

Higher-Order Nonlinear Priors for Surface Reconstruction

Tolga Tasdizen, *Member, IEEE*, and Ross Whitaker, *Member, IEEE*

Abstract—For surface reconstruction problems with noisy and incomplete range data, a Bayesian estimation approach can improve the overall quality of the surfaces. The Bayesian approach to surface estimation relies on a likelihood term, which ties the surface estimate to the input data, and the prior, which ensures surface smoothness or continuity. This paper introduces a new high-order, nonlinear prior for surface reconstruction. The proposed prior 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. An exact solution would require solving a fourth-order partial differential equation (PDE), which can be difficult with conventional numerical techniques. Our approach is to solve a cascade system of two second-order PDEs, which resembles the original fourth-order system. This strategy is based on the observation that the generalization of image processing to surfaces entails filtering the surface normals. We solve one PDE for processing the normals and one for refitting the surface to the normals. Furthermore, we implement the associated surface deformations using level sets. Hence, the algorithm can accommodate very complex shapes with arbitrary and changing topologies. This paper gives the mathematical formulation and describes the numerical algorithms. We also show results using range and medical data.

Index Terms—Surface reconstruction, robust estimation, anisotropic diffusion, level sets.

1 INTRODUCTION

THIS paper addresses the problem of Bayesian 3D surface reconstruction. Specifically, we consider reconstruction from multiple registered range images, but we also demonstrate the application of the methods developed in this paper to surface reconstructions from magnetic resonance imaging (MRI) data. The importance of high-quality surface reconstructions from range data is growing as range measurement technologies become more accurate and affordable. Despite the increased accuracy in these measurement devices, significant challenges to surface reconstruction remain. In particular, surface reconstruction is complicated by measurement noise and variations in measurement density. For instance, the measurement density is high at regions of the surface that are visible in multiple images. By contrast, parts of the surface that are occluded in all of the images have no data.

We model the uncertainties arising from measurement noise, overlapping measurements, and occlusion in a Bayesian framework, which formulates surface reconstruction as the maximization of a posterior probability function. According to Bayes rule, maximum a posteriori (MAP) estimators maximize the product of two distinct probabilities: the likelihood of the measurement data conditioned on the surface model and the prior probability distribution for the model. We refer to these two probabilities, respectively, as the *likelihood* and the *prior*. A major challenge of Bayesian surface reconstruction is the determination of these two quantities.

This paper focuses on the prior probability distribution. Specifically, we attempt to correctly generalize statistical image reconstruction strategies [1], [2], [3], [4], [5] to surface reconstruction, and we propose higher-order nonlinear feature-preserving surface priors as the solution to this problem.

In a variational optimization framework, penalty functions on the surface curvature give rise to fourth-order partial differential equations (PDE). In previous work [6], we propose that the natural generalization of image processing to surfaces occurs via the *surface normal vectors* and show results for postprocessing of noisy surfaces. For example, a smooth surface is one that has smoothly varying normal vectors. Our strategy is to use a two-step approach: 1) operate on the normal map of a surface to reduce its curvature and 2) fit a surface to the processed normals and the range data. Iterating this two-step process, we can efficiently implement a flow that resembles a fourth-order PDE. In this light, the differences between surface processing and image processing are threefold. Normals exist on a manifold (the surface) and cannot be processed using a flat metric, as is typically done with images. Processing techniques must also accommodate vector valued normals that are constrained to be unit length. Because normals are coupled with the surface shape, the normals should drag the surface along as their values are modified during processing.

The remainder of this paper is organized as follows: Section 2 presents a discussion of the related work in the literature. Section 3 gives a brief summary of Bayesian image and surface reconstruction formulations. In Section 4, we formulate higher-order feature preserving surface energy functions that give rise to a useful family of priors and, in Section 5, we introduce a strategy to solve them. Section 6 demonstrates results and compare them to results from previous methods, and discuss some properties of the

• The authors are with the School of Computing, University of Utah, 50 S. Central Campus Drive, Salt Lake City, UT 84112-9205.
E-mail: {tolga, whitaker}@cs.utah.edu.

Manuscript received 22 Oct. 2002; revised 10 Sept. 2003; accepted 11 Nov. 2003.

Recommended for acceptance by J.R. Beveridge.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 117162.

proposed approach. Section 7 presents conclusions and directions for future research possibilities.

2 RELATED WORK

Surface reconstruction methods can roughly be classified into high-level and low-level approaches. The high-level approaches are generally formulated as a problem of finding sets of geometric primitives that best represent the objects being measured [7], [8], [9], [10], [11]. High-level descriptions of scenes are useful for certain computer vision tasks such as object recognition; however, scenes described by simple primitives have limited complexity. Low-level methods are based on either explicit models such as surface meshes or implicit models such as level sets. The literature demonstrates methods for zipping together meshes that describe different sides of an object [12], and fitting deformable models that expand inside a sequence of range images [13]. These approaches are suitable for use with high quality range images with relatively small amounts of noise. Several authors propose volumetric methods [14], [15], [16], [17], [18] that combine range measurements in a volumetric medium. The reconstructed surfaces are the zero-sets, or isosurfaces, of such volumes.

In this paper, we utilize the level set-based Bayesian range map registration and surface reconstruction framework developed by Whitaker and Gregor [17], [19], [18]. This strategy uses maximum likelihood parameter estimation to register the views before combining multiple range images via a level set implementation that can represent any solid object, regardless of shape and topology. The combination of a Bayesian formulation with level set methods is effective for complex, noisy scenes. The Bayesian formulation requires the determination of the likelihood probability density function, which depends on the range sensor properties, and the prior probability density function, which models the space of the physical objects being measured. The role of the likelihood term is to force the fitted models to be accurate representations of the measured data. Whitaker and Gregor derive the likelihood term from a line-of-sight error formulation, which is shown to accurately model laser range finders [17], [18]. This paper seeks to improve the prior term which serves to eliminate the measurement noise and other artifacts in the fitted model by requiring it to adhere to certain expectations about the application domain or scene.

