Anisotropic diffusion of surface normals for feature preserving surface reconstruction

Tolga Tasdizen *Ross Whitaker*

University of Utah, School of Computing
Technical Report UUCS-03-007

Also submitted to the 4th Int. Conference on 3-D Digital Imaging and Modeling for review

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

April 18, 2003

Abstract

For 3D surface reconstruction problems with noisy and incomplete range data measured from complex scenes with arbitrary topologies, a low-level representation, such as level set surfaces, is used. Such surface reconstruction is typically accomplished by minimizing a weighted sum of data-model discrepancy and model smoothness terms. This paper introduces a new nonlinear model smoothness term for surface reconstruction based on variations of the surface normals. A direct solution requires solving a fourth-order partial differential equation (PDE), which is very difficult with conventional numerical techniques. Our solution is based on processing the normals separately from the surface, which allows us to separate the problem into two second-order PDEs. The proposed method can smooth complex, noisy surfaces, while preserving sharp, geometric features, and it is a natural generalization of edge-preserving methods in image processing, such as anisotropic diffusion.

Chapter 1

Introduction

The precision of range measurement systems, such as time-of-flight laser range finders, has been increasing while their price drops. If combined with a well-founded method for surface reconstruction, these improvements could make capturing 3D shape as ubiquitous as photography. However, significant challenges to surface reconstruction, such as measurement noise and variations in measurement density, remain. This paper addresses the problem of preserving geometric features, i.e. edges, corners and junctions on surfaces, in *full 3D reconstructions* of complex scenes from multiple, noisy range images. Figure 1.1(a) illustrates a typical range image. Measurement noise and occlusions (between the file cabinet and the chair) can be observed in this data. Figure 1.1(b) illustrates the reconstruction with the proposed method from a similar view point. The result is an improvement over the state-of-the-art full 3D reconstructions for complex scenes: creases and corners at the interchapters of the various planes in the scene have been preserved while measurement noise has been effectively eliminated.

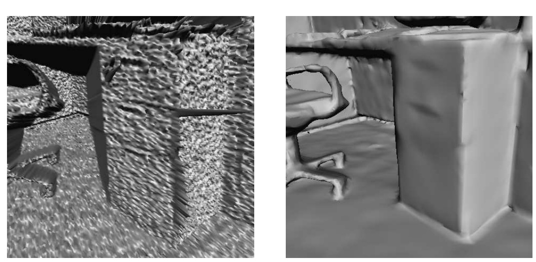

(a) (b)

Figure 1.1: (a) A range image surface, and (b) feature preserving surface reconstruction.

Full 3D reconstruction recovers a view-independent surface model from multiple registered range images. This problem is distinct from *depth reconstruction*, also known as $2\frac{1}{2}$D reconstruction, which recovers structure from only one point of view or a stereo pair of images [1, 2, 3]. Depth reconstruction does not produce a model that makes sense when viewed from different viewpoints or when determining inherently 3D properties, such as volume. For instance, occluded portions of the scene in Figure 1.1(a) are present in the reconstructed model because multiple range images are used in full 3D reconstruction. This result would not have been possible with depth reconstruction methods. The full 3D problem is not a mere extension of depth reconstruction because it lacks the following properties of the latter: the depth map has a well-defined topology (a function of two variables) and there is a one-to-one correspondence between the measurements and the positions on the model.

Recovering a full 3D model from a set of noisy 2D range images is an ill-posed inverse problem. Hence, the estimator can not depend solely on the input data, and requires regularization. Regularization reduces the effects of measurement noise and fills surfaces in a plausible manner where there is no data from any of the range images by placing additional constraints on the reconstructed surfaces. This problem has been approached in the computer vision literature mainly as a problem of finding sets of geometric primitives that best represent the objects being measured [4, 5, 6, 7]. Primitives typically have only a few shape parameters, *i.e.*, height and radius for a cylinder; therefore, impose their own structure on to the data. In this case regularization is inherent to the surface model. Such approaches are suitable for higher-level tasks of object recognition and decomposition into parts; however, they are limited to modeling relatively simple objects.

An alternative is to use level set surfaces [8], which are a non-parametric shape representation. Level set surfaces can be used to recover scenes with arbitrarily complex geometry and topology. The reconstructed level set models are not limited to prescribed topologies, and can adapt to the topology of the measured scenes automatically. However, level set surfaces do not have a rigid shape structure, and therefore, regularization must be performed explicitly.

We formulate surface estimation (reconstruction) in a variational energy optimization framework. Variational methods typically minimize an energy function, which is a weighted sum of an input data-model discrepancy term and a model smoothness term, with respect to the model. Then the surface estimator is defined as

$$\hat{S}(D, \alpha) = \arg\inf_{S} \left[F(S, D) + \alpha P(S)\right] \tag{1.1}$$

where α determines the relative weights of the terms. The input data-model discrepancy term, $F(S, D)$, forces the surface estimator to be “close” to the measured data. The model smoothness term, $P(S)$, provides regularization.

