{k+1}}(u_k - u_{k-1}) \quad (9)$$

Successive Over-Relaxation by Means of Lagged-Diffusivity. The Lagged-diffusivity approach was first employed by Vogel and Oman [20] for image deblurring and denoising. By keeping the diffusivity $g(x)$ fix over a number of iterations one can solve the resulting sparse linear equation system $\operatorname{div}(g(x)\nabla u) = 0$ with numerical solvers like Jacobi, Gauss-Seidel or Successive Over-Relaxation (SOR).

3.2. Implementation

In order to solve the optimization problem in Eq. (4) with one of the three optimization methods from above we proceed as follows: We compute one or more iterations of our optimization algorithm and then project the current solution back to the convex set \mathcal{C} of functions with a pre-described volume.

Projection Scheme. The orthogonal projection of any function u' onto \mathcal{C} can be described as the following optimization problem:

$$\min_u \frac{1}{2} \int_S \|u - u'\|^2 dx \quad \text{s.t.} \quad \int_S u dx = V. \quad (10)$$

By introducing the Lagrange multiplier $\lambda \in \mathbb{R}$ and calculating the partial derivatives of the corresponding Lagrangian function we obtain the following extremality conditions:

$$0 = u - u' + \lambda \quad \forall x \in S \quad (11)$$

$$0 = \int_S u dx - V \quad (12)$$

Inserting (11) into (12) yields

$$u = u' + \left(\frac{V - \int_S u' dx}{\int_S dx} \right) \cdot \mathbf{1}_S \quad (13)$$

as a simple update scheme for the volume projection in which $\mathbf{1}_S$ is the one-function being 1 at every point $x \in S$. Intuitively, this means that the residual volume is evenly distributed over all function values of u in S .

Boundary Conditions. In order to guarantee silhouette consistency induced by the set S , we apply Dirichlet boundary conditions at the silhouette boundary ∂S and Neumann boundary conditions if the silhouette coincides with the image boundary $\partial\Omega$:

$$u(x) = 0 \quad \forall x \in \partial S; \quad \frac{du}{dx}(x) = 0 \quad \text{if } x \in \partial\Omega. \quad (14)$$

In this way silhouette consistency is ensured and objects touching the image boundary are cut orthogonal to the image plane. Intuitively, this means that object surfaces continue uniformly at image boundaries rather than dropping to zero.

Parallelization. All minimization methods described in Sec. 3.1 have been parallelized on recent graphics hardware. This includes the projection step since it can be applied to each pixel independently once the difference between target and current volume is known. For parallelization of the SOR method a Red-Black scheme has been employed.

3.3. Weighted Minimal Surfaces

Without adding further constraints to the solution, the problem in (4) tends to be smooth by definition. In order to enable our method to reconstruct non-smooth objects, we can add local weights to the energy functional. More formally Eq. (3) extends to

$$E(u) = \int_S \rho(x) \sqrt{1 + |\nabla u|^2} dx. \quad (15)$$

Fortunately, the introduction of the weighting function $\rho : S \mapsto \mathbb{R}^+$ does not affect the convexity of the energy.

Proposition 2. *The two-dimensional fixed volume minimal surface problem defined in (4) extended with the weighting function as shown in Eq.(15) is convex.*

Proof. The proof is a straight-forward extension of the one from proposition 1. \square

Further, this extension is easily integrated into the optimization methods described above. Adding weights to

the surface considerably extends the class of possible reconstructions. Setting all weights $\rho(x) = 1$ leads to the original formulation in Eq. (3). In the implementation we use this as a default setting, however, the user can locally adapt this surface parameter.

4. Experimental Results

We tested our method on several real-world images, compared the results with three other state of the art methods and evaluated visual appearance, runtime and amount of user input.

Since one cannot obtain true depth values from a single image we do not strive for a comparison with ground truth data. We rather focus on plausibility and pleasantness of the reconstructions. Moreover, as the backsides of objects are naturally invisible in the input image, we assume the reconstructions to be symmetric. This can be achieved by simply mirroring the computed depth values along the image plane. With this heuristic, we are able to obtain closed object representations from depth maps.