The use of priors have been extensively investigated in related works in image processing. For instance, Mumford and Shah formulate the problem of image segmentation in a variational framework with a Bayesian rationale [1], [2]. Nordstrom [3] investigates the relationship of the Perona & Malik (P&M) anisotropic diffusion approach to edge detection [20] to the Mumford-Shah variational strategy. Other authors [21], [5], [22] also present an unified view of the reconstruction, nonlinear diffusion, and robust statistics approaches. Our method is motivated by Nordstrom's *biased anisotropic image diffusion*.

In contrast to the research in image processing, investigation of more effective surface priors has not been emphasized in previous surface reconstruction literature. This is partly due to the prevalent strategy of fitting shape primitives to data in computer vision. Primitives usually have only a few shape parameters, i.e., height and radius for a cylinder. They impose their own structure on to the data and act as our prior

belief about the contents of a scene and, hence, there is no need for a separate prior probability term in the reconstruction. Loosely structured explicit models such as surface meshes and implicit models such as level set surfaces of volumes, can represent complicated scenes; however, these low-level models do not impose a rigid shape structure. Therefore, prior probability terms must be included in surface reconstruction with such models to reduce the effects of noise and fill in surfaces where there is no data.

Surface area penalty serves as a simple prior for surface reconstruction [17], [19], [18], and a gradient descent on the surface area energy results in mean curvature flow (MCF). However, in the context of surface reconstruction, MCF suffers from several problems including volume shrinkage and elimination of sharp features (creases). A great deal of research focuses on modified second-order flows that produce better results than MCF. Using level set methods, several authors have proposed smoothing surfaces by weighted combinations of principal curvatures. For instance, Whitaker [23] proposes a nonlinear reweighting scheme that favors the smaller curvature and preserves cylindrical structures. Lorigo et al. [24] propose a smoothing by the minimum curvature.

A similar set of curvature-based algorithms have been developed for surface meshes. For instance, Taubin [25] proposes a nonshrinking Gaussian smoothing. Clarenz et al. [26] propose a modified MCF as an anisotropic diffusion of the surface. They threshold a weighted sum of the principal curvatures to determine the surface locations where edge sharpening is needed. Tangential displacement is added to the standard MCF at these locations for sharpening the edges. Another mesh-based modified MCF is proposed in [27] where a threshold on the mean curvature is used to prevent over-smoothing. Taubin proposes a "linear anisotropic Laplacian operator" for meshes that is based on a separate processing of the normals [28]. Although these flows produce results that tend to preserve sharp features, they are not a strict generalization of P&M anisotropic diffusion [20] from images to surfaces. Because they are based on reweightings of curvature, these methods always smooth the surface in some direction. Thus, they do not exhibit a sharpening of details, which is achieved by the P&M equation (for images) through an inverse diffusion process.

Chopp and Sethian [29] derive the intrinsic Laplacian of curvature for a 2D implicit curve, and solve the resulting fourth-order PDE. They argue that the numerical methods used to solve fourth-order flows are not practical, because they lack numerical stability. They propose several new numerical schemes, but none are satisfactory due to their slow computation and inability to handle singularities. In related works [30], [31], approximations of higher-order geometric surface flows have been applied to surface fairing in computer graphics.

3 BAYESIAN RECONSTRUCTION

This section provides a brief discussion of the Bayesian formulation for the image/surface reconstruction framework.

3.1 Surface Reconstruction

Let S and M be a surface and the collective set of measured data, respectively. We will assume that M consists of a set of

n registered range images, $\{R^{(j)}\}_{j=1}^n$. For all the examples in this paper, we used the registration methods described in [19], [18]. Then, the posterior probability of \mathcal{S} given the data is

$$P(\mathcal{S} | M) = \frac{P(M | \mathcal{S})P(\mathcal{S})}{P(M)}, \quad (1)$$

where $P(M | \mathcal{S})$ is the likelihood term and $P(\mathcal{S})$ is the prior. Because the goal is to find the surface that maximizes the posterior, the denominator $P(M)$, which is a constant normalization factor that is independent of \mathcal{S} , can be dropped. Typically, MAP estimators are implemented as minimizations of the negative logarithm of the posterior probability

$$\hat{\mathcal{S}} = \arg \inf_{\mathcal{S}} [-\ln P(M | \mathcal{S}) - \ln P(\mathcal{S})], \quad (2)$$

where $\hat{\mathcal{S}}$ is the estimator.

Using the independence of the range images, the log-likelihood can be expressed as a sum $\ln P(M | \mathcal{S}) = -\sum_{j=1}^n \ln P(R^{(j)} | \mathcal{S})$. Whitaker et al. [17], [19], [18] formulate the conditional likelihood of a range image as a volume integral over the object Ω enclosed by \mathcal{S} . In this way, the gradient descent for the log-likelihood term is expressed as a function, $f(\mathbf{x})$, where \mathbf{x} denotes surface locations in the volume. This function provides a weighted sum of the effects of measurements from different scans on the point \mathbf{x} . Details of $f(\mathbf{x})$ are beyond the scope of this paper, for a more in depth discussion, we refer the reader to [17].

Consequently, the gradient descent that minimizes (2) is described by the surface motion

$$\frac{\partial \mathcal{S}}{\partial t} = -f(\mathcal{S})N - \delta P(\mathcal{S}), \quad (3)$$

where δ denotes the Euler-Lagrange operator and N is the surface normal. Notice that the effects of the likelihood and the prior terms on the gradient descent are additive; hence, they can be analyzed separately. The priors we introduce in this paper are not particular to the range data reconstruction problem; they are suitable for use in any Bayesian surface reconstruction problem. For example, the results of this paper can also be applied to tomographical surface reconstruction [32]. Section 4.2 presents results with MRI data as well as range data.

3.2 Image Reconstruction

In this section, we discuss the use of priors in image reconstruction and their relationship to P&M anisotropic diffusion. Mumford and Shah [1], [2] propose an energy minimization approach for image reconstruction/segmentation. The Mumford-Shah energy is the sum of three terms: 1) the quadratic data-model discrepancy, 2) a quadratic penalty on the variation of the piecewise smooth model over the image domain except on a set of discontinuities which are modeled by a separate binary model, and 3) the length of the discontinuities. The first term is the data likelihood term, and the latter two terms act as the prior in this energy. This minimization problem is difficult to solve because of its dependence on binary functions.

