A continuous surface reconstruction method on point cloud captured from a 3D surface photogrammetry system

Wenyang Liu

Department of Bioengineering, University of California, Los Angeles, California 90095

Yam Cheung, Pouya Sabouri, Tatsuya J. Arai, and Amit Sawant

Department of Radiation Oncology, University of Texas Southwestern, Dallas, Texas 75390

Dan Ruan^{a)}

Department of Bioengineering, University of California, Los Angeles, California 90095 and Department of Radiation Oncology, University of California, Los Angeles, California 90095

(Received 25 March 2015; revised 16 August 2015; accepted for publication 28 September 2015; published 20 October 2015)

Purpose: To accurately and efficiently reconstruct a continuous surface from noisy point clouds captured by a surface photogrammetry system (VisionRT).

Methods: The authors have developed a level-set based surface reconstruction method on point clouds captured by a surface photogrammetry system (VisionRT). The proposed method reconstructs an implicit and continuous representation of the underlying patient surface by optimizing a regularized fitting energy, offering extra robustness to noise and missing measurements. By contrast to explicit/discrete meshing-type schemes, their continuous representation is particularly advantageous for subsequent surface registration and motion tracking by eliminating the need for maintaining explicit point correspondences as in discrete models. The authors solve the proposed method with an efficient narrowband evolving scheme. The authors evaluated the proposed method on both phantom and human subject data with two sets of complementary experiments. In the first set of experiment, the authors generated a series of surfaces each with different black patches placed on one chest phantom. The resulting VisionRT measurements from the patched area had different degree of noise and missing levels, since VisionRT has difficulties in detecting dark surfaces. The authors applied the proposed method to point clouds acquired under these different configurations, and quantitatively evaluated reconstructed surfaces by comparing against a high-quality reference surface with respect to root mean squared error (RMSE). In the second set of experiment, the authors applied their method to 100 clinical point clouds acquired from one human subject. In the absence of ground-truth, the authors qualitatively validated reconstructed surfaces by comparing the local geometry, specifically mean curvature distributions, against that of the surface extracted from a high-quality CT obtained from the same patient.

Results: On phantom point clouds, their method achieved submillimeter reconstruction RMSE under different configurations, demonstrating quantitatively the faith of the proposed method in preserving local structural properties of the underlying surface in the presence of noise and missing measurements, and its robustness toward variations of such characteristics. On point clouds from the human subject, the proposed method successfully reconstructed all patient surfaces, filling regions where raw point coordinate readings were missing. Within two comparable regions of interest in the chest area, similar mean curvature distributions were acquired from both their reconstructed surface and CT surface, with mean and standard deviation of ($\mu_{\text{recon}} = -2.7 \times 10^{-3} \text{ mm}^{-1}$, $\sigma_{\text{recon}} = 7.0 \times 10^{-3} \text{ mm}^{-1}$) and ($\mu_{\text{CT}} = -2.5 \times 10^{-3} \text{ mm}^{-1}$, $\sigma_{\text{CT}} = 5.3 \times 10^{-3} \text{ mm}^{-1}$), respectively. The agreement of local geometry properties between the reconstructed surfaces and the CT surface demonstrated the ability of the proposed method in faithfully representing the underlying patient surface.

Conclusions: The authors have integrated and developed an accurate level-set based continuous surface reconstruction method on point clouds acquired by a 3D surface photogrammetry system. The proposed method has generated a continuous representation of the underlying phantom and patient surfaces with good robustness against noise and missing measurements. It serves as an important first step for further development of motion tracking methods during radiotherapy. © 2015 American Association of Physicists in Medicine. [<http://dx.doi.org/10.1118/1.4933196>]

Key words: surface reconstruction, implicit representation, photogrammetry, variational formulation