4.1. Comparison to Related Methods

In Fig. 4 we visually compare our results to the ones obtained with methods by Zhang et al. [21], Prasad et al. [14] and Töppe et al. [18]. For comparison we used our own implementation of [18] and the implementation from [21]. We do not have an implementation of Prasad et al. and therefore used the results presented in [13].

The method by Zhang et al. sticks out in this comparison because it is restricted to depth map reconstruction while the other methods focus on curved objects. Except for method [18] all approaches are globally optimal and compute reconstructions at interactive frame rates. The methods mostly differ in the necessary amount of user input.

With the method of Zhang et al. the user has a variety of choices for surface manipulations such as position and normal constraints, discontinuity constraints, planar region constraints and manual mesh-subdivision. Usually many of these constraints are necessary for reasonable reconstructions leading to modeling times of several minutes to hours even for experienced users. We, as moderately experienced users, spent 20-40 minutes for each of the examples shown in Fig. 4.

Similarly, the method by Prasad et al. needs concise input and expert knowledge. The user has to assign parts of contour lines to lines in the parameter space, which becomes harder for objects of higher genus. As a result, the topology is restricted to genus two. Still, objects of higher genus exhibit over-oscillation of the surface as seen in the teapot example in Fig. 4. Moreover, for volume inflation the user needs to define a set of interpolation constraints. In subsequent steps the user may need to add further constraints for allowing surface creases. On the other hand and

in contrast to our approach, Prasad et al. can cope better with some images, in which the symmetry plane of the object is not parallel to the image plane. An example for this is represented by the donut.

Töppe et al. [18] minimize a similar energy and need the same amount of user input as our method, which is considerably less compared to the other reconstruction methods. Several examples in Figures 4 and 5 compare both approaches. Since for our method less memory and computation time is needed, it is feasible to use input images with considerably higher resolution. This results in higher detailed silhouettes and reconstructions as can be seen in the plane example in Fig. 7. Also, results of our method appear smoother as we compute continuous depth values (see e.g. the balloon). In contrast, in the approach of Töppe et al. [18] memory and runtime scale poorly with the size of the input image as they have to maintain a voxel field.

4.2. Experimental Evaluation of our Approach

Figures 5 and 4 show reconstruction results of our method for various input images. The examples represent objects of very different quality reaching from natural to man-made objects. One can see that the reconstructions appear quite plausible.

In general, since we compute a minimal surface, reconstructions will often exhibit a balloonish appearance. However, the final minimal surface strongly depends on the shape of the input silhouette. With regard to this, a strength of our approach is that volume is inflated naturally in correspondence to silhouette compactness. Examples for this favorable behavior are the bird, the stone arch and also the teapot in Fig. 4. They show that parts of the silhouette that are compact inflate more, whereas thinner structures are inflated less.

All the examples in Figures 5 and 4 come without smoothness adaption (see Sec. 3.3). In these cases, the only parameter of our approach is the volume of the reconstruction. Fig. 6 visualizes how changes of the target volume V intuitively affect the shape of the reconstruction.

In the other cases the user changed the smoothness of the surface locally. User scribbles define the locations for which the weighting factor $\rho(x)$ of Eq. (15) can be set to a user defined value. Setting $\rho(x)$ to less than 1 locally allows for sharp edges and surface extrusions like the airplane wings in Fig. 7, while values larger than 1 have the opposite effect of creating indentations. We employed this mechanism as part of an interactive feature in our single view reconstruction tool. Remember that $\rho(x) = 1$ everywhere the user did not specify weighted regions.

Fig. 7 shows some results for which the user altered the smoothness locally. One can see that non-smooth reconstructions can be achieved intuitively. Next to each reconstruction the corresponding user scribbles are shown.




Figure 3. Runtime comparison of different algorithms minimizing Eq. (4) measured on the teapot example without user-scribbles.

example		Töppe et al. [18]	our method	speedup
teapot	size	131x101x58	131x101	13.0
	time	1.82s	0.14s	
arch	size	179x137x79	179x137	6.3
	time	6.24s	0.99s	
ladybug	size	151x122x27	151x122	10.8
	time	1.62s	0.15s	
bird	size	157x244x4	157x244	10.6
	time	2.12s	0.2s	
balloon	size	82x97x44	82x97	17.7
	time	2.65	0.15s	

Table 1. Runtime comparison of the method in [18] with our method for the examples depicted in Figures 5, 6.