Nordstrom [3] established the connection between the Mumford-Shah energy minimization approach to segmentation and the P&M anisotropic diffusion approach to edge detection [20]. This connection is made by observing P&M diffusion from a variational perspective. The P&M diffusion PDE is the gradient descent derived by the first variation of

$$\int_U G(\|\nabla I\|^2) dx dy, \quad (4)$$

where U is the image domain. In its original form, P&M diffusion was introduced to replace Laplacian smoothing, which is equivalent to the solution of the heat equation $\partial I / \partial t = \nabla \cdot \nabla I$, with a nonlinear PDE

$$\partial I / \partial t = \nabla \cdot [g(\|\nabla I\|^2)\nabla I], \quad (5)$$

where I is the gray-level image and g , the derivative of G with respect to $\|\nabla I\|^2$, is the edge stopping function. Perona & Malik suggest using $g = e^{-\|\nabla I\|^2 / 2\mu^2}$, where μ is a positive, free parameter that controls the level of contrast of edges that can affect the smoothing process. Notice that $g(\|\nabla I\|^2)$ approaches 1 for $\|\nabla I\| \ll \mu$ and 0 for $\|\nabla I\| \gg \mu$. Edges are generally associated with large image gradients and, thus, diffusion across edges is prevented while relatively flat regions undergo smoothing. A mathematical analysis shows that solutions to (5) can actually exhibit an inverse diffusion near edges and enhance or sharpen smooth edges that have gradients greater than μ [33]. These properties make P&M diffusion a good candidate for edge detection.

Nordstrom's *biased anisotropic diffusion* converts the variational form of anisotropic diffusion to a reconstruction energy by adding a data term to (4). This yields the variational energy

$$\int_U [(I_o - I)^2 + G(\|\nabla I\|^2)] dx dy, \quad (6)$$

where I_o is the input image. The nonlinear anisotropic penalty on the variation of the model, $G(\|\nabla I\|^2)$, acts as the prior. This penalty term is practically equivalent to the sum of the quadratic model variation and the linear length of discontinuities terms in the Mumford-Shah formulation. Hence, the addition of the data-model discrepancy term transforms P&M diffusion into a variational framework without the need for an explicit binary image for modeling the discontinuities. This is an important practical advantage over minimizing the Mumford-Shah energy.

In this paper, we use energies of the same form as Nordstrom's biased anisotropic diffusion. We propose that the correct generalization of P&M anisotropic diffusion and related image processing methods to surfaces entails the use of higher-order nonlinear priors. We also demonstrate that penalty functions on the surface normals can be used to generate this family of higher-order priors.

4 GEOMETRIC SURFACE PRIORS

Using the Gibbs distribution, probability distributions for priors are commonly expressed in terms of energy functions. Let $E(\mathcal{S})$ be a function that maps surfaces to scalar energy values. Then, the prior can be constructed as

$$P(\mathcal{S}) = \frac{1}{Z} e^{-\alpha E(\mathcal{S})}, \quad (7)$$

where α is a parameter that controls the sharpness of the distribution and Z is a normalizing constant which ensures that the probability distribution function integrates to unity [34]. According to (7), surfaces with lower energy states have a higher probability of occurrence. Geiger and Yuille [34] observe that the parameter α reflects the strength of the

prior. All surfaces are equally likely when $\alpha = 0$; whereas, for $\alpha \rightarrow \infty$, only those surfaces with the lowest possible energy have nonzero probabilities of occurrence. Substituting the Gibbs prior (7) for the prior term in (3) gives

$$-\delta \ln P(\mathcal{S}) = -\delta \ln \frac{1}{Z} e^{-\alpha E(\mathcal{S})} = \alpha \delta E(\mathcal{S}), \quad (8)$$

where α now appears as a relative weight on the surface energy term in (3) with respect to the log-likelihood term. Because the gradient descent of the negative logarithm of the prior is the Euler-Lagrange of the energy function $E(\mathcal{S})$, the rest of this paper is concerned with constructing and solving higher-order feature preserving energy functions that model real surfaces more accurately than previous energy functions investigated in the literature.

Surface area is a commonly utilized surface prior energy based on the underlying assumption that, among surfaces that represent a data set equally well, those with smaller area are relatively simpler and, therefore, have a higher probability of occurrence. A gradient descent minimization of surface area gives the mean curvature flow PDE [33]

$$\delta E(\mathcal{S}) = \text{HN} = \left(\frac{\kappa_1 + \kappa_2}{2} \right) N, \quad (9)$$

where κ_1, κ_2 are the principal curvatures, H is the mean curvature of the surface \mathcal{S} , and N is the surface normal. Despite the simplicity of its solution, surface area is not a realistic shape prior. Difficulties associated with this approach will be demonstrated in Section 6. This implies that a successful model of the smoothness constraints on realistic surfaces requires a general, higher-order and nonlinear prior.

A second-order energy function is the integral of *total curvature*, the sum of the squares of the principal curvatures,

$$\int_{\mathcal{S}} \kappa_1^2 + \kappa_2^2 dS \quad (10)$$

which has been shown to deform surfaces into spheres when minimized [35]. For surfaces with a fixed topology, total curvature is equivalent to the square of the mean curvature by the Gauss-Bonnet theorem. A more general energy function is

$$\int_{\mathcal{S}} G(\kappa_1^2 + \kappa_2^2) dS, \quad (11)$$

where G is a general nonlinear function of total curvature. Minimizing this general energy together with the data likelihood term requires solving fourth-order surface PDEs.

5 A SPLITTING STRATEGY FOR HIGHER-ORDER PRIORS

Several methods for solving fourth-order PDEs have been proposed in the literature. Chopp and Sethian [29] derive the intrinsic Laplacian of curvature for a 2D implicit curve and solve the resulting fourth-order PDE. They investigate several numerical schemes, but find none to be satisfactory in terms of stability. This exact solution is also computationally expensive.

A two-step approximate solution to the intrinsic Laplacian of mean curvature flow for meshes is proposed in [31]. However, that approach can only be applied to meshes and relies on analytic properties of the steady-state solutions for that specific surface flow, $\Delta H = 0$, by fitting

surface primitives that have such properties. Thus, the formalism does not generalize to variational formulations where the solution need not satisfy $\Delta H = 0$.

Another splitting strategy can be found in [37], where the authors penalize the smoothness of a vector field while simultaneously forcing the gradient directions of a gray-scale image to closely match the vector field. The penalty function on the normal field is proportional to the divergence of the normal vectors. This forms a high-order interpolation function, which is shown to be useful for image inpainting—recovering missing patches of data in 2D images. The strategy of simultaneously penalizing the divergence of a normal field, and the mismatch of this field with the image gradient is an approximation to the original fourth-order system. Similarly, our approach also splits the fourth-order system using normal vectors. However, we emphasize the processing of normals on an arbitrary surface manifold (rather than the flat geometry of an image) and use a pair of cascaded PDEs instead of solving them simultaneously. Our cascading approach is closely related to the full fourth-order system, but does not result in an exact solution. A detailed discussion of this relationship can be found in Appendix A.

Our proposed two-step solution allows the surface shape to lag the normals as they are filtered and then refitted by a separate process. This is found to be computationally more efficient than simultaneously solving the two equations and coupling them with a penalty term as in Ballester et al. [37]. Fig. 1 shows this three step process graphically in 2D—shapes give rise to normal maps, which, when filtered, give rise to new shapes.

5.1 Notation

To facilitate the discussion in this section, we employ the Einstein notation convention, where subscripts indicate tensor indexes, and repeated subscripts within a product represent a summation over the index (across the dimensions of the underlying space). Additionally, we use the convention that subscripts on quantities represent derivatives, except where they are in parenthesis, in which case they refer to a vector-valued variable. Thus, ϕ_i is the gradient vector of a scalar quantity $\phi: \mathbb{R}^n \mapsto \mathbb{R}$. The Hessian is ϕ_{ij} , and the Laplacian is ϕ_{ii} . A vector field is $v_{(i)}$, where $v: \mathbb{R}^n \mapsto \mathbb{R}^n$, and the divergence of that field is $v_{(i)i}$. The magnitude of the gradient of ϕ is $\|\phi_i\| = \sqrt{\phi_i \phi_i}$.

Level set surface models rely on the notion of a regular surface, a collection of 3D points with a topology that allows each point to be modeled *locally* as a function of two variables. We can describe the deformation of such a surface using the 3D velocity of each of its constituent points, i.e., $\partial s_{(i)}(t)/\partial t$ for all $s_{(i)} \in \mathcal{S}$. We represent the deformable surface implicitly as

$$\mathcal{S} = \{s_{(i)}(t) \mid \phi(s_{(i)}(t), t) = 0\}, \quad (12)$$

where t is a time parameter. Surfaces defined in this way divide a volume into two parts: inside ($\phi > 0$) and outside ($\phi < 0$). It is common to choose ϕ to be the signed distance transform of \mathcal{S} , or an approximation thereof.

The surface remains a level set of ϕ over time and, thus, taking the total derivative with respect to time gives

$$\frac{\partial \phi}{\partial t} = -\phi_j \frac{\partial s_{(j)}}{\partial t}. \quad (13)$$




Fig. 1. Shown here in 2D, the surface flow process begins with a shape and constructs a normal map from the distance transform (left), modifies the normal map according to a PDE derived from a penalty function (center), and refits the shape to the normal map (right).

Because the gradient ϕ_j is proportional to the surface normal, $\partial s_{(i)}/\partial t$ affects ϕ only in the direction of the surface normal—motion in any other direction results merely in a reparameterization. To compute the appropriate speed term $\partial\phi/\partial t$, we rewrite the energy defined in (11) in a level set framework

$$\mathcal{G}(\phi) = \int_U G(\kappa_1^2 + \kappa_2^2) \|\phi_i\| dx, \quad (14)$$

where $U \subset \mathbb{R}^3$ is the domain of ϕ . The $\|\phi_i\|$ multiplicative factor provides a per-level-set weighting, and minimization of this energy defines $\partial\phi/\partial t$. A discussion of $\partial\phi/\partial t$ follows.

5.2 Curvature from Normal Maps

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. Let $N_{(i)}$ be the normal map, which is a field of normals that are everywhere perpendicular to the family of embedded isosurfaces of ϕ —thus, $N_{(i)} = \phi_i / \|\phi_k\| = \phi_i / \sqrt{\phi_k \phi_k}$.¹ The 3×3 projection matrix for the implicit surface normal is $P_{(ij)} = \phi_i \phi_j / \phi_k \phi_k$, and $P_{(ij)} V_{(i)}$ returns the projection of $V_{(i)}$ onto $N_{(i)}$. Let $I_{(ij)}$ be the identity matrix. Then, the projection onto the plane that is perpendicular to the vector field $N_{(i)}$ is the *tangent projection operator*, $T_{(ij)} = I_{(ij)} - P_{(ij)}$.

The shape matrix [38] of a surface describes its curvature independent of the parameterization. For an implicit surface, it is obtained by differentiating the normal map and projecting the derivative, $N_{(ij)}$, onto the surface tangent plane. The Euclidean norm of the shape matrix is the sum of squared principal curvatures

$$\kappa^2 = \left\| N_{(ij)} T_{(jkl)} \right\|^2. \quad (15)$$

We now express the level set energy function defined in (14) in terms of the normals of the level set surface as a volumetric integral

$$\mathcal{G}(N_{(i)}) = \int_U G \left(\left\| N_{(ij)} T_{(jkl)} \right\|^2 \right) \|\phi_i\| dx. \quad (16)$$

The first variation of this energy with respect to the normals is a second order PDE. It is crucial to observe that, even

1. We use different indices, i and k , for the gradient and the magnitude of the gradient, respectively, to make it clear that the division is performed after the summation in the denominator.

though the projection operator $T_{(jkl)}$ is a function of ϕ , it is independent of the normals because ϕ is fixed as we process the normals. Hence, $T_{(jkl)}$ does not increase the order of the first variation of (16). In contrast, taking the first variation of (14) with respect to ϕ directly would yield a fourth order PDE on ϕ , resulting in a much harder system to solve.

As we process the normal map to minimize (16), allowing ϕ to lag, we must ensure that the normal vectors maintain the unit length constraint. Solutions to constrained optimization problems defined on nonflat manifolds are discussed in [39], [40]. Using the method of Lagrange multipliers, the first variation of the constrained energy becomes

$$\frac{d\mathcal{G}}{dN_{(ij)}} = 2 \left(I_{(ij)} - N_{(i)} N_{(j)} \right) \left[g(\kappa^2) \left(N_{(ijm)} T_{(mkl)} \right) \right]_k, \quad (17)$$

where g is the derivative of G with respect to κ^2 , and κ^2 is as defined in (15). The subscript k outside the square brackets denotes the divergence operation. The unit length constraint for the normal vectors introduces the projection operator $I_{(ij)} - N_{(i)} N_{(j)}$ to (17), i.e., the changes are perpendicular to the unit normal vectors. In a numerical implementation with finite time steps, a separate, explicit normalization step is also needed. A gradient descent of this metric $\partial N_{(ij)}/\partial t = -d\mathcal{G}/dN_{(ij)}$ results in a PDE that minimizes (16). We will discuss several choices for G in Section 6.

5.3 Surface Refitting

We have shown how to evolve the normals to minimize functions of total curvature; however, the final goal is to process the surface, which requires deforming ϕ . Therefore, the next step is to relate the deformation of the level sets of ϕ to the evolution of $N_{(i)}$. Suppose that we are given the normal map $N_{(i)}$ to some set of surfaces, but not necessarily level sets of ϕ —as is the case if we filter $N_{(i)}$ and let ϕ lag. We can manipulate ϕ to fit the normal field $N_{(i)}$ by minimizing a penalty function that quantifies the discrepancy between the gradient vectors of ϕ and the target normal map. This penalty function is

$$\mathcal{D}(\phi) = \int_U \left[\sqrt{\phi_i \phi_i} - \phi_i N_{(i)} \right] dx. \quad (18)$$

The first variation of this penalty function with respect to ϕ is

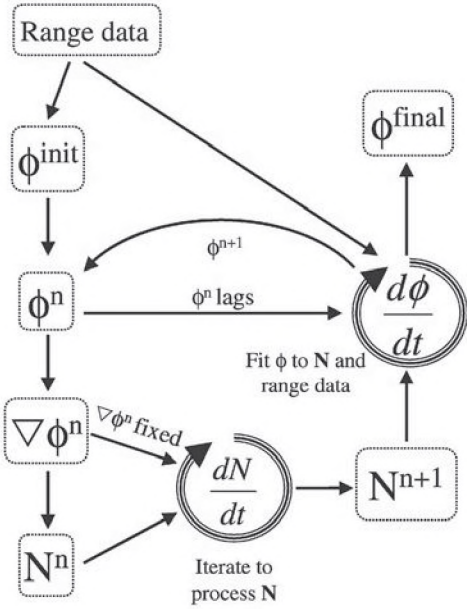


Fig. 2. Flow chart.

$$\frac{dD}{d\phi} = - \left[\frac{\phi_i}{\sqrt{\phi_k \phi_k}} - N_{(i)} \right]_i = - (H^\phi - H^N), \quad (19)$$

where H^ϕ is the mean curvature of the level set surface and H^N is half the divergence of the normal map. The gradient descent PDE that minimizes (18) is $d\phi/dt = -\|\phi_k\|dD/d\phi$. The factor of $\|\phi_k\|$, which is typical with level set formulations, comes from manipulation of the shape of the level set, which is embedded in ϕ . The surface moves according to the difference between its own curvature and that of the normal field.

Let

$$R_{(i)} = \frac{\phi_i}{\sqrt{\phi_k \phi_k}} - N_{(i)}$$


denote the residuals at steady state. Because we must have $\frac{d\phi}{d\phi} = 0$ at steady state, we observe that $R_{(i)} = 0$. This result states that the residual vectors between the processed normals and the unit gradients of the refitted surface are a divergence free field.

The gradient descent for the MAP surface estimator (3) is a weighted sum of the refitting term, derived in (19), and a data term. Therefore, the final update rule for ϕ is


$$\frac{\partial \phi}{\partial t} = -f(\mathcal{S})N + \alpha \|\nabla \phi\| [H^\phi - H^N], \quad (20)$$

where α is a free parameter that determines the relative weight of the prior term with respect to the data term. We will refer to α as the prior weight.

The normal processing stage and the surface refitting stage can now be combined. The entire process is depicted in Fig. 2. We iterate dN/dt for a fixed number, 20 in our implementation, so that the difference between ϕ and the surface defined by the normals remain small. The surface fitting to the combined normal map and data terms is formulated as a gradient descent in (20). This process comprises the $d\phi/dt$ loop in Fig. 2. The overall algorithm shown in Fig. 2 repeats these two steps to minimize the penalty functions in terms of the surface. This main loop iterates until the root mean square (rms) change in ϕ as a



(a)



(b)

Fig. 3. (a) A heavy tailed Gaussian probability distribution function. This function can be scaled to integrate to 1 because, in our discrete implementation, κ has a finite upper bound. (b) Graph of G versus κ with $\mu = 0.2$. G is also the logarithm of (a).

result of the application of $d\phi/dt$ becomes small, less than 10^{-6} in our implementation, which signals convergence. The overall system behavior resembles the fourth-order PDE that minimizes surface curvature. This mathematical relationship is explored in further detail in Appendix A.

6 EXPERIMENTS

In Section 5, we developed a framework allowing the use of priors that are general functions of the variations of the surface normals. We now explore two such functions and compare them to the surface area prior. Total curvature of a surface as defined in (10) can be obtained in our formulation by choosing G to be the identity function, i.e., $G(\kappa^2) = \kappa^2$. The derivative of G with respect to κ^2 is $g(\kappa^2) = 1$. We refer to this choice of G as the *isotropic prior* because it corresponds to solving the heat equation on the normal map with a constant, scalar conduction coefficient. Minimizing the total curvature of surface normals works well for smoothing surfaces and eliminating noise, but, like Gaussian blurring, it also deforms or removes important features.

The generalization of P&M anisotropic diffusion to surfaces is achieved from variational principles by choosing the appropriate function of the total curvature in (16), such as

$$G(\kappa^2) = 2\mu^2 \left(1 - e^{-\frac{\kappa^2}{2\mu^2}} \right), \text{ and } g(\kappa^2) = e^{-\frac{\kappa^2}{2\mu^2}}. \quad (21)$$

We name this choice of G , graphed in Fig. 3b, the *anisotropic prior*. The minimization of the penalty function in (16) with the




Fig. 4. A noisy range image of a sphere plotted as a depth map.

anisotropic prior gives a vector-valued anisotropic diffusion on the level set surface—a generalization of P&M diffusion to surfaces—that preserves/enhances *creases*. Creases are the generalization of edges in images to surfaces. The preservation of creases is achieved by the flat tails of G , that limit the penalty on high curvature areas. Note that μ is fixed at 0.2 for all the experiments in this paper. In contrast to its behavior 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. However, in surface reconstruction, the units are in curvature as can be observed in (21). The only dependence is the scale of interest, i.e., the curvatures of interest. This makes it possible to choose a μ value that gives consistent results over a range of surfaces that are approximately of the same scale.

The anisotropic prior can also be analyzed from a statistical point of view. Typically, robust estimators use heavy tailed Gaussian probability distributions such as the one illustrated in Fig. 3a. Such probability distribution functions assign nonzero probabilities to outliers, large κ values in our case. The logarithm of the heavy tailed Gaussian has the same form as the anisotropic G we have chosen. In fact, using the Gibbs distribution (7), we observe that the prior probability distribution achieved by $G(\cdot)$ is proportional to $P = e^{-G(\cdot)}$. The parameter μ in (21) controls the κ values for which G and P flatten out.

Section 6.1 presents a discussion of the quantitative differences between the isotropic and anisotropic priors and the surface area prior using synthetic data as ground truth. We also investigate the effects of the prior weight on the results. Section 6.2 presents surface reconstruction examples from real data measured by laser range finders and magnetic resonance imaging devices.

6.1 Experiments with Synthetic Data

The experiments presented in this section use geometric shapes for which we can construct analytical distance transforms. We use the following experiment setup:

1. Build range images from the analytical distance transform using the model for the laser range finder located at several positions.
2. Add independent Gaussian noise to the range images to simulate measurement noise.
3. Reconstruct a surface model from the noisy range images.

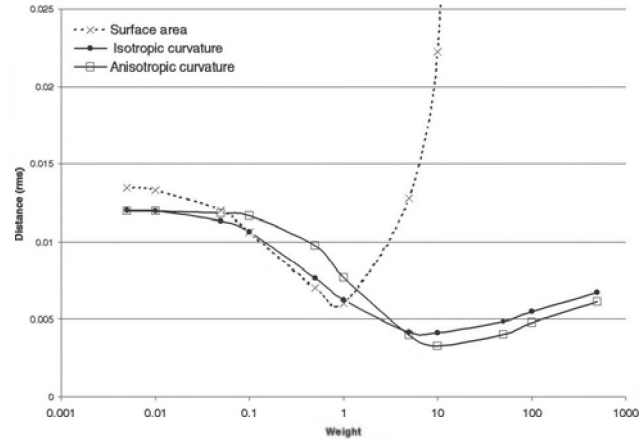


Fig. 5. Rms geometric distance between the analytical sphere and the reconstructed surface.

4. Compare the resulting surface model to the analytical shape by computing the rms geometric distance between the two surfaces.

We first examine 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 that is 10 percent of the sphere's radius (0.1 units), is added to each range image. One of the noisy range images is shown in Fig. 4 as a depth map. We reconstruct surface models from these noisy range images using the three priors under investigation with a range of weights, α . For each choice, we run the algorithm described until it reaches convergence as described in Section 5. We then calculate the rms distance between the original model and the reconstructed model and denote this distance by \mathcal{E} .

Fig. 5 plots \mathcal{E} against the $\log \alpha$ for the different priors under examination. The units on the y -axis are the same as the units used to describe the size of the shape. It can be observed from Fig. 5 that the limit of \mathcal{E} as $\alpha \rightarrow 0$ is approximately 0.0125. This limit is the error obtained if surface reconstruction is performed without a prior (relying only on the data). It 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 priors at their optimal weight provide a 75 percent reduction on this error. On the other hand, surface area provides slightly better than a 50 percent reduction at its optimal weight. These best reconstructions are illustrated in Fig. 6. We contend that the shapes of the error plots are more important than the results at optimal choices of weight. The plot for the surface area prior dips down sharply around $\alpha = 1$ which indicates a small range of useful weights for this prior. The surface area prior performs especially poorly as α is increased beyond 1. This is due to the fact shrinkage in the surface models caused by the surface area prior. In practice, this will cause difficulties for the user in choosing a weight that works for different reconstruction scenarios. In contrast, both of the higher-order priors have relatively flat error plots because their limiting behaviors match the model in this example. Isotropic reconstruction performs as well as the anisotropic reconstruction because the sphere does not contain creases.

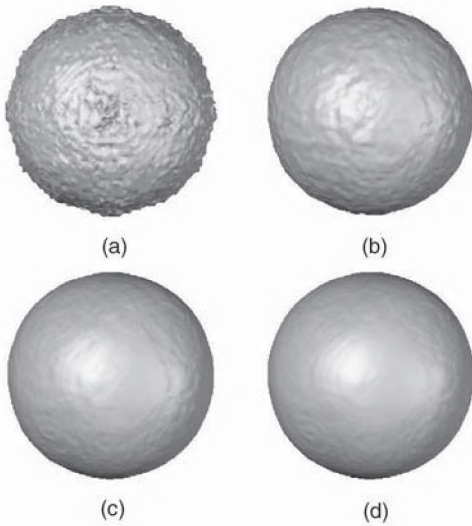


Fig. 6. (a) Surface model initialization from noisy data. Resulting surface model for the (b) surface area prior with $\alpha = 1$, (c) the isotropic prior with $\alpha = 5$, and (d) the anisotropic prior with $\alpha = 10$.






Fig. 8. (a) Surface model initialization from noisy data. Resulting surface model for (b) the surface area prior with $\alpha = 1$, (c) the isotropic prior with $\alpha = 1$, and (d) the anisotropic prior with $\alpha = 10$.



(a)



(b)

Fig. 7. Rms geometric distance between the reconstructed surface and the analytical models for (a) the cube and (b) the “3D-plus-shape.”

To further examine the differences between the priors, we experiment with a cube and another piecewise planar shape which we call the “3D-plus-shape.” In these experiments, we used eight range finder locations (one in each octant). Fig. 8a shows the the surface initialization from the noisy range images of the cube with sides 1 unit long. Independent Gaussian noise with standard deviation 0.1 was added to the simulated data to create the noisy range images. The results (see Fig. 8) with the surface area and isotropic priors are noisier and have rounded corners. In contrast, the results with anisotropic reconstruction are successfully denoised and have sharp creases. Using higher α values for the surface area and isotropic priors to eliminate the noise results in further shape distortion. The last example, shown in Fig. 9a, amplifies the differences observed with the cube experiment. It is not possible to denoise the surface without causing severe distortion to the shape with the other priors; see Figs. 9c, 9d, 9e, and 9f. Anisotropic curvature reconstruction does not suffer from this trade off between smoothing and distortion. It provides a very good (denoised and relatively undistorted) reconstruction across a range of values for the prior weight; see Figs. 9g and 9h. The error plots for the cube and the “3D-plus-shape” shown in Fig. 7 confirm that the anisotropic reconstruction error flattens out as α increases beyond its optimal value.

6.2 Surface Reconstruction from Measured Range Data

In this section, we first investigate these surface priors in the context of the reconstruction of an office scene shown in Fig. 10. This example involves 12 range scans of a room which were registered using the methods described in [18]. One of these 12 range images is shown in Fig. 10a. The model initialization is shown in Fig. 10b. In addition to measurement noise, this initialization contains artifacts such as multiple disconnected pieces seen in the upper right corner. Fig. 11 illustrates the results obtained with the surface area, isotropic and anisotropic priors. As in Section 6.1, we allow the algorithms to run until convergence is reached. The prior term weights were selected as those that produced the




Fig. 9. (a) Analytical “3D-plus-shape” and (b) surface model initialization from noisy data. Resulting surface model for the surface area prior with (c) $\alpha = 1$ and (d) $\alpha = 10$. For the isotropic prior with (e) $\alpha = 1$ and (f) $\alpha = 10$. For the anisotropic prior with (g) $\alpha = 10$ and (h) $\alpha = 100$.

qualitatively best results. The results with the anisotropic prior remain close to the actual surface while eliminating noise. On the other hand, the surface area prior that produces a comparable amount of denoising demonstrates shape distortion such as the breaking of the arms of the chairs and the rounding of the creases on the desk and computer equipment. The isotropic curvature prior does not cause any severe shape distortions, but it smooths edges, creases, and corners as expected. This experiment illustrates the importance of the anisotropic prior in reconstructions involving scenes with high curvature features and sharp creases.

Fig. 12 demonstrates the robustness of the different priors by a detailed examination of a small portion of the scene. Figs. 12a and 12b are the initialization and the visible imagery for one of the chairs in the scene. Figs. 12c, 12e, and 12g illustrate the results obtained by qualitatively choosing successful values for α . Figs. 12d, 12f, and 12h illustrate the results if α is chosen to be 10 times this value. These results show that the anisotropic curvature is least sensitive to the choice of the prior weight. Also, observe that the beam connecting the base to the seat is being pinched-off in Fig. 12d.




Fig. 10. (a) One of the range images used in the surface reconstruction experiment and (b) surface initialization.

Fig. 13 illustrates the reconstruction of a military vehicle with the anisotropic curvature prior. This reconstruction also used 12 range images. Both reconstructions require approximately one hour of computation on a *Intel 1.7 GHz Xeon processor*.

As discussed before, the prior term and the methods presented in this paper are not specific to reconstruction from range data. We now demonstrate results of our method used in conjunction with an algorithm for level set surface segmentation from volumetric MRI data. For this purpose, we adopt the data term introduced in [41], [42]. This data term deforms the surface model in such a way that voxels falling between a low and a high intensity threshold will be contained in the interior of the surface. Fig. 14a shows one slice of the the volumetric MRI data. We initialize the surface as a box that is slightly smaller than the size of the data. The intersection of this initial surface with the slice is shown as a light gray contour.

In the lack of a prior term, we the data term forces the surface to match every detail in the data, creating a noisy segmentation, see Fig. 15. Using the same anisotropic prior as before, we obtain a smoother, less noisy, and feature preserving segmentation. Fig. 14b shows a slice of the resulting surface. This segmentation requires approximately two hours of computation on an *Intel 1.7 GHz Xeon processor*.

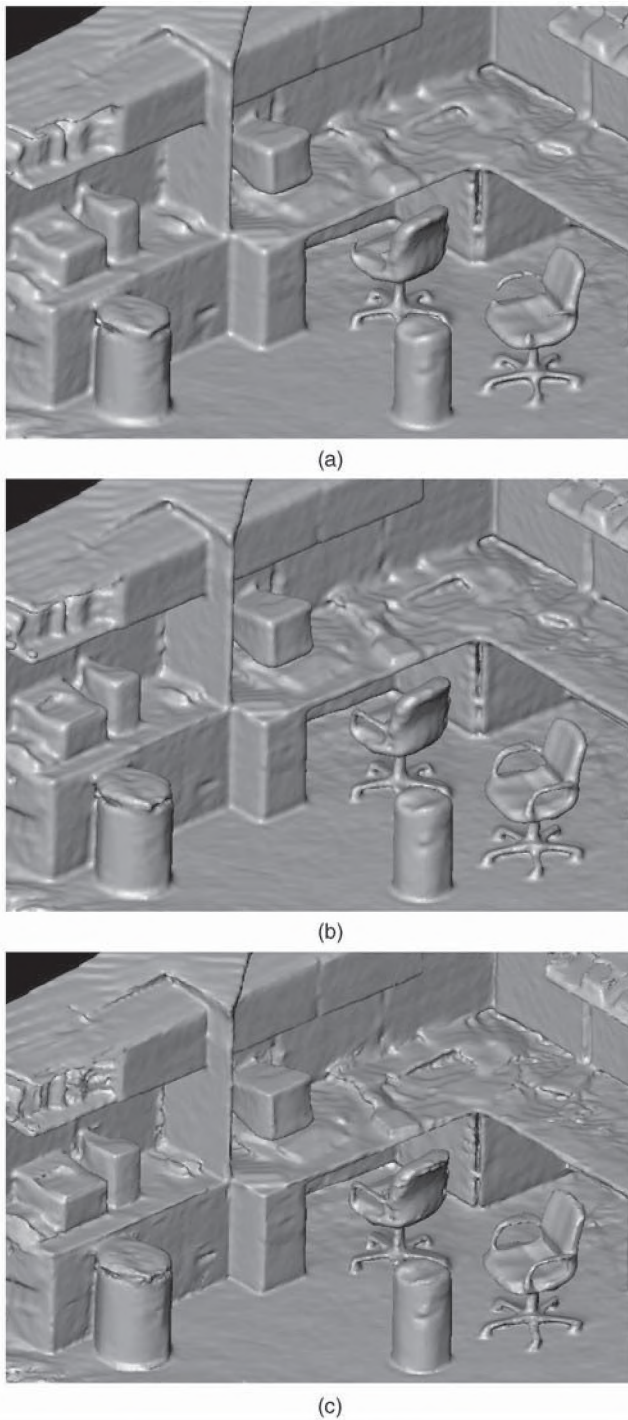


Fig. 11. Results with (a) surface area, (b) isotropic curvature, and (c) anisotropic priors.

7 CONCLUSION

Anisotropic diffusion and image reconstruction techniques based on robust metrics have been shown to be very useful tools in image processing. We generalize these methods to surface reconstruction by considering a general family of penalty functions of curvature. The minimization of these second-order penalty functions require solving fourth-order PDEs on level sets. We avoid an exact solution because this would require very small time steps in a numerical implementation [29]. Instead, we propose a cascaded pair of PDEs that resemble the original fourth-order PDE.




Fig. 12. (a) Initial surface model for chair and (b) a visible light image of the chair taken from a similar point of view. Results for the surface area prior with weights (c) 1 and (d) 10. Results for the isotropic prior with weights (e) 1 and (f) 10. Results for the anisotropic prior with weights (g) 1 and (h) 10.

The splitting of the original fourth-order PDE into a cascade system is achieved via the surface normals. This method is based on the proposition that the *natural* generalization of




Fig. 13. Reconstruction of a military vehicle (with the anisotropic prior) from range data.

image processing to surfaces occurs via the normals. Normals are processed separately from the surface using a metric on the surface manifold, rather than a simple, flat metric. By processing the normals separately from the surface, we can solve a pair of cascaded second-order equations instead of a fourth-order equation. Typically, we allow one equation (the surface) to lag the other. This method is numerically more




Fig. 14. A sagittal slice from a head MRI volume. The highlighted contours represent the intersection of the level set surface with the slice: (a) initialization and (b) segmented model with the anisotropic prior.



(a)



(b)

Fig. 15. Segmentation of the facial surface (a) with a data term only and (b) with the data term and the anisotropic prior.

stable and computationally less expensive than solving the fourth-order PDE directly and resembles the original fourth-order system. However, it is not mathematically equivalent to it. We solve these equations using implicit surfaces, representing the implicit function on a discrete grid and modeling the deformation with the method of level sets. The method applies equally well to surfaces that can be represented in a volume.

We have shown that an anisotropic penalty on curvature computed from the surface normals leads to a surface reconstruction prior that preserves creases while denoising the input. This process demonstrates important quantitative and qualitative advantages over processes that use the surface

area prior. The data term can be chosen independently from the prior; therefore, the ideas introduced in this paper can be applied to other forms of surface reconstruction such as applications in tomography.

The main shortcoming of this method is the computation time, which is significant. However, the current process lends itself to parallelism. The advent of cheap, specialized, stream-processing hardware promises significantly faster implementations when the inherent parallelism in the process is exploited [43], [41], [42]. Multithreading could also be utilized to exploit the parallelism in the process. Regarding the algorithm itself, the use of adaptive or multiresolution level set strategies could also improve the processing time.

APPENDIX A

MATHEMATICAL FOUNDATION

The proposed algorithm for imposing second-order priors on surfaces does not, strictly speaking, produce a fourth-order flow. We have avoided a true fourth-order implementation because such flows impose small time steps on the corresponding numerical algorithms. The proposed method is a second-order flow which has a fundamental relationship to the second-order penalty function described in Section 4. As noted by Ambrosio and Masnou [36], the relationship between such second-order systems (see also [37]) and the corresponding fourth-order variational problem is not fully understood.

The purpose of this appendix is to describe, in a more precise way, the relationship between the proposed second-order system and the corresponding fourth-order flow. The argument is that if the normals are the normals of the corresponding level-sets of ϕ and the normals are processed by an infinitesimal amount, then the refitting process creates a flow on the level-set surfaces which is the same as the corresponding fourth-order flow. This does not imply that the proposed method approximates the fourth-order flow, because we do not constrain the normals and the surface to move together in a coupled manner. