

Anisotropic Range Image Integration

Christopher Schroers¹, Henning Zimmer², Levi Valgaerts³, Andrés Bruhn⁴,
Oliver Demetz¹, and Joachim Weickert¹

¹ Mathematical Image Analysis Group, Faculty of Mathematics and Computer Science,
Campus E1.7, Saarland University, 66041 Saarbrücken, Germany
{schroers,demetz,weickert}@mia.uni-saarland.de

² Computer Graphics Laboratory, Department of Computer Science,
ETH Zurich, 8092 Zurich, Switzerland
hzimmer@inf.ethz.ch

³ Max-Planck-Institut für Informatik,
Campus E1.4, 66123 Saarbrücken, Germany
valgaerts@mpi-inf.mpg.de

⁴ Institute for Visualization and Interactive Systems,
Universitätsstraße 38, University of Stuttgart, 70569 Stuttgart, Germany
bruhn@vis.uni-stuttgart.de

Abstract. Obtaining high-quality 3D models of real world objects is an important task in computer vision. A very promising approach to achieve this is given by variational range image integration methods: They are able to deal with a substantial amount of noise and outliers, while regularising and thus creating smooth surfaces at the same time. Our paper extends the state-of-the-art approach of Zach et al. (2007) in several ways: (i) We replace the isotropic space-variant smoothing behaviour by an anisotropic (direction-dependent) one. Due to the directional adaptation, a better control of the smoothing with respect to the local structure of the signed distance field can be achieved. (ii) In order to keep data and smoothness term in balance, a normalisation factor is introduced. As a result, oversmoothing of locations that are seen seldom is prevented. This allows high quality reconstructions in uncontrolled capture setups, where the camera positions are unevenly distributed around an object. (iii) Finally, we use the more accurate closest signed distances instead of directional signed distances when converting range images into 3D signed distance fields. Experiments demonstrate that each of our three contributions leads to clearly visible improvements in the reconstruction quality.

1 Introduction

Range image integration aims at combining multiple range images, also referred to as depth maps, into a single 3D model. During the last few years, the topic of range image integration has attracted an increasing amount of attention because range images are becoming more readily available through devices such as the Kinect or time-of-flight cameras [10]. Furthermore, it is possible to compute depth maps from stereo image

pairs using an existing real time or high accuracy stereo method [12]. In this way, one can employ range image integration in a *multi-view stereo* setting [13].

Often intermediate volumetric representations are used to integrate range images because they allow handling meshes of arbitrary genus. Such a volumetric range image integration has first been used to fuse range images captured by active sensors [3, 7, 16]. In this early work, the range images are converted into 3D signed distance fields and combined into a cumulative signed distance field using an averaging scheme. The 3D model can then be obtained using an isosurface polygonisation method; see e.g. [9].

It is known that averaging without regularisation leads to inconsistent surfaces due to frequent sign changes within the cumulative signed distance field [8]. Zach et al. address this problem in [17] by computing the cumulative signed distance field as the global minimiser of a suitable energy functional that incorporates a total variation (TV) [11] smoothness term along with a robust L^1 data term.

Our Contribution. Our work extends the variational range image integration approach presented by Zach et al. [17] in three aspects. First, the isotropic (space-variant) diffusion term is replaced by an anisotropic (direction-dependent) one, which is designed to smooth along the evolving surface and evolving ridges in the cumulative signed distance field but not across. This way it is possible to obtain very smooth surfaces from noisy range images while preserving ridges and corners. Second, a normalisation is introduced in the data term to maintain balance with the smoothness term. As a result, oversmoothing of locations that are seen seldom is prevented and it is possible to obtain high quality reconstructions of objects that have been captured unevenly often from different sides. Third, compared to Zach et al., we do not use signed distances along the line of sight when converting range images into 3D distance fields. Instead, we compute the closest signed distance to the range surface. We can show for every adjustment that it leads to reconstructions of superior quality.