Runtime Comparison As described in Sec. 3, we employed a gradient descent scheme, FISTA and SOR for solving problem (4). All experiments have been done on a PC with a 2.27GHz Intel Xeon CPU, 12GB RAM equipped with a NVIDIA GeForce GTX480 graphics card running a recent Linux distribution. For comparing run-times of the respective optimization algorithms, we ran each on a reconstruction example until convergence. We then plotted for each time step t the distance $d(u_t, u^*)$ of the intermediate result u_t to the precomputed converged result u^* .

$$d(u_t, u^*) := \int_{(x,y) \in \Omega} (u_t(x,y) - u^*(x,y))^2 dx dy \quad (16)$$

The convergence criterion for all experiments has been set to

$$|(E(u_{t-1}) - E(u_t))/E(u_t)| < \theta \quad (17)$$

with $\theta = 10^{-15}$. Fig. 3 shows the results for the three optimization schemes. As can be clearly seen, the lagged diffusivity SOR approach of section 3 is the most efficient

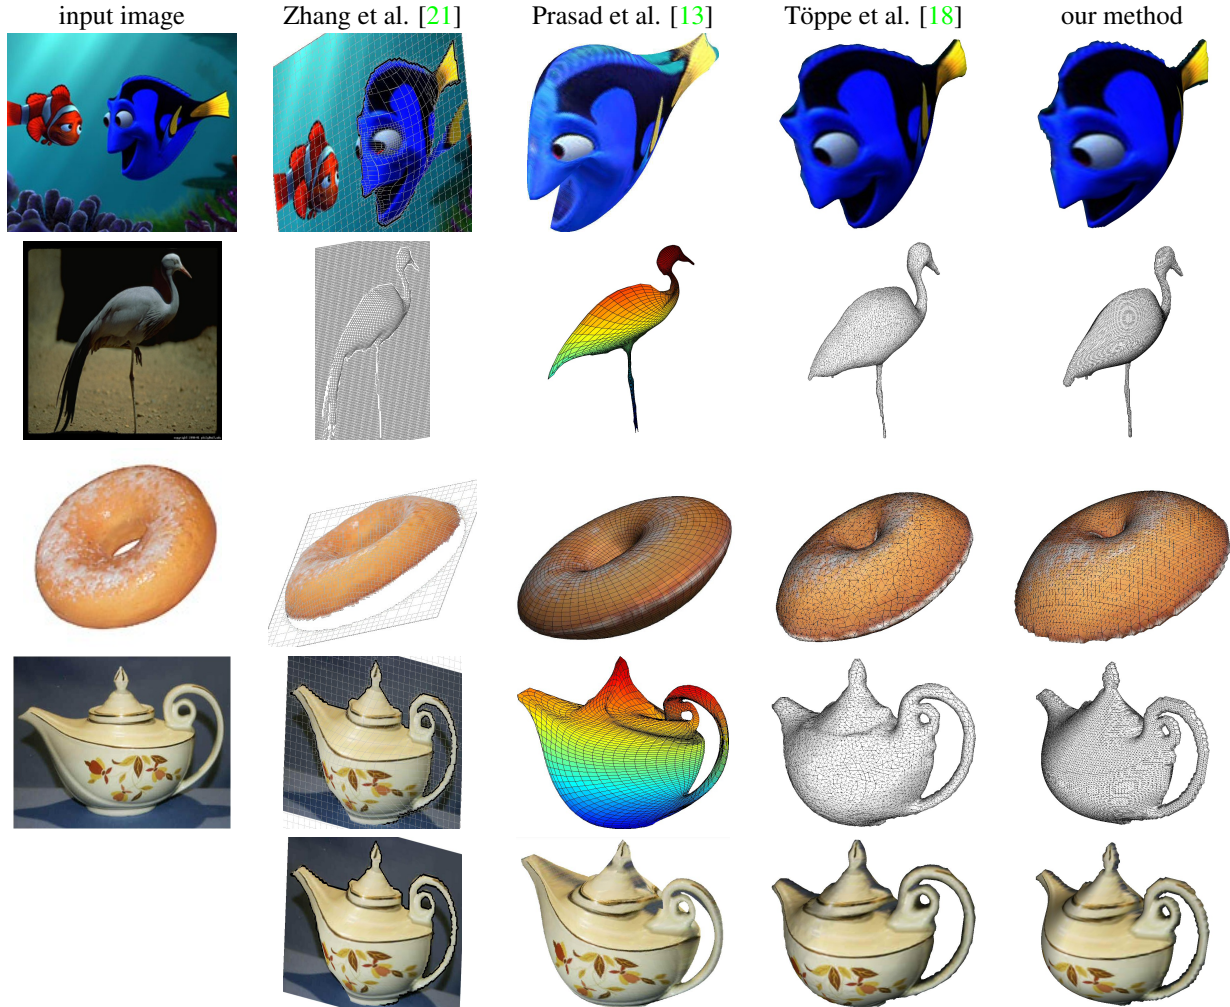


Figure 4. Comparison of reconstruction results for several single-view methods. Qualitatively our methods (*right column*) keeps up with state-of-the-art methods and sometimes even compares favorable over them.

algorithm in terms of time to convergence. The FISTA algorithm is only slightly faster than gradient descent. This is due to the fact that for differentiable functionals the algorithm degrades to a gradient descent. The performance gain stems from the adaptive over-relaxation step. Note that due to the constraints on the feasible set, we have no proof that SOR converges to the global optimum (see *Projection Scheme*). However, the results of SOR were almost equal to results from methods attaining the global optimum.

In order to evaluate the overall computational efficiency of our method we measured the computation times of the fastest optimization scheme until convergence and compared them with the method by Töppe et al. [18]. Table 1 shows detailed runtime comparisons for all experiments in Fig. 5. Since both methods optimize a convex energy, the results are independent of the initialization. The number of iterations needed until convergence, however, is not.

For all experiments the empty surface, respectively the

empty volume, has been used for initialization. When the user changes the target volume on a computed result, we can initialize the recomputation cycle with the previously computed solution. This will effectively result in a faster convergence. For input silhouettes with large areas, like the stone arch, the diffusion process has to propagate along longer distances, which leads to the relatively high runtime.

Generally, Table 1 clearly shows that our method is significantly faster than the one of Töppe et al. [18]. This difference mainly stems from the additional dimension that is used in the latter in order to discretize the depth values while our approach directly computes continuous solutions.

5. Conclusion

We have presented a novel method for single view reconstruction based on surface area minimization for pre-defined volumes. In contrast to a recently published *vol-umetric* approach, the proposed 2D solution has three ad-