This paper studies the model smoothness term. Surface area is a simple choice and has been used extensively in previous work [9, 10]. This is based on the assumption that among surfaces with similar data-model discrepancy measures, those that have smaller area are simpler than surfaces of larger area, and therefore have a higher chance of occurrence in reality. Despite its simplicity, surface area as a measure of smoothness has significant drawbacks, such as pinching of thin structures. We argue that measuring the variation of the surface normal vectors offers a better and more flexible alternative. Specifically, this family of smoothness functions offers an elegant and mathematically correct generalization of Perona & Malik's (P&M) anisotropic diffusion [12] to surface reconstruction. This generalization allows us to preserve important shape features such as edges, corners and junctions in reconstructed scenes while effectively eliminating measurement noise and other artifacts.

The remainder of this paper is organized as follows. Chapter 2 discusses the related work in the literature. Chapter 3 discusses level set surface reconstruction and proposes a generalization of P&M's edge detection method as a level set surface regularization term. Chapter 4 solves the proposed regularization term as a level set motion. Chapter 5 demonstrates the quantitative advantages of the proposed method and provides examples of reconstruction of real, complex scenes from noisy range data. Chapter 6 summarizes the contributions of this paper and discusses possibilities for future research directions.

Chapter 2

Related Work

As mentioned earlier, computer vision researchers approach surface reconstruction either as a depth reconstruction problem [1, 2, 3] or a view independent problem. Earlier literature on the view independent problem focuses mainly on high-level approaches that fit various geometric primitives to the data [4, 5, 6, 7]. Both problems are different from a view independent reconstruction of complex scenes with low-level shape representations. In response to the advances in 3D range sensing device, researchers in a variety of fields have started to study this problem. In computer graphics, the accuracy of the data exceeds the requirements of the application, and therefore the problem is treated as a problem of assembling pieces of noiseless information. For instance, Turk and Levoy [13] propose a "zippering" algorithm for combining triangle meshes of range maps of an object from different points of view. Curless and Levoy [14] take into account measurement noise by averaging range information in a volumetric representation. However, their method is not based on statistics the scanner and model geometry.

Several authors [9, 10] demonstrate the advantages of using level set methods for reconstructing complex shapes. They use mean curvature flow, a second-order level set partial differential equation (PDE), to obtain a smooth solution. This flow is the gradient descent for the first variation of surface area [15]. Hence, for regularization purposes, both approaches are formulated with the surface area model smoothness term in the estimator (1.1). Mean curvature flow suffers from several problems including volume shrinkage, pinching of thin structures, and elimination of sharp features [11]. These problems can be observed in the results presented in Chapter 5.

To alleviate the problems associated with mean curvature flow, several authors have proposed smoothing level set surfaces by modified second-order flows that use weighted combinations of principle curvatures. For instance, Lorigo *et al.* [16] propose a smoothing flow that uses the minimum curvature for tubular structures. Clarenz *et al.* [17] propose an anisotropic surface mesh diffusion as a modified second order flow, but this modified PDE lacks a variational basis, and therefore is not useful for surface reconstruction. Stevenson and Delp propose minimizing total curvature, which is the surface integral of the sum of the squared principal curvatures, for regularizing depth reconstruction and parametric Monge patches [18]. This method does not apply to level set surfaces. Furthermore, the gradient descent flow for total curvature is a fourth-order PDE that is computationally expensive and unstable to compute. Hence, Stevenson and Delp use a non-geometric thin plate approximation.

In previous work [11], we propose a two-step approach to surface smoothing: (1) operate on the normal map of a surface, and (2) refit a surface to the processed normals. In this paper, we show that an quadratic measure on the variations in surface normals variations is equivalent to total curvature. Unlike using calculus of variations on the total curvature of the surface directly, our formulation in terms of the normals yields a computationally tractable and stable gradient descent flow. Furthermore, we also propose a robust measure of surface normal variations that allows a novel generalization of P&M feature preserving anisotropic image diffusion [12] to surface reconstruction.

Our energy optimization approach can also be stated in terms of Bayesian maximum a posteriori (MAP) estimation. According to Bayes rule, MAP estimators maximize the logarithm of the product of two distinct probabilities: the likelihood of the measurement data conditioned on the surface model and the prior probability distribution of the model. In the estimator given by (1.1), the input data-model discrepancy term corresponds to the logarithm of the conditional likelihood and the model smoothness term corresponds to the logarithm of the prior [10]. This brings up a connection with other approaches that use shape priors for active contour and level set models [19, 20, 21, 22]. However, in all of these works, the "prior" is a global description of expected shape(s) learned from a training set. In the segmentation and/or registration stage, the shape model is forced to be a rigid transformation of the learned prior shape with some tolerance for local variations. This approach is useful for reconstructing/segmenting specific classes of shapes, such as cortical surfaces from head MRI data; however, learned priors can not be used for regularization of reconstructed models in a general setting. The commonality between general shapes is not on a global level, but on a lower level, such as probability distributions for surface normal variations, which we use in this paper. The quadratic and robust penalty term on surface normal variations yield generic isotropic and anisotropic low level shape priors, respectively. These priors are not learned from a training set.