1. INTRODUCTION

Understanding motion and deformation of the internal anatomy is of great importance in radiotherapy. It helps to accurately reconstruct delivered dose distribution and to generate more accurate image basis for adaptive planning and delivery. However, direct monitoring of internal motion remains a very difficult task, due to limitations in spatiotemporal resolutions in imaging development (e.g., 4D MRI and 4DCT) and other factors such as concomitant dose considerations. Very often, an indirect approach is taken where a most accessible surrogate (or partial) measurement is used in combination with a model that relates the internal anatomy to the surrogate. Examples of such indirect monitoring include electronic portal imaging (EPI),¹ optical marker tracking,² and fluoroscopic/cone-beam computed tomography (CBCT)-based monitoring.³ Optical monitoring has the benefit of radiation-free and enjoys high frame rates. As an advancement from optical marker tracking, the development of photogrammetry techniques provides a data-intensive yet noninvasive way for monitoring patient external surface motion in real time. Such systems, exemplified by VisionRT, determine 3D coordinate of points on the patient surface based on measurements made from two or more photographic images taken from different positions/angles. To monitor motion and thoracic setup, chest skin surface is monitored at 10 Hz and reconstructed. Potentially, such surface data can be aligned to a reference data volume or longitudinally to the measurement acquired at a preceding time to facilitate setup and real-time monitoring.^{4,5}

However, sensor properties, lighting conditions, and view occlusions result in noise and patches of absent measurement during photogrammetry acquisition and challenge the accurate monitoring and interpretation of surface motion. Previous validation studies have revealed large registration errors in the order of centimeters due to the poor quality of photogrammetry-reconstructed surface, based on direct point cloud representation followed by the iterative closest point (ICP) matching procedure⁶—an approach currently employed by the commercial VisionRT system.

From a physics perspective, ICP registration between point clouds implicitly assumes the existence of one-to-one correspondence between (a majority of) members of the two point clouds, which does not hold with randomly projected patterns in VisionRT over time. Such model mismatch may be tolerated when only very simple 1D respiratory surrogate is required, because such random errors tend to “average out”—which is the scenario of current clinical practice. As richer information being extracted, the severity of this issue becomes more pronounced: it is the major culprit for the large errors for rigid alignment (five degrees of freedom).⁶ In the active pursuit of even higher dimensional models of volumetric description of internal anatomical geometry from surface information, registering inconsistent point clouds risks significantly misleading results. To this end, an alternative and more robust surface reconstruction method is needed.

Surfaces can be represented either explicitly or implicitly. The explicit approaches, such as parametrical surfaces⁷ and triangulated surfaces,^{8,9} often have requirements on data

dimensionality or regularity on the distribution of control points, making their adoption challenging in handling noise and tracking large deformations.¹⁰ The implicit surface representations, on the other hand, usually embeds a surface into a smooth scalar function and enjoys computational convenience and topological flexibility,^{11,12} with a typical example being the level set approach.¹³

We adopt the level set representation in this study and propose to reconstruct *continuous* instantaneous surfaces based on the acquired point clouds from VisionRT by optimizing a regularized fitting energy. The proposed method avoids building/maintaining explicit point correspondence and offers extra robustness to noise and missing measurements. A narrowband solving scheme is implemented to further improve the efficiency. To the best of our knowledge, this is the first time a surface reconstruction method has been applied toward point clouds from VisionRT system, and it serves as an important first step to facilitate the development of sophisticated high dimensional models for motion tracking.

2. METHOD

2.A. Variational formulation of the weighted minimal surface model

The level set method has been widely applied to surface representation due to its flexibility and computational convenience.¹⁴ As an implicit scheme, it offers a continuous representation of underlying curves/surfaces and avoids constructing explicit point correspondence for capturing and tracking deformations. Specifically, a surface C is implicitly represented as the zero level set of its corresponding level set function ϕ in a space that is one dimension higher. In this study, given a point cloud P , we aim to construct a continuous surface by optimizing the following regularized energy:^{10,15}

$$E(\phi(\mathbf{x})) = \int_{\Omega} d(\mathbf{x}) \delta(\phi(\mathbf{x})) |\nabla \phi(\mathbf{x})| d\mathbf{x}, \quad (1)$$