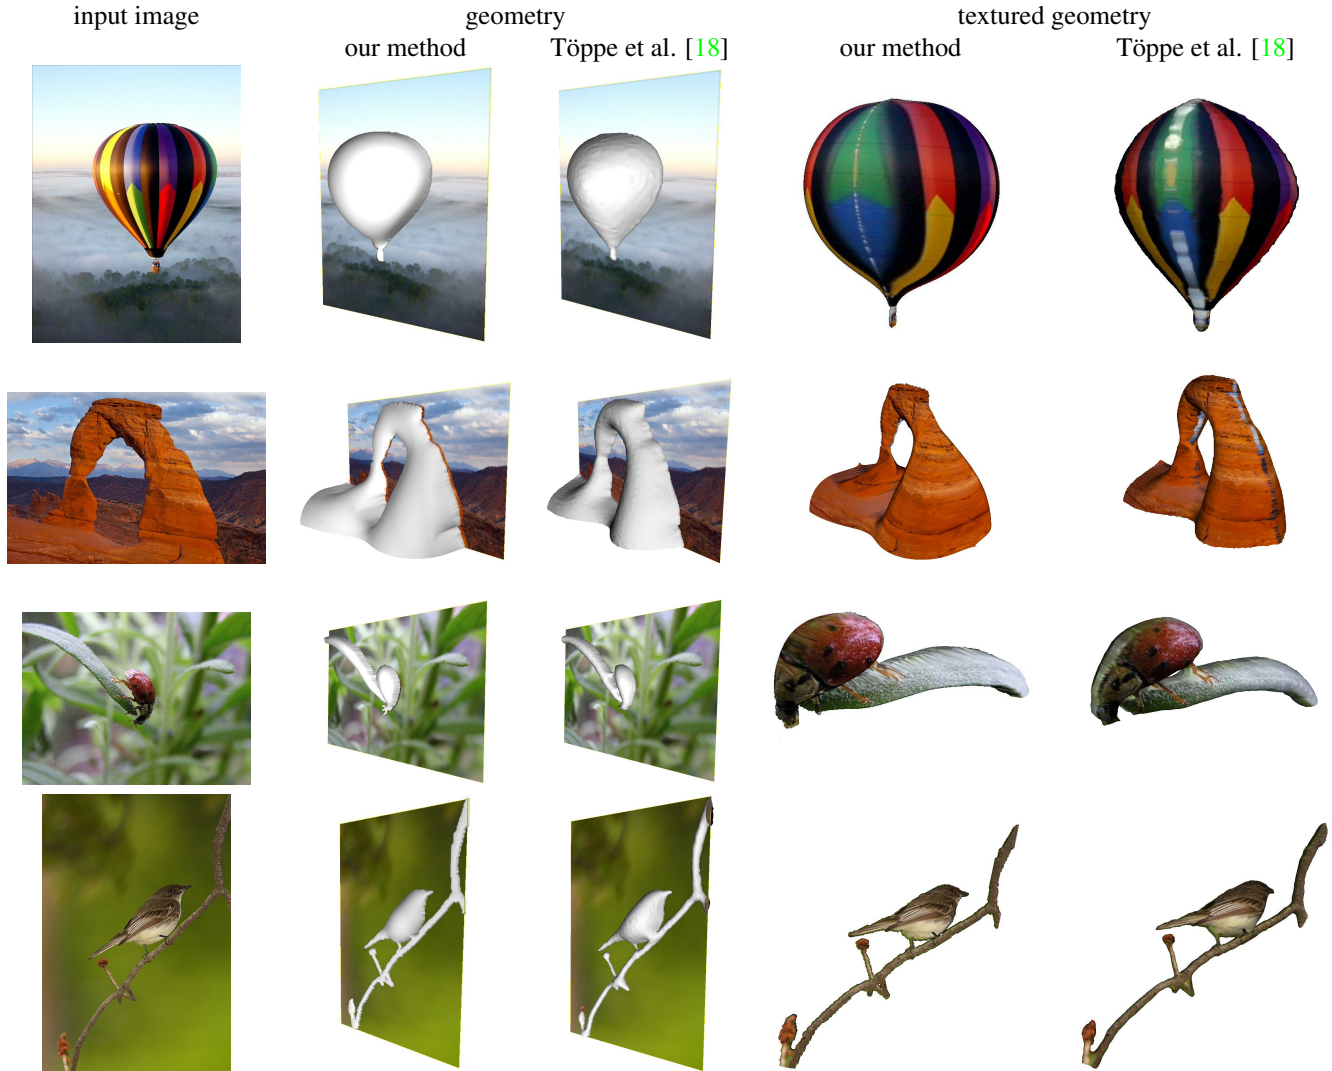


Figure 5. Our method works well on a variety of input images. Reconstruction results are similar to Töppe et al. [18] but in contrast are obtained for higher resolutions, less memory, lower computation times and higher precision.



Figure 6. Influence of the volume parameter on the reconstruction for our method. The volume distributes naturally, with more volume on compact silhouette parts and less on thin silhouette structures. The input image for this reconstruction is the arch depicted in Fig. 5

vantages: First, the resolution in the third dimensions has double-valued precision, secondly, the computed solution is provably optimal (rather than suboptimal). Thirdly, the 2D formulation drastically reduces memory and computation time (by about an order of magnitude). For a large variety of objects and good image resolutions, plausible re-

constructions are computed in fractions of a second making this method well suited for interactive 3D modeling from images. Since the target volume is the only free parameter in our model the proposed method also outperforms many previous approaches with respect to the necessary amount of user input.



Figure 7. Reconstruction results with user input altering the local smoothness of the surface. Next to the reconstructions the input images are shown with the respective user scribbles. User scribbles (yellow) decrease the surface smoothness locally.

Acknowledgments

We thank Thomas Pock and Richard Hartley for fruitful discussions on single view reconstruction. Also, we are grateful to Mukta Prasad, Li Zhang and their respective co-authors for providing their data and software.

References

- [1] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2:183–202, Mar 2009. 3
- [2] Y. Boykov and M.-P. Jolly. Interactive organ segmentation using graph cuts. In *Medical Image Computing and Computer Assisted Interventions*, volume 1935 of *LNCS*, pages 276–286. Springer, 2000. 2
- [3] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proc. International Conference on Image Processing*, pages 168–172, 1994. 3
- [4] P. Concus. Numerical solution of the minimal surface equation. *Mathematics of Computation*, 21:340–350, 1967. 3
- [5] P. Favaro and S. Soatto. A geometric approach to shape from defocus. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(3):406–417, 2005. 1
- [6] D. Hoiem, A. A. Efros, and M. Hebert. Automatic photo pop-up. *ACM Transactions on Graphics*, 24(3):577–584, 2005. 1
- [7] Y. Horry, K.-I. Anjyo, and K. Arai. Tour into the picture: using a spidery mesh interface to make animation from a single image. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, pages 225–232, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co. 1
- [8] K. Ikeuchi and B. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17:141–185, 1981. 1
- [9] P. Joshi and N. Carr. Repoussé: Automatic inflation of 2d art. In *Eurographics Workshop on Sketch-Based Modeling*, 2008. 1
- [10] D. Liebowitz, A. Criminisi, and A. Zisserman. Creating architectural models from images. In *Proc. Eurographics*, volume 18, pages 39–50, 1999. 1
- [11] J. Malik and R. Rosenholtz. Computing local surface orientation and shape from texture for curved surfaces. *International Journal of Computer Vision*, 23(2):149–168, 1997. 1
- [12] P. Perona and J. Malik. Scale-space and edge-detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990. 3
- [13] M. Prasad. *Class-based Single View Reconstruction*. PhD thesis, University of Oxford, Jul 2009. 4, 6
- [14] M. Prasad, A. Zisserman, and A. W. Fitzgibbon. Single view reconstruction of curved surfaces. In *Proc. International Conference on Computer Vision and Pattern Recognition*, pages 1345–1354, 2006. 1, 4
- [15] N. Sochen, R. Kimmel, and R. Malladi. A general framework for low level vision. *IEEE Transactions on Image Processing*, 7:310–318, 1997. 3
- [16] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry-seeking models and 3d object reconstruction. *International Journal of Computer Vision*, 1:211–221, 1987. 1
- [17] T. Hassner and R. Basri. Example based 3d reconstruction from single 2d images. In *Proc. International Conference on Computer Vision and Pattern Recognition, Beyond Patches Workshop*, page 15. IEEE Computer Society, Jun 2006. 1
- [18] E. Töppe, M. R. Oswald, D. Cremers, and C. Rother. Image-based 3d modeling via cheeger sets. In *Proc. Asian Conference on Computer Vision*, Queenstown, New Zealand, Nov 2010. 1, 2, 3, 4, 5, 6, 7
- [19] M. Unger, T. Pock, D. Cremers, and H. Bischof. Tvseg - interactive total variation based image segmentation. In *British Machine Vision Conference*, Leeds, UK, Sep 2008. 2
- [20] C. R. Vogel and M. E. Oman. Fast, robust total variation-based reconstruction of noisy, blurred images. *IEEE Transactions on Image Processing*, 7:813–824, 1998. 3
- [21] L. Zhang, G. Dugas-Phocion, J.-S. Samson, and S. M. Seitz. Single view modeling of free-form scenes. In *Proc. International Conference on Computer Vision and Pattern Recognition*, pages 990–997, 2001. 1, 4, 6