Organisation of the Paper. In Section 2 we describe how to obtain a more accurate signed distance field from a given range image. Section 3 explains how the signed distance fields are integrated into a globally optimal cumulative signed distance field following the idea of Zach et al. [17]. Moreover, it derives a new anisotropic smoothing behaviour as well as a meaningful normalisation factor for the data term. Subsequently, Section 4 describes implementation aspects before we display experimental results in Section 5. We conclude the paper with a summary in Section 6.

2 Signed Distance Fields

A range image maps each location of the image domain $\Omega_2 \subset \mathbb{R}^2$ to a depth value, which describes the distance from the camera centre to the surface of the scene along the corresponding optical ray. Let us assume that a range image $r : \Omega_2 \rightarrow \mathbb{R}_+$ and the corresponding camera projection $\pi : \mathbb{R}^3 \rightarrow \Omega_2$ are given.

In volumetric range image integration methods [3, 17], a range image r is converted into a 3D signed distance field f by computing the signed distance ℓ of a point




Fig. 1. Generally, the directional signed distance $\ell(\mathbf{x})$ overestimates the closest signed distance $d(\mathbf{x})$ to the range surface \mathcal{S} and is a less accurate approximation of the true distance to the object.

$\mathbf{x} \in \Omega_3 \subset \mathbb{R}^3$ along the line of sight. Computing this *directional signed distance* is computationally inexpensive because it can directly be evaluated as

$$\ell(\mathbf{x}) = r(\pi(\mathbf{x})) - |\mathbf{x} - \mathbf{c}|, \quad (1)$$

where \mathbf{c} corresponds to the camera centre, and locations in front of the surface are arbitrarily given a positive value. However, Figure 1 illustrates that the directional signed distance $\ell(\mathbf{x})$ generally overestimates the closest distance. Although directional distance and closest distance can coincide at certain locations, e.g. at the range surface itself, problems occur when averaging and regularising multiple directional signed distance values. Therefore, we propose to use a more accurate approximation of the closest distance to the object by computing the *closest signed distance* to the range surface \mathcal{S} :

$$d(\mathbf{x}) = \text{sgn}(\ell(\mathbf{x})) \cdot \inf_{\mathbf{y} \in \mathcal{S}} |\mathbf{x} - \mathbf{y}|. \quad (2)$$

Since the range image and its corresponding projection is given, the range surface \mathcal{S} can directly be evaluated. The sign of the closest distance is determined by the sign of the directional signed distance. Alternatively, the sign could also be determined using range surface normals as in [16].

We follow [17] for the remaining part of this section by scaling the signed distance values with a factor of $1/\delta$ and truncating them to the interval $[-1, 1]$:

$$f(\mathbf{x}) = \psi(d(\mathbf{x})) \quad \text{with} \quad \psi(d) = \begin{cases} \text{sgn}(d) & \text{if } |d| \geq \delta \\ d/\delta & \text{else.} \end{cases} \quad (3)$$

The parameter δ thus reflects the expected uncertainty of the depth values. We also use a binary weight $w : \Omega_3 \rightarrow \{0, 1\}$ associated with the signed distance field f in order to assign low confidence to f at locations behind the surface where $d(\mathbf{x}) < -\eta$. The parameter $\eta > 0$ thus specifies how much of the occluded region behind a surface is assumed to be solid.

3 Variational Signed Distance Field Integration

The cumulative signed distance function $u : \Omega_3 \rightarrow \mathbb{R}$ is computed as the minimiser of an energy functional of type

$$E(u) = \int_{\Omega_3} \left(D(\mathbf{f}, \mathbf{w}, u) + \alpha S(\nabla u) \right) d\mathbf{x}, \quad (4)$$

containing n signed distance fields $\mathbf{f} = (f_1, \dots, f_n)^\top$ and the associated weights $\mathbf{w} = (w_1, \dots, w_n)^\top$. The *data term* $D(\mathbf{f}, \mathbf{w}, u)$ models the assumption that u should be similar to all signed distance fields \mathbf{f} , while the *smoothness term* or *regulariser* enforces u to be smoothly varying in space by penalising large gradients of u . Its influence is steered by the smoothness weight $\alpha > 0$. The desired surface geometry is then given by the zero level set of the global minimiser u .

Zach et al. [17] employ a robust L^1 data term along with a total variation (TV) smoothness term, such that data term and smoothness term are not continuously differentiable and the resulting energy is not strictly convex. Therefore, they introduce an auxiliary variable and solve a convex approximation of this energy using a numerical scheme that combines the duality principle for the TV term with a point-wise optimisation step.

3.1 Minimisation by Gradient Descent

Alternatively, it is possible to replace the absolute value function by the continuously differentiable and strictly convex approximation $\Psi_D(s^2) = \Psi_S(s^2) = \sqrt{s^2 + \epsilon^2}$ with a small regularisation constant $\epsilon > 0$ yielding the data and the smoothness term

$$D(\mathbf{f}, \mathbf{w}, u) = \sum_{i=1}^n w_i \Psi_D((u - f_i)^2) \quad \text{and} \quad S(\nabla u) = \Psi_S(|\nabla u|^2), \quad (5)$$

respectively. The resulting energy approximates the TV- L^1 energy by Zach et al. and is strictly convex. Thus, its minimiser can be found as the steady state ($t \rightarrow \infty$) of the gradient descent equation

$$\partial_t u = \alpha \operatorname{div} (S_{\nabla u}(\nabla u)) - D_u(\mathbf{f}, \mathbf{w}, u). \quad (6)$$

When introducing the abbreviations $\Psi'_{i,D} := \Psi'_D((u - f_i)^2)$ and $\Psi'_S := \Psi'_S(|\nabla u|^2)$, data and smoothness term derivatives are given by

$$D_u(\mathbf{f}, \mathbf{w}, u) = 2 \left(u \sum_{i=1}^n w_i \Psi'_{i,D} - \sum_{i=1}^n w_i \Psi'_{i,D} f_i \right) \quad \text{and} \quad S_{\nabla u}(\nabla u) = 2 \Psi'_S \nabla u. \quad (7)$$

It is known that TV regularisation leads to minimal surfaces because it penalises the perimeter of the level sets of u [2]. Increasing the smoothness weight α thus results in reducing isolated small scale features and generating low-genus isosurfaces instead of an increased smoothing of u . However, the space-variant diffusivity Ψ'_S ignores the surface orientation. Incorporating an orientation dependent behaviour requires anisotropic smoothing [15], which is discussed next.

3.2 Anisotropic Regularisation

In order to obtain an anisotropic smoothing behaviour, we modify the diffusion term in the gradient descent equation (6) by replacing the smoothness term derivative with

$$S_{\nabla u}(\nabla u) = 2 \Psi'_S(\mathbf{J}_{\rho, \sigma}) \nabla u. \quad (8)$$




Fig. 2. Visualisation of diffusion tensors as ellipsoids for different local structures.

This essentially lifts the idea of Zimmer et al. [19], who modeled an anisotropic disparity-driven stereo vision, to three dimensions. The matrix-valued function Ψ'_S is an extension of a scalar-valued function that is applied only to the eigenvalues while leaving the eigenvectors unchanged, and

$$\mathbf{J}_{\rho,\sigma} := K_\rho * (\nabla u_\sigma \nabla u_\sigma^\top) \quad (9)$$

is the *structure tensor* [4]. Here, $u_\sigma := K_\sigma * u$, where $*$ denotes a convolution with a Gaussian K_σ of standard deviation σ .

Apparently, the structure tensor $\mathbf{J}_{\rho,\sigma}$ extends the tensor product $\nabla u \nabla u^\top$ in two aspects: first, u is regularised by a Gaussian convolution with standard deviation σ . In this context, σ can be regarded as a *noise scale* because the low-pass effect of Gaussian convolution attenuates high frequencies. Although ∇u_σ is already useful for edge detection, it is sensitive to noise if σ is chosen too small. On the other hand, cancellation effects are introduced if σ is chosen too large. This is overcome by the second aspect, which is an additional Gaussian convolution of the tensor entries with standard deviation ρ , also referred to as *integration scale*. The integration scale describes the window size over which the orientation is analysed.

Let us now discuss how the anisotropic smoothing behaviour adapts to the local structure by considering the eigenvalues of the diffusion tensor $\Psi'_S(\mathbf{J}_{\rho,\sigma})$ for the following four cases: (a) In homogeneous regions, all eigenvalues are equally large, which causes homogeneous smoothing in all three directions. (b) At smooth surfaces, one eigenvalue is close to zero, which leads to anisotropic smoothing along the surface but not across. (c) At ridges, i.e. oriented 1D structures in 3D space, only one eigenvalue is large, resulting in smoothing along the ridge. (d) At corners, all eigenvalues vanish, which prevents smoothing. Figure 2 visualises the diffusion tensors as ellipsoids, where the eigenvectors correspond to the semi-principal axes and the eigenvalues to the respective equatorial radii.

3.3 Data Term Normalisation

The data term is not in balance with the smoothness term because using more range images effectively reduces smoothing as the influence of the data term grows larger. At a first glance it seems that a remedy is given by dividing the data term by the number of cameras used. However, it often happens that different locations are visible unequally often due to the camera positioning and the surface of the object. In this case, locations

that have been seen very seldom will be excessively smoothed, whereas locations that have been seen many times only receive very little smoothing. Obviously, this cannot be overcome by adapting the data term globally. Instead, one has to consider how often each location has been observed by normalising the data term accordingly:

$$D(\mathbf{f}, \mathbf{w}, u) = \left(\sum_{i=1}^n w_i + \gamma \right)^{-1} \sum_{i=1}^n w_i \Psi_D((u - f_i)^2). \quad (10)$$

If a location has not been seen by any camera, all weights are zero and the small positive constant γ prevents division by zero. In this case, the data term evaluates to zero such that information is filled in solely based on the smoothness term.

4 Implementation

An axis aligned bounding box that contains all range surfaces is chosen as domain of integration Ω_3 . It can be discretised by choosing a number of equidistant samples $\mathbf{n} = (n_1, n_2, n_3)^\top$ in each direction, resulting in the sampling distances $\mathbf{h} = (h_1, h_2, h_3)^\top$.

Efficient Computation and Storage of Signed Distance Fields. In order to set up one of multiple signed distance fields, one has to compute the distance from a point to a triangle $n_1 \cdot n_2 \cdot n_3 \cdot m$ times when assuming that a range surface is discretised by m triangles. This complexity causes severe problems because common resolutions of 200^3 voxels and $6 \cdot 10^5$ triangles require almost $5 \cdot 10^{12}$ computations.

We use two strategies that help to reduce the computational effort when setting up the signed distance fields. First, we accelerate the computation of the closest distance for a single voxel by organising the range surface in a bounding volume hierarchy. Second, we use the directional signed distance as a heuristic for closeness to the surface and only compute the closest distance if $|\ell_i(\mathbf{x})| < c \cdot \delta$ for some $c > 1$.

Storing all signed distance fields and their associated weights directly requires a huge amount of memory. In order to reduce the memory requirement, one can either employ a voxelwise runlength encoding [17] or a coarser quantisation of the signed distances leading to a histogram based approach [18]. In some cases, adjusting the data term to enforce similarity to the pointwise weighted median can also yield acceptable results and drastically reduce memory requirement. As the signed distance fields are truncated to the interval $[-1, 1]$, it is possible to implicitly encode a weight of zero by using a value outside this interval.

Numerical Solution of the PDE. We discretise the gradient descent equation (6) on a regular grid using finite differences and solve it efficiently using the recently proposed fast explicit diffusion (FED) [6]. When ignoring the smoothness term, the minimiser u is given by the pointwise weighted median of the signed distance fields. In general, convergence can be strongly accelerated when using this as an initialisation compared to an initialisation with a constant value $z \in [-1, 1]$. If a voxel has never been seen and the bounding box is chosen rather tight, it is most probable that the voxel lies inside the




Fig. 3. From left to right: (a)-(c) A sphere was reconstructed from synthetically generated range images with $\delta = 0, |\mathbf{h}|/4, |\mathbf{h}|$ using closest signed distances. (d) Error visualisation for (c). (e) Reconstruction using directional signed distances and $\delta = |\mathbf{h}|$. (f) Error visualisation for (e).

object such that it is reasonable to initialise it with -1 . Alternatively, a coarse-to-fine approach can be employed to speed up the convergence. It depends on the input data and the choice of the regularisation parameter α , which strategy yields faster convergence.

5 Experimental Results

Guidelines for Choosing the Parameters. The parameters δ and η denote the relevant region close to the surface and the occluded region behind the surface, respectively. When choosing $\delta < |\mathbf{h}|$, subvoxel accuracy that was originally present in the range image is lost in the signed distance field due to the truncation of distances. Therefore, one can see increasing staircasing artifacts when δ goes from $|\mathbf{h}|$ towards zero (see Figure 3 (a)-(c)) and it is advisable to choose $\delta \geq |\mathbf{h}|$. Additionally, it makes sense to adapt δ in such a way that it reflects the expected measurement error in the depth maps. Choosing η involves a tradeoff: On the one hand, η should be as small as possible to avoid influencing surfaces on the other side. On the other hand, η has to be large enough to allow for sign changes in the signed distance field.

Directional Distance vs. Closest Distance. In Figure 3 (e) one can see artifacts on a sphere that was reconstructed from 48 ground truth range images using directional signed distance values. Figure 3 (f) shows a visualisation of the error values, where blue corresponds to a negative, red to a positive error value and green corresponds to an error of zero. The error for a vertex is given by its distance to the ground truth according to the error measure for accuracy used in the Middlebury Benchmark [13], and the minimum and maximum error values have been mapped to blue and red, respectively. When using the closest distance to the range surface, one can obtain an accurate reconstruction without artifacts (see Figure 3 (c),(d)).

Non-normalised Vs. Normalised Data Term. It is common to downweight the data term by the number of range images used as it is done in the experimental section of [18]. However, depending on the geometry of the object and the camera placement, different locations might not be seen equally often such that a location based normalisation as proposed by us is required. In order to demonstrate the surface evolution

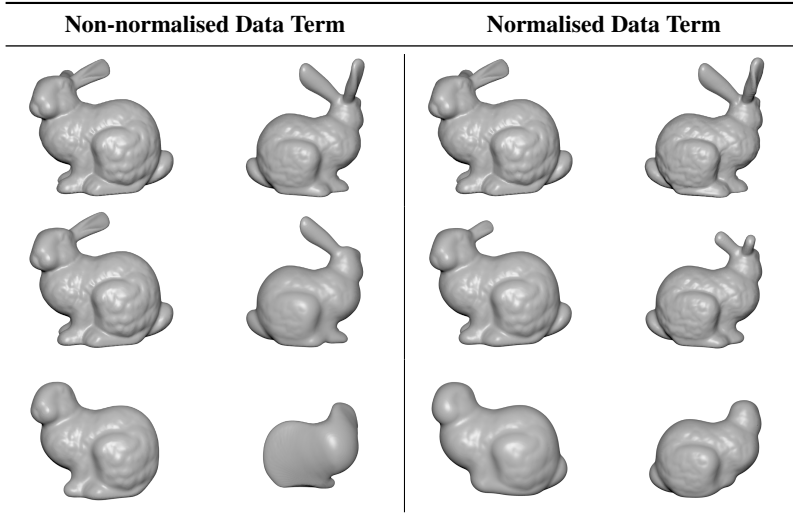


Fig. 4. With α increasing from top to bottom, one can see that the normalisation ensures an equal smoothing of front and back. The variant without normalisation excessively smoothes the back.

under increasing α , the Stanford bunny taken from the Stanford 3D scanning repository [1] is used to generate 42 range images taken from the front and only 6 from the back. Looking at Figure 4, the experiments verify that using the normalisation smoothes front and back equally as desired. On the other hand, the non-normalised variant excessively smoothes the back while performing almost no smoothing in the front.

Total Variation Vs. Anisotropic Regularisation. A stair like object has been reconstructed from 48 synthetically generated range images. In order to account for measurement errors as they often occur in modern depth cameras, a small amount of noise has been added along the line of sight. Figure 5 shows that the anisotropic diffusion term achieves a superior result when compared to the isotropic one. In the magnification it is clearly visible that the flat parts of the stairs are extremely smooth while the ridges are preserved. The TV smoothness term is also able to preserve the ridges but cannot achieve a comparable smoothness in the flat regions. This also holds for the reconstructions of the full dino dataset from the Middlebury benchmark depicted in Figure 6 (b) and (c). We have used closest distances when converting the range images into signed distance fields, and we have computed depth maps according to the method of Valgaerts et al. [14]. Compared to the method of Zach [18], we are able to significantly improve the accuracy from 0.55mm to 0.33mm. In fact, only Furukawa and Ponce [5] are able to obtain a higher accuracy of 0.32mm at this time. However, the reconstruction of Furukawa is not able to achieve a similar smoothness, such that our reconstruction is visually closer to the ground truth (see Figure 6 (a)). With 4.5 hours, the runtime of our CPU-based implementation lies between those of Furukawa (5.75 h) and Zach (3 min) (see also <http://vision.middlebury.edu/mview/>). Furthermore, our FED scheme is well suited for parallelisation on GPUs.




Fig. 5. Stairs were reconstructed from 48 synthetically generated range images with a small amount of noise added along the line of sight. *From left to right: (a) One noisy range surface. (b) Reconstruction with TV smoothness term. (c) Reconstruction with anisotropic diffusion term.*




Fig. 6. *Top row:* Reconstructions for the dino full dataset from the Middlebury benchmark with the errors 0.32mm, 0.55mm and 0.33mm, respectively. *Bottom row:* Magnification.

6 Conclusions and Future Work

We have extended the variational range image integration method of Zach et al. [17] w.r.t. three aspects. First, an anisotropic smoothing behaviour that outperforms the existing isotropic one has been proposed. It can produce smoother surfaces while preserving ridges and corners. Second, a normalisation of the data term ensures that portions of an object which have only been captured seldom are not overly smoothed. It allows for high quality reconstructions in setups where the camera positions are unevenly distributed. Third, the signed distance fields were generated from range images by computing the closest signed distance to the range surface instead of evaluating the directional signed distance along the line of sight. In the experimental section, we could show that all three modifications were able to improve the reconstruction quality.

Future work could aim at increasing the quality of reconstructions even further by recognising depth discontinuities and treating them in a more sophisticated way.

Acknowledgements. Funding by the Cluster of Excellence *Multimodal Computing and Interaction* is gratefully acknowledged. This work has been performed while all authors were affiliated members of the Mathematical Image Analysis Group.

References

1. The Stanford 3D Scanning Repository, <http://graphics.stanford.edu/data/3Dscanrep/>
2. Chan, T.F., Esedoglu, S.: Aspects of total variation regularized L^1 function approximation. *SIAM Journal on Applied Mathematics* 65(5), 1817–1837 (2004)
3. Curless, B., Levoy, M.: A volumetric method for building complex models from range images. In: *Proceedings of SIGGRAPH 96*. vol. 3, pp. 303–312 (1996)
4. Förstner, W., Gülch, E.: A fast operator for detection and precise location of distinct points, corners and centres of circular features. In: *Proc. ISPRS Intercommission Conference on Fast Processing of Photogrammetric Data*. pp. 281–305. Interlaken, Switzerland (Jun 1987)
5. Furukawa, Y., Ponce, J.: Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32(8), 1362–1376 (2010)
6. Grewenig, S., Weickert, J., Bruhn, A.: From box filtering to fast explicit diffusion. In: Gesele, M., Roth, S., Kuijper, A., Schiele, B., Schindler, K. (eds.) *Pattern Recognition, Lecture Notes in Computer Science*, vol. 6376, pp. 533–542. Springer, Berlin (2010)
7. Hilton, A., Stoddart, A., Illingworth, J., Windeatt, T.: Reliable surface reconstruction from multiple range images. In: Buxton, B., Cipolla, R. (eds.) *Computer Vision ECCV '96, Lecture Notes in Computer Science*, vol. 1064, pp. 117–126. Springer Berlin (1996)
8. Hornung, A., Kobbelt, L.: Robust reconstruction of watertight 3D models from non-uniformly sampled point clouds without normal information. In: *Proc. Fourth Eurographics Symposium on Geometry Processing*. pp. 41–50. Eurographics Association (2006)
9. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. In: *Proceedings of SIGGRAPH 87*. vol. 21, pp. 163–169 (Jul 1987)
10. Newcombe, R.A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A.J., Kohli, P., Shotton, J., Hodges, S., Fitzgibbon, A.: Kinectfusion: Real-time dense surface mapping and tracking. *IEEE International Symposium on Mixed and Augmented Reality* 7(10), 127–136 (2011)
11. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D* 60, 259–268 (1992)
12. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision* 47(1-3), 7–42 (2002)
13. Seitz, S., Curless, B., Diebel, J., Scharstein, D., Szeliski, R.: A comparison and evaluation of multi-view stereo reconstruction algorithms. In: *Proc. 2006 IEEE Conference on Computer Vision and Pattern Recognition*. pp. I: 519–528. IEEE Computer Society Press, New York, NY (Jun 2006)
14. Valgaerts, L., Bruhn, A., Weickert, J.: A variational model for the joint recovery of the fundamental matrix and the optical flow. In: Rigoll, G. (ed.) *Pattern Recognition. Lecture Notes in Computer Science*, vol. 5096, pp. 314–324. Springer, Berlin (2008)
15. Weickert, J.: *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart (1998)
16. Wheeler, M.D., Sato, Y., Ikeuchi, K.: Consensus surfaces for modeling 3D objects from multiple range images. In: *Proc. Sixth International Conference on Computer Vision*. pp. 917–924. IEEE Computer Society (1998)
17. Zach, C., Pock, T., Bischof, H.: A globally optimal algorithm for robust TV- L^1 range image integration. In: *Proc. Ninth International Conference on Computer Vision*. pp. 1–8. IEEE Computer Society Press, Rio de Janeiro, Brazil (Oct 2007)
18. Zach, C.: Fast and high quality fusion of depth maps. In: *Proc. Fourth International Symposium on 3D Data Processing, Visualization and Transmission*. pp. 1–8 (2008)
19. Zimmer, H., Bruhn, A., Valgaerts, L., Breuß, M., Weickert, J., Rosenhahn, B., Seidel, H.P.: PDE-based anisotropic disparity-driven stereo vision. In: Deussen, O., Keim, D., Saupe, D. (eds.) *Proceedings of Vision, Modeling, and Visualization 2008*. pp. 263–272. Akademische Verlagsgesellschaft Aka, Konstanz, Germany (Oct 2008)