where $\phi(\mathbf{x})$ represents the level set function, δ is the Dirac delta function, and $d(\mathbf{x})$ is defined as a distance function to the point cloud $d(\mathbf{x}) = \text{dist}(\mathbf{x}, P)$. Note that $\{\mathbf{x}: \phi(\mathbf{x}) = 0\} = \emptyset$ provides a collection of trivial global minimum for the above objective. In this study, we are only interested in finding local minimum solutions that correspond to the minimal surface attached to the point cloud and minimizing the energy in Eq. (1) with the gradient descent method following the variational flow,

$$\frac{\partial \phi}{\partial t} = |\nabla \phi| \left(\nabla d \cdot \frac{\nabla \phi}{|\nabla \phi|} + d \nabla \cdot \frac{\nabla \phi}{|\nabla \phi|} \right). \quad (2)$$

The evolution of the above PDE balances between the data fidelity attraction $\nabla d \cdot (\nabla \phi / |\nabla \phi|)$ and the surface tension regularization $d \nabla \cdot (\nabla \phi / |\nabla \phi|)$. The natural scaling of d will adaptively increase the relative regularization weight on the low sampling or missing region from the point cloud,¹⁵ which is desirable for our application.

2.B. Narrowband evolving scheme

The level set evolution is computationally expensive, with evolution time step limited by the Courant–Friedrichs–Lewy (CFL) condition. Given a d dimensional reconstruction task, the complexity is $O(n^d)$ per time step, where n is the total number of points on a regular spatial computation grid. Notice that, however, the interface of interest is represented as the zero level of the level set function, updating points within a small neighborhood of the evolving interface would suffice. In this study, we adopt a narrowband evolving scheme¹⁶ to improve the efficiency. Specifically, we restrict computations to be within a narrowband Ω_w about the current zero level set with width w , and we update the narrowband region at each time step t as $\Omega_w = \{\mathbf{x} : |\phi(\mathbf{x}, t)| \leq w/2\}$. This narrowband scheme reduces the computational complexity from $O(n^d)$ to $O(w n^{d-1})$, since at each time step, we only need to calculate the derivatives and update the level set function at those $w n^{d-1}$ grid points within the narrowband region instead of the full domain.

2.C. Implementation details

2.C.1. Point cloud padding

One typical point mesh captured by VisionRT system is shown in Fig. 1(a), where an open upper surface of the patient torso is captured. This open surface presents a challenge to the level set optimization in Eq. (1), easily leading to annihilation of the evolving surfaces (the trivial global minimum). As a practical remedy, we artificially attach a flat point slab to define the back boundary, normally at the couch level, as shown in Fig. 1(b). This attachment has no impact on the front surface reconstruction but provides the desired stability and robustness to the level set evolution.

2.C.2. Distance function to the point cloud

The distance function $d(\mathbf{x})$ to the point cloud P is obtained by numerically solving the following Eikonal equation¹⁰ with fast marching method:^{13,17}

$$|\nabla d(\mathbf{x})| = 1, \quad d(\mathbf{x}) = 0, \quad \mathbf{x} \in P. \quad (3)$$

FIG. 1. Point cloud padding scheme: (a) the original point mesh acquired by VisionRT; (b) a flat point slab is padded to define the back boundary.

ALGORITHM I. Surface reconstruction.

Input: padded point cloud P , multilevel number S , narrowband width w , final grid r
Initialize ϕ_S^0 as a rectangular bounding box enclosing point cloud at the coarsest level
for $s = S$ to 1 **do**
 $P_s \leftarrow$ Evaluating P on grid r downsampled by 2^{s-1}
 Construct distance function on P_s by solving the Eikonal equation:
 $|\nabla d_s(\mathbf{x})| = 1, \quad d_s(\mathbf{x}) = 0, \quad \mathbf{x} \in P_s.$
 $k \leftarrow 0$
 while $|\phi_s^{k+1} - \phi_s^k| < \text{tol}$ **do**
 Update narrowband: $\Omega_w = \{\mathbf{x} : |\phi_s^k(\mathbf{x})| \leq \frac{w}{2}\}$
 Evolution within the updated narrowband Ω_w :
 $\phi_s^{k+1} = \phi_s^k + \Delta t \{ |\nabla \phi_s^k| (\nabla d_s \cdot \frac{\nabla \phi_s^k}{|\nabla \phi_s^k|} + d_s \nabla \cdot \frac{\nabla \phi_s^k}{|\nabla \phi_s^k|}) \}$
 end while
 if $s > 1$ **then** $\phi_{s-1}^0 \leftarrow \text{Interpolate}(\phi_s^{k+1})$
 end if
end for
Output: final reconstructed surface $\phi \leftarrow \phi_s^{k+1}$

2.C.3. Initialization and multiresolution

Initialization is important in level-set based methods. A good initialization can significantly speed up the convergence and prevent undesirable local/global minima. In our surface reconstruction task, an ideal initialization would be an outer surface that tightly bounds the exterior of the point cloud. Surface $\mathcal{T} = \{\mathbf{x} : d(\mathbf{x}, P) = \epsilon\}$, where ϵ is a small positive number, serves as a good candidate. However, the existence of multiple isolated regions (islands) due to occlusion in the point cloud could render such initialization misleading. Detecting and tagging the region of the surface defined¹⁵ are potentially useful, but such method is empirically sensitive to the choice of relaxation parameters and its robustness depends heavily on the sampling density of the point cloud. In this study, we directly tackle this problem with a multiresolution framework.¹⁰ Specifically, at the coarsest grid resolution, a rectangular bounding box is placed to enclose the point cloud. The reconstruction result of the coarser level is propagated as the initialization for the finer-grid-level evolution. Algorithm I presents implementation details of our proposed method.

2.C.4. Parameter settings and computation complexity

Our surface reconstruction was implemented on a rectangular grid size of $200 \times 150 \times 75$. A two-level grid resolution ($S = 2$) was used for the evolution. The coarse level was implemented on a grid size of $40 \times 30 \times 15$. Our narrowband evolving scheme was implemented based on LSMLIB (<http://ktchu.serendipityresearch.org/software/lsmlib/>), an open source level set library. The narrowband width was set to $w = 4$ for all experiments, with time step $\Delta t = 0.01$. Our method was tested with MATLAB 2013b on Mac OS 10.9.4, running on a macbook laptop with quad-core I7 2.3 GHz, 8GB RAM. The total running time for reconstructing one surface was about 30 s, without using any specific code optimization for this particular pilot study on methodology. The CT images used in this study were acquired on the abdomen region with a Philips Brilliance Big Bore 16-slice CT simulator. Imaging parameters were $\text{FOV} = 600 \times 600 \text{ mm}^2$, matrix size = 512×512 , and resolution = $1.17 \times 1.17 \times 2 \text{ mm}^3$. Imaging protocols in this study were approved by local Institutional Review Board. The enrolled subject was provided informed consent in accordance with institutional policy.

3. EXPERIMENTS AND RESULTS

To quantitatively evaluate the performance of the proposed method especially on point clouds with noise and missing measurements, we first made a series of static acquisitions on one chest phantom under various configurations with different black patches placed on the abdomen area. With the same acquisition/lighting condition and time scale, we compared the reconstructed surfaces under different testing configura-

tions against the reference surface reconstructed from phantom without any patches, and quantitatively evaluated the reconstruction accuracy with respect to root mean squared error (RMSE).

We then tested the proposed method on clinical point clouds captured from one patient during radiotherapy. In the absence of ground-truth instantaneous surface, we assume that surface extracted from high-quality CT provides a faithful representation of the geometric characteristics of the patient surface. Note that CT acquisition is at a different time scale than the 3D optical measurement, if not at a completely different time instance, so the two surfaces should not be compared on a pointwise fashion. We assessed the geometric characteristics by evaluating the local curvature behaviors on comparable regions of interest (ROIs). In particular, we selected two comparable ROIs around the chest area from our reconstructed surfaces and the CT surface, respectively, and compared the mean curvature distributions within those two ROIs.

3.A. Surface reconstruction on phantom

The phantom configurations of the reference measurement and testing measurements with different black patches are shown in Fig. 2. Specifically, we used a chest phantom without any artificial black patches and acquired a high-quality scan as the reference surface. For different configurations, we placed different black patches on the abdomen to simulate different acquisition conditions, since VisionRT system has difficulties in detecting those dark areas due to reflective conditions. Figure 3 shows example point clouds captured under these different configurations, where missing measurements usually occurred at positions with black patches. We reconstructed 50

FIG. 2. Different measurement configurations of the chest phantom scans: (a) phantom configuration for the reference measurement, (b) phantom configuration with black patch on the lower left side of the abdomen, (c) phantom configuration with two adjacent patches on the lower abdomen, and (d) phantom configuration with two disjoint patches on the lower abdomen.

FIG. 3. Example point clouds acquired under different phantom configurations. (a) Reference configuration, (b) lower-left-patch configuration, (c) adjacent-double-patches configuration, and (d) disjoint-double-patches configuration.

surfaces from each of the testing configurations and compared them against the reference surface reconstructed under the reference configuration. We represented each reconstructed surface as a height function in a Cartesian coordinates: $z = f(x, y)$ for the convenience of further evaluation. Example reconstructed surfaces are compared in Fig. 4, where smooth and

continuous surfaces were reconstructed despite the presence of different noise and missing measurement levels in their corresponding point clouds. We quantitatively evaluated the reconstruction accuracy by comparing the reconstructed surfaces under different testing configurations against the reference reconstructed surface based on RMSE. Specifically, on a regular

FIG. 4. Example reconstructed surfaces under different phantom configurations (represented as height functions). (a) Reference configuration, (b) lower-left-patch configuration, (c) adjacent-double-patches configuration, and (d) disjoint-double-patches configuration.

TABLE I. Error statistics of the reconstructed surfaces under different testing configurations.

	Lower-left-patch configuration	Adjacent-double-patches configuration	Disjoint-double-patches configuration
RMSE (mm)	0.54	0.60	0.55
S.D. (mm)	0.50	0.55	0.57

grid with size $m \times n$, we calculated RMSE and standard deviation as the following: $\text{RMSE} = \sqrt{(1/mn) \sum_i^m \sum_j^n e_{i,j}^2}$ and $\text{S.D.} = \sqrt{(1/mn) \sum_i^m \sum_j^n (e_{i,j} - ((\sum_i^m \sum_j^n e_{i,j})/mn))^2}$, where $e_{i,j}$ is the pointwise error defined as $e_{i,j} = z_{\text{ref}}(i,j) - z_{\text{test}}(i,j)$, with z_{ref} and z_{test} being the height functions of the reconstructed reference and testing surface, respectively. The error statistics of the reconstructed surfaces under different testing configurations are reported in Table I. The reconstruction errors from different configurations were all within 1 mm based on RMSE, demonstrating quantitatively the ability of the proposed method in faithfully reconstructing underlying surfaces even under noise and missing measurements.

3.B. Surface reconstruction on human subject

As shown in Fig. 5(a), missing measurements often occur during VisionRT acquisition due to lighting conditions and/or

the presence of irregular patient body surface structure such as dense hair. Figures 5(b)–5(d) illustrate the reconstruction process of our proposed method on one example point cloud. A bounding box was first placed and evolved at a coarse grid level, whose result was used as the initialization for the fine level evolution. The final surface reconstruction result was shown in Fig. 5(d), where a smooth and continuous surface was reconstructed.