Chapter 3

Variational Implicit Surface Reconstruction

A deformable surface, $S(t)$, can be represented as the zero level set of a higher dimensional embedding function $\phi : \mathbb{R}^3 \times \mathbb{R} \to \mathbb{R}$, $S(t) = \{\mathbf{x}(t) \in \mathbb{R}^3 \mid \phi(\mathbf{x}(t), t) = 0\}$, where t is the evolution parameter (time). Surfaces defined in this way divide a volume into two parts: inside ($\phi > 0$) and outside ($\phi < 0$). The family of PDEs that describe motions of S via $\partial\phi/\partial t$, and the upwind scheme for solving them on a discrete grid is the methods of *level sets* [8].

The surface reconstruction energy of (1.1) can be expressed as a function of the data D, and the level set model ϕ. Whitaker [10] formulates the input data-model discrepancy part of this energy as a volumetric integral

$$F(\phi, D) = \int_\Omega G(\mathbf{x}, D) H(\phi(\mathbf{x}))\, d\mathbf{x}, \tag{3.1}$$

where Ω is the volumetric domain swept out by the range images, G is an accurate line-of-sight error function, and H is the Heaviside function [23]. Minimizing (3.1) by itself would correspond to a *maximum likelihood estimator*, but a *maximum apriori estimator* is implemented by using surface area as a model smoothness term:

$$P(\phi) = \int_\Omega ||\nabla\phi(\mathbf{x})||\ d\mathbf{x}. \tag{3.2}$$

Then, the gradient descent for ϕ is

$$\partial\phi/\partial t = ||\nabla\phi(t)|| \left(G(\mathbf{x}, D) + \alpha H(\mathbf{x})\right), \tag{3.3}$$

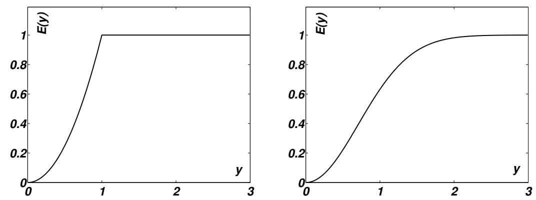

(a) (b)

Figure 3.1: (a) A quadratic penalty term with a hard cutoff, (b) an exponential penalty term corresponding to P&M's anisotropic diffusion.

where $H = \nabla \cdot (\nabla\phi / ||\nabla\phi||)$ is mean curvature [15]. Using discrete time steps, the model is evolved as $\phi(t + \Delta t) = \phi(t) + \Delta t \; \partial\phi/\partial t$. The steady-state solution of this evolution is the surface estimator: $\hat{S} = \left\{\mathbf{x} \in \mathbb{R}^3 \mid \lim_{t\to\infty} \phi\left(\mathbf{x}, \mathbf{t}\right) = 0\right\}$.

In this paper, we use the same data-model discrepancy term and the initialization method for $\phi(t = 0)$. However, we propose a better model smoothness term based on measures on the variation of the surface normal vectors, $\mathbf{N}$,

$$P(\phi) = \int_{\Omega} E\left(||\nabla_\phi \mathbf{N}(\mathbf{x})||_{fr}\right) ||\nabla\phi(\mathbf{x})|| \; d\mathbf{x}, \tag{3.4}$$

where $\nabla_\phi \mathbf{N}$ is the matrix whose rows are the gradient vectors of the respective components of $\mathbf{N}$ intrinsic to the isosurfaces of ϕ. The *Frobenius matrix norm*, i.e. the square root of the sum of squares of the matrix elements, is denoted by $||\cdot||_{fr}$. Notice that if $E(y) = 1$, (3.4) reduces to surface area (3.2). On the other hand, if we choose a quadratic measure, $E(y) = y^2$, we obtain the sum of squared principal curvatures (total curvature) [11]. The resulting gradient descent PDE is analogous to running the heat equation PDE on the surface normals. Figure 4.1 (a) illustrates a noisy field of unit vectors which are the gradient directions for the noisy distance transform of the polygon shown. Figure 4.1 (b) demonstrates the result of smoothing with the PDE derived above. Similar to the heat equation, this flow eliminates both noise and sharp discontinuities in the data. Our goal is to preserve these discontinuities while penalizing noisy variations of the normals elsewhere. This is similar to segmenting the normal map.

Mumford and Shah formulate the problem of image segmentation in a variational framework [24]. The Mumford-Shah energy is the sum of three terms: (i) the data-model discrepancy, (ii) quadratic penalty on model image smoothness over the image domain except on a set of discontinuities modeled by a binary image, and (iii) the length of the discontinuities in that binary image. The sum of the latter two terms correspond to using the penalty function, $E(y)$, shown in Figure 3.1(a). The existence of a binary model poses serious difficulties in the estimation process. Nordstrom [25] and Black *et al.* [26] show that P&M anisotropic diffusion approach to edge detection [12] is equivalent to using a corresponding robust penalty term in the Mumford-Shah segmentation framework. This penalty term,

$$E(y) = 1 - e^{-\frac{y^2}{\mu^2}}, \tag{3.5}$$

shown in Figure 3.1(b), avoids using a binary discontinuity model. The parameter μ controls the degree of edge preservation. The above result can readily be generalized to normal vectors by substituting $y = ||\nabla_\phi \mathbf{N}(\mathbf{x})||_{fr}$. Figure 4.1(c) illustrates the result of smoothing the noisy image of normals with this penalty term. The discontinuities in the normal directions between the quadrants are preserved while the noise is smoothed. Minimization of energies of the form (3.4) require solving fourth-order level set PDEs, a computationally unstable and expensive task. The next chapter introduces a method for breaking the solution into two second-order PDEs that can be efficiently solved.

Chapter 4

Level set motion via normal map diffusion

In equation (3.4), $\nabla_\phi \mathbf{N}$ is the matrix whose rows are the gradient vectors of the respective components of $\mathbf{N}$ intrinsic to the isosurfaces of ϕ. By intrinsic, we mean that when using implicit representations one must account for the fact that derivatives of functions defined on the surface are computed by projecting their 3D derivatives onto the surface tangent plane. The 3×3 projection matrix for the implicit surface normal is $\mathbf{P} = \nabla\phi\otimes\nabla\phi/\left\|\nabla\phi\right\|^2$, where $\otimes$ is the tensor product. The projection matrix onto the surface tangent plane is $\mathbf{I}-\mathbf{P}$, where $\mathbf{I}$ is the identity matrix. Then the intrinsic gradient of the normals can be defined using this projection operator and the regular Euclidean gradient $\nabla_\phi \mathbf{N} = \nabla\mathbf{N}\left(\mathbf{I}-\mathbf{P}\right)$.

Given an initialization for ϕ with the methods described in [10], we compute the surface normals $\mathbf{N} = \nabla\phi/\left\|\nabla\phi\right\|$. Then, to avoid solving fourth-order level set PDEs directly,

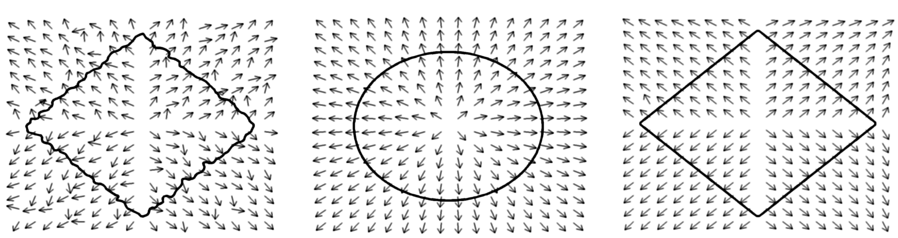

(a) (b) (c)

Figure 4.1: (a)Normals computed from a noisy distance transform of a polygon, normals smoothed with (b) the quadratic penalty term, and (c) the robust penalty term.

we decouple $\mathbf{N}$ from ϕ. In other words, we fix ϕ (the surface shape) as we process the normals to minimize the energy given by (3.4). Solutions to constrained minimization for unit vectors on an implicit surface are discussed in [27]; however, the goal in that work is to not to smooth surfaces, but to diffuse general vector functions on surfaces. We implement the constrained minimization with the following second-order PDE:

$$\frac{\partial \mathbf{N}}{\partial t} = -\left(\mathbf{I} - \mathbf{N} \otimes \mathbf{N}\right) \nabla \cdot \left[E'\left(||\nabla_\phi \mathbf{N}||_{fr}\right) \nabla_\phi \mathbf{N}\right] \tag{4.1}$$

where $E'(y)$ is the derivative of $E(y)$. For the penalty term given in (3.5), $g = e^{-\frac{y^2}{\mu^2}}$ up to a constant multiplicative factor. Figure 4.1 (a) illustrates a noisy field of unit vectors. Figure 4.1 (b) and (c) demonstrate results of smoothing with the choices of $E(y)$ discussed in Chapter 3.

The next step is to relate the deformation of the level sets of ϕ to the evolution of $\mathbf{N}$. Once more, using a variational approach, we can manipulate ϕ so that it fits the new normal field by minimizing a penalty function,

$$\mathcal{D}(\phi) = \int_\Omega \left[\sqrt{\nabla\phi \cdot \nabla\phi} - \nabla\phi \cdot \nabla\mathbf{N}\right] d\mathcal{S}, \tag{4.2}$$

that quantifies the discrepancy between the gradient vectors of ϕ and the target normal map Ballester *et al.* [28] use the same function for filling in missing regions in images by joint interpolation of the image intensity and its gradient. The first variation of this function with respect to ϕ is

$$\frac{d\mathcal{D}}{d\phi} = -\nabla \cdot \left[\frac{\nabla\phi}{||\nabla\phi||} - \mathbf{N}\right] = -\left[H^\phi - H^{\mathbf{N}}\right] \tag{4.3}$$

where H^ϕ is the mean curvature of the level set surface and H^{N} is half the divergence of the normal map. Figure 4.1(b) and (c) illustrate the curves refitted to the smoothed normal fields with this approach. Finally, the gradient descent for the surface reconstruction with the model smoothness energy (3.4) is

$$\partial\phi/\partial t = ||\nabla\phi(t)|| \left(G(\mathbf{x}, D) + \alpha\left(H^\phi(\mathbf{x}) - H^{\mathbf{N}}(\mathbf{x})\right)\right), \tag{4.4}$$

which is similar to (3.3), but has a different smoothing term.

The flow chart for the algorithm is shown in Figure 4.2. We have derived a gradient descent for the normal map that minimize the energy functions of the form (3.4). The normals processing stage of the algorithm computes the gradient descent for the normals defined in (4.1) for a fixed number of iterations (25 for the experiments in this paper). Hence, we avoid evolving evolving the normals too far away from their initialization from ϕ. The surface fitting to the the combined normal map and data terms is given as a gradient descent in (4.4).

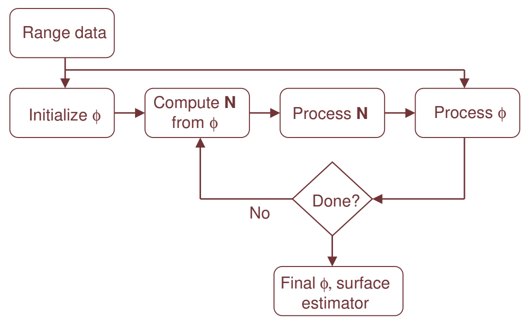

Figure 4.2: Surface reconstruction flow chart.

This stage of the algorithm is run until the discrepancy measure (4.2) between the new normals and ϕ ceases to decrease, which signals the need for another round of processing the normal vectors. The overall algorithm repeats these two steps to minimize the surface reconstruction energy in terms of ϕ until the RMS value for $\partial\phi/\partial t$ becomes small (less than 10^{-6}), which signals convergence. This algorithm consists of solving two second-order PDEs in series instead of a direct fourth-order PDE, which makes it computationally tractable. We show the relationship of this algorithm to solving the direct fourth-order PDE in [11].

Chapter 5

Experiments

In this chapter, we compare reconstructions with proposed the model smoothness energies against reconstructions with the standard surface area energy. For the proposed family of energies (3.4), we will call the choice of $E(y) = y^2$ and $E(y) = 1 - e^{-\frac{y^2}{\mu^2}}$, the isotropic and anisotropic reconstructions, respectively. Note that μ is fixed at 0.2 for all the experiments. Unlike, in P&M image diffusion, this parameter does not need to be changed for different surface reconstructions. In the context of P&M image diffusion, the units of μ are in gray levels; consequently, the optimal choice of μ is image dependent. In surface reconstruction, the units are in curvature which is data independent. This makes it possible to choose a μ value that gives consistent results over a broad range of surfaces.

We first experiment with geometric shapes for which we can construct analytical distance transforms. We use the following experiment setup:

1. Build range images by simulating the laser range finder located at several positions,
2. Add independent Gaussian noise to the range images,
3. Reconstruct a model, and
4. Compute the root mean square geometric distance, between the resulting model and the analytical shape.

The first shape we examine is a sphere with radius 1 unit. All other distances are relative to this measurement unit. For this experiment we simulate six range finders located at a distance of 3.5 units from the center of the sphere along the six cardinal directions. Independent Gaussian noise with a standard deviation of 0.1 units, is added to each range image. Figure 5.1(a) plots the RMS error, $\mathcal{E}$, against the logarithm of the regularization weight, $\log \alpha$, for the different reconstructions. The units on the *y-axis* are the same as the units used to described the size of the shape. It can be observed from Figure 5.1 that the limiting value for $\mathcal{E}$ as $\alpha \to 0$ is approximately 0.0125. This limit is the error obtained if surface reconstruction is performed without regularization. This error level is smaller than the noise added to the range images because of the averaging effect of using multiple range images. The anisotropic and the isotropic curvature priors at their optimal weight provide a 75% reduction on this error. On the other hand, surface area provides slightly better than a 50% reduction at its optimal weight. The shapes of the error plots is more important than the results at optimal choices of weight. The surface area prior performs poorly as α is increased beyond 1; this is due to the fact that the surface area prior causes shrinkage in the surface models. In practice, this will mean difficulties for the user in choosing a weight for surface area regularization that works different reconstruction scenarios. In contrast, the proposed reconstructions have relatively flat error plots. Isotropic reconstruction is as good as the anisotropic reconstruction because the sphere does not contain creases.

To examine the differences between isotropic and anisotropic reconstruction further, we experiment with a cube shape. In this experiment, we use 8 range finder locations (one in each octant). Figure 5.2 (a) and (b) show the original cube with sides 1 unit long, and the surface initialization from the noisy range images, respectively. Independent Gaussian noise with standard deviation 0.1 was added to the simulated data to create the noisy range images. The results (see Figure 5.2) with isotropic reconstruction have rounded corners compared to the successfully denoised, approximately piecewise planar results obtained with the anisotropic reconstruction.

The next example involves 12 real range scans of a room which were registered using the methods described in [10]. A close-up view of a portion of one of the range images and the result of anisotropic reconstruction are shown in Figure 1.1. The anisotropic reconstruction of the entire scene is shown in Figure 5.3. We now examine reconstructions of one of the chairs in this scene. Figure 5.4 (a), (c) and (e) illustrate the results obtained by qualitatively choosing good values for α. Figure 5.4(b), (d) and (f) illustrate the results if α is chosen to be 10 times this value. These results show that the anisotropic reconstruction produces the best results and is least sensitive to the choice of α. Another well known problem with surface area reconstruction can easily be observed in Figure 5.4(b); the beam connecting the base to the seat is being pinched-off. This experiment illustrates the importance of the anisotropic reconstruction for scenes with high curvature features and sharp creases.

Figure 5.5 illustrates anisotropic reconstructions of a vehicle. This example further illustrates the success of the anisotropic reconstruction in denoising data with sharp features.

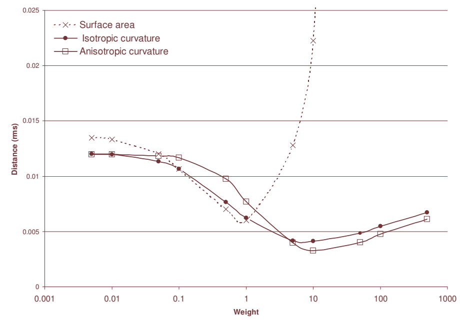

(a)

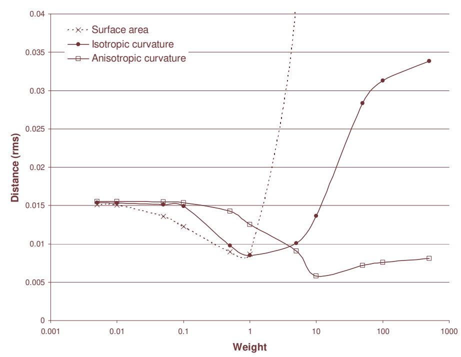

(b)

Figure 5.1: Rms distance between the reconstructed and the analytical surface for (a) the sphere, and (b) the cube.

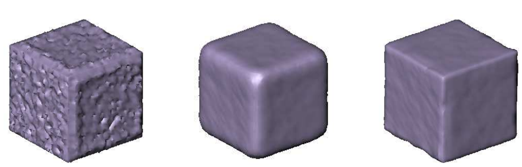

(a) (b) (c)

Figure 5.2: (a) Initialization from noisy data. Resulting model for (b) isotropic reconstruction, and (c) anisotropic reconstruction with $\alpha = 10$.

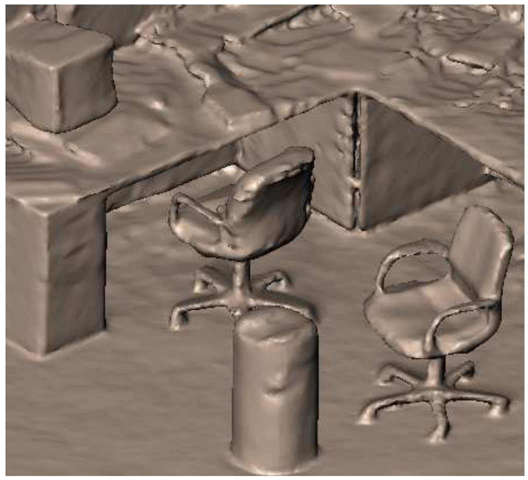

Figure 5.3: Anisotropic reconstruction.

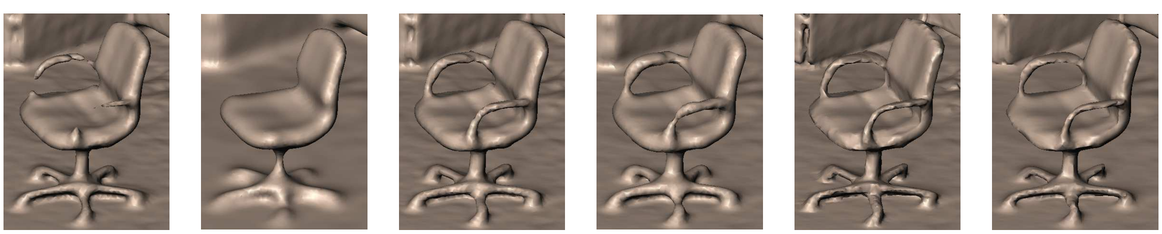

(a) (b) (c) (d) (e) (f)

Figure 5.4: Results for the surface area reconstruction with weights (a) 1, and (b) 10, isotropic reconstruction with weights (c) 1, and (d) 10, anisotropic reconstruction with weights (e) 1, and (f) 10.

The current shortcoming of this method is the computational speed which was approximately one hour on a *Intel Xeon 1.7Ghz Proc.* for the examples presented.

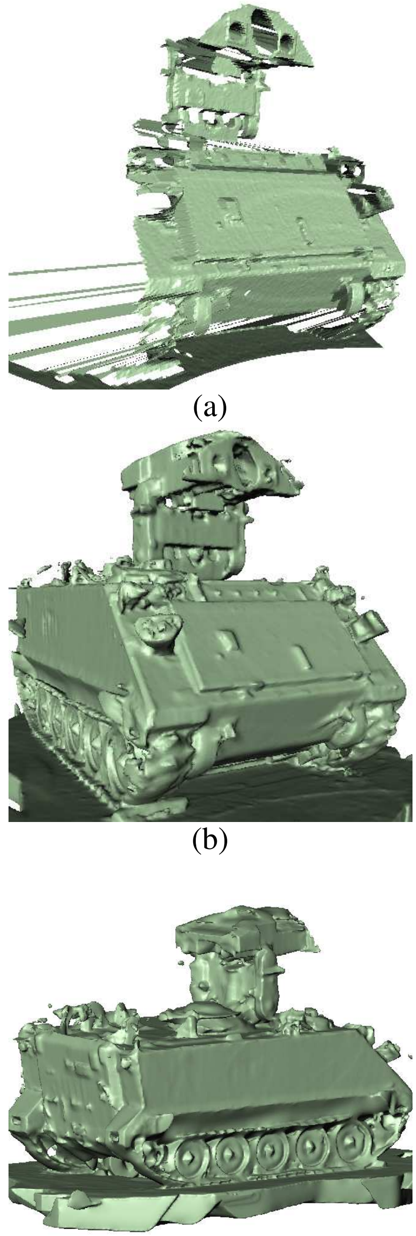

Figure 5.5: (a) One of 12 range images used in experiment, (b) result of feature preserving reconstruction from a similar view point, and (c) a different view point.

Chapter 6

Conclusion

We derive a variational generalization of P&M anisotropic diffusion for feature preserving surface reconstruction. This generalization is based on a robust penalty on surface normal vector variations, which is shown to have important advantages over using surface area and the quadratic penalty on surface normal vector variations for regularization. The data term is independent of the prior, the ideas introduced in this paper can be applied to other forms of surface reconstruction such as applications in tomography [29]. We use implicit surfaces, representing the implicit function on a discrete grid, modeling the deformation with the method of level sets. Therefore, the method applies equally well to surfaces that can be represented in a volume. The results shown in this paper are not possible with previous methods in the literature.

Measures on surface normal variations require solving fourth-order PDEs on level sets. However, by processing the normals separately from the surface, we can solve a pair of second-order equations instead of a fourth-order equation. This method is numerically more stable and computationally less expensive than solving the fourth-order PDE directly. The shortcoming of this method is the computation time; however, the process lends itself to parallelism, and therefore, the use of multi-threading.

Acknowledgements

This work is supported by the Office of Naval Research under grant #N00014-01-10033 and the National Science Foundation under grant #CCR0092065.

Bibliography

[1] W. E. L. Grimson, “From images to surfaces: A comp. study of the human early vision system”, MIT Press, 1981.

[2] D. Terzopoulos, “Multiresolution Computation of Visible-Surface Representations”, Ph.D. Dissertation, MIT, 1984.

[3] A. Blake and A. Zisserman, “Visual Reconstruction”, MIT Press, 1987.

[4] R. M. Bolle and D. B. Cooper, “On Optimally Combining Pieces of Information, with Application to Estimating 3-D Complex-object position from range data”, IEEE T. PAMI, 8, pp. 619–638, 1986.

[5] R. Bajcsy and F. Solina, “Three Dimensional Object Representation Revisited”, Proc. ICCV, pp. 231–240, 1987.

[6] A. P. Pentland, “Recognition by Parts”, Proc. ICCV, pp. 612–620, 1987.

[7] D. DeCarlo and D. Metaxas, “Adaptive Shape Evolution Using Blending”, Proc. ICCV, pp. 834–839, 1995.

[8] S. Osher and J. Sethian, ”Fronts Propogating with Curvature-Dependent Speed: Algorithms Based on Hamilton-Jacobi Formulations”, J. Comp. Physics, 79, pp. 12–49, 1988.

[9] R. Malladi, J. A. Sethian, and B. C. Vemuri, “Shape Modeling with Front Propagation: A Level Set Approach, IEEE T. PAMI, 17(2), pp. 158–175, 1995.

[10] R. T. Whitaker, “A Level-Set Approach to 3D Reconstruction From Range Data”, IJCV, 29(3), pp. 203–231, 1998.

[11] T. Tasdizen, R. Whitaker, P. Burchard, and S. Osher, “Geometric surface smoothing via anisotropic diffusion of normals”, Proc. IEEE Visualization, pp. 125–132, 2002.

[12] P. Perona and J. Malik, ”Scale space and edge det. using anisotropic diff.”, IEEE T. PAMI, 12(7), pp. 629–639, 1990.

[13] G. Turk and M. Levoy, “Zippered polygon meshes from range images”, Proc. SIGGRAPH’94, pp. 311–318, 1994.

[14] B. Curless and M. Levoy, ”A volumetric method for building complex models from range images”, Proc. SIGGRAPH’96, pp. 303–312, 1996.

[15] G. Sapiro, “Geometric Partial Differential Equations and Image Analysis”, Cambridge U. Press, 2001.

[16] L. Lorigo, O. Faugeras, E. Grimson, R. Keriven, R. Kikinis, A. Nabavi, and C.-F. Westin, “Co-Dimension 2 Geodesic Active Contours for the Segmentation of Tubular Strucures”, Proc. IEEE CVPR, 2000.

[17] U. Clarenz, U. Diewald, and M. Rumpf, “Anisotropic Geometric Diffusion in Surface Processing”, Proc. IEEE Visualization, pp. 397–405, 2000.

[18] R. L. Stevenson and E. J. Delp, “Viewpoint Invariant Recovery of Visual Surfaces from Sparse Data”, IEEE T. PAMI, 14(9), pp. 897-909, 1992.

[19] M. Leventon, E. Grimson and O. Faugeras, “Statistical Shape Influence in Geodesic Active Contours, Proc. IEEE CVPR, pp. I:316–322, 2000.

[20] Y. Chen, S. Thiruvenkadam, F. Huang, D. Wilson, and E. Geiser, “On the incorporation of shape priors into geometric active contours”, IEEE Workshop on Variational and Level Set Methods, pp. 145–152, 2001.

[21] A. Tsai, A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, A. Grimson, and A. Willsky, “Model-based Curve Evolution Technique for Image Segmentation”, Proc. IEEE Computer Vision and Pattern Recognition, pp. I:463–468, 2001.

[22] M. Rousson and N. Paragios, “Shape Priors for Level Set Representations”, ECCV, pp. 78–92, 2002.

[23] H.-K. Zhao, T. Chan, B. Merriman, and S. Osher, “ A Variational Level Set Approach to Multiphase Motion”, J. Comp. Physics, 127, pp. 179–195, 1996.

[24] D. Mumford and J. Shah, “Boundary Detection by Minimizing Functionals”, IEEE CVPR, 1985.

[25] N. Nordstrom, “Biased Anisotropic Diffusion—A unified Regularization and Diffusion Approach to Edge Detection”, Image and Vision Comp., 8(4), pp. 318–327, 1990.

[26] M. J. Black, G. Sapiro, and D. H. Marimont, “Robust Anisotropic Diffusion”, IEEE T. Image Processing, 7(3), pp. 421–432, 1998.

[27] M. Bertalmio, L.-T. Cheng, S. Osher, and G. Sapiro, “Variational methods and partial differential equations on implicit surfaces”, J. Comp. Physics., pp. 759–780, 2001.

[28] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, and J. Verdera, “Filling-In by Joint Interpolation of Vector Fields and Gray Levels”, IEEE T. Image Processing, 10(8), pp. 1200–1211, 2001.

[29] V. Elangovan and R. Whitaker, ”From Sinograms to Surfaces: A Direct Approach to Segmenting Tomographic Data”, MICCAI’01. pp. 213–223, 2001.