For local geometric characteristic comparison between the reconstructed surface and the CT surface, the local chest ROI was also represented as a height function in a Cartesian coordinates: $z = f(x, y)$, as illustrated in Figs. 6(a) and 6(b). Numerically, we calculated the mean curvature value at each grid location as $k = ((1 + f_y^2)f_{xx} - 2f_x f_y f_{xy} + (1 + f_x^2)f_{yy}) / 2(1 + f_x^2 + f_y^2)^{3/2}$. The mean curvature distributions of this example surface and CT surface are reported and compared in Figs. 6(c) and 6(d). The empirical mean and standard deviation of those two distributions are $(\mu_{\text{recon}} = -2.7 \times 10^{-3} \text{ mm}^{-1}, \sigma_{\text{recon}} = 7.0 \times 10^{-3} \text{ mm}^{-1})$ and $(\mu_{\text{CT}} = -2.5 \times 10^{-3} \text{ mm}^{-1}, \sigma_{\text{CT}} = 5.3 \times 10^{-3} \text{ mm}^{-1})$, respectively. We further repeated the same process on 100 snapshot point cloud measurements and thus recovered 100 reconstructed surface ROIs. Figure 6(e) illustrates the population distribution of the mean curvature values from the 100 reconstructed surfaces, with $(\mu_{\text{pop}} = -2.7 \times 10^{-3} \text{ mm}^{-1}, \sigma_{\text{pop}} = 6.9 \times 10^{-3} \text{ mm}^{-1})$. The agreement among these distributions demonstrates indirectly that the proposed method faithfully captures the geometric properties of the underlying surface.

FIG. 5. Example surface reconstruction process: (a) point cloud captured by the VisionRT system, (b) initialization at the coarse level, (c) evolution result at the coarse level, and (d) final surface reconstruction result at the fine level.

FIG. 6. Comparison of height functions and mean curvature distributions from human subject: (a) height function of the example reconstructed surface, (b) height function of the CT skin surface, (c) mean curvature distribution from the example reconstructed surface, (d) mean curvature distribution from the CT skin surface, and (e) mean curvature distribution from our 100 reconstructed surfaces.

4. DISCUSSION AND CONCLUSION

We have developed an accurate and reliable surface reconstruction method by optimizing a regularized fitting energy from point cloud measurements subject to noise and occlusion. On phantom data, we have achieved submillimeter reconstruction accuracy under different configurations, demonstrating quantitatively the faith of the proposed method in preserving local structural properties of the underlying surfaces despite the presence of noise and missing measurements. On point clouds from human subject, the reconstructed instantaneous surfaces have agreed well with patient surface extracted from high-quality CT in its local geometric characteristics.

The reconstruction of a continuous representation of the underlying patient surface from an unstructured (and possibly

randomly acquired) point cloud measurement eliminates the explicit indexing of the point orders. For the further purpose of dynamic motion monitoring and tracking, constructing correspondence between continuous surfaces is a more amicable and robust task¹⁸ than registering two point clouds that are intrinsically misaligned in acquisition. Our method can be naturally and conveniently extended to incorporate further regularizations such as learnt surface priors^{19,20} or surface normal information.¹³

The computational burden of the PDE evolution and the consequent inefficiency, due to the constraint imposed by the CFL condition, is alleviated partially with the narrow-band evolution scheme. Practically, the CFL condition can be further relaxed without incurring numerical instability, yet the upper bound of such relaxation required further testing

with more extensive data set. Our current work does not quite meet the rate for real-time application yet, but is still useful for constructing dynamic shape libraries and understanding the manifold which captures the surface variations from one's respiration.²⁰ We are also actively working on GPU-based implementation and parallel computation^{21,22} for further speedup.

ACKNOWLEDGMENTS

This work is supported in part by NIH Grant No. R01 CA169102-02. We also acknowledge the support from Vision RT.

^{a)}Author to whom correspondence should be addressed. Electronic mail: druan@mednet.ucla.edu

¹M. G. Herman, J. M. Balter, D. A. Jaffray, K. P. McGee, P. Munro, S. Shalev, M. Van Herk, and J. W. Wong, "Clinical use of electronic portal imaging: Report of AAPM Radiation Therapy Committee Task Group 58," *Med. Phys.* **28**(5), 712–737 (2001).

²T. H. Wagner, S. L. Meeks, F. J. Bova, W. A. Friedman, T. R. Willoughby, P. A. Kupelian, and W. Tome, "Optical tracking technology in stereotactic radiation therapy," *Med. Dosim.* **32**(2), 111–120 (2007).

³D. Létourneau, J. W. Wong, M. Oldham, M. Gulam, L. Watt, D. A. Jaffray, J. H. Siewerdsen, and A. A. Martinez, "Cone-beam-CT guided radiation therapy: Technical implementation," *Radiother. Oncol.* **75**(3), 279–286 (2005).

⁴J. L. Peng, D. Kahler, J. G. Li, S. Samant, G. Yan, R. Amdur, and C. Liu, "Characterization of a real-time surface image-guided stereotactic positioning system," *Med. Phys.* **37**(10), 5421–5433 (2010).

⁵O. Gopan and Q. Wu, "Evaluation of the accuracy of a 3D surface imaging system for patient setup in head and neck cancer radiotherapy," *Int. J. Radiat. Oncol., Biol., Phys.* **84**(2), 547–552 (2012).

⁶G. Gao, S. Tarte, A. King, Y. Ma, P. Chinchapatnam, T. Schaeffter, R. Razavi, D. Hawkes, D. Hill, and K. Rhode, "Validation of the use of photogrammetry to register pre-procedure MR images to intra-procedure patient position for image-guided cardiac catheterization procedures," *Proc. SPIE* **6918**, 69181Q (2008).

⁷D. F. Rogers, *An Introduction to NURBS: With Historical Perspective* (Elsevier, Amsterdam, 2000).

⁸N. Amenta and M. Bern, "Surface reconstruction by Voronoi filtering," *Discrete Comput. Geom.* **22**(4), 481–504 (1999).

⁹H. Edelsbrunner and E. P. Mücke, "Three-dimensional alpha shapes," *ACM Trans. Graphics* **13**(1), 43–72 (1994).

¹⁰H.-K. Zhao, S. Osher, and R. Fedkiw, "Fast surface reconstruction using the level set method," in *Proceedings IEEE Workshop on Variational and Level Set Methods in Computer Vision* (IEEE, Vancouver, BC, 2001), pp. 194–201.

¹¹C. Bajaj, *Introduction to Implicit Surfaces* (Morgan Kaufmann, Burlington, MA, 1997).

¹²S. J. Ruuth and B. Merriman, "A simple embedding method for solving partial differential equations on surfaces," *J. Comput. Phys.* **227**(3), 1943–1961 (2008).

¹³S. Osher and R. Fedkiw, *Level Set Methods and Dynamic Implicit Surfaces* (Springer Science & Business Media, Berlin, 2003), Vol. 153.

¹⁴S. Osher and J. A. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *J. Comput. Phys.* **79**(1), 12–49 (1988).

¹⁵H.-K. Zhao, S. Osher, B. Merriman, and M. Kang, "Implicit and nonparametric shape reconstruction from unorganized data using a variational level set method," *Comput. Vis. Image Understanding* **80**(3), 295–314 (2000).

¹⁶D. Adalsteinsson and J. A. Sethian, "A fast level set method for propagating interfaces," *J. Comput. Phys.* **118**(2), 269–277 (1995).

¹⁷J. A. Sethian, "A fast marching level set method for monotonically advancing fronts," *Proc. Natl. Acad. Sci. U. S. A.* **93**(4), 1591–1595 (1996).

¹⁸W. Liu and D. Ruan, "Estimating nonrigid motion from inconsistent intensity with robust shape features," *Med. Phys.* **40**(12), 121912 (13pp.) (2013).

¹⁹J. E. Solem and F. Kahl, "Surface reconstruction using learned shape models," *Adv. Neural Inf. Process. Syst.* **17**, 1305–1312 (2004).

²⁰W. Liu and D. Ruan, "A unified variational segmentation framework with a level-set based sparse composite shape prior," *Phys. Med. Biol.* **60**(5), 1865–1877 (2015).

²¹A. E. Lefohn, J. E. Cates, and R. T. Whitaker, "Interactive, GPU-based level sets for 3D segmentation," in *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2003* (Springer, Berlin, 2003), pp. 564–572.

²²M. Roberts, J. Packer, M. C. Sousa, and J. R. Mitchell, "A work-efficient GPU algorithm for level set segmentation," in *Proceedings of the Conference on High Performance Graphics* (Eurographics Association, Geneva, 2010), pp. 123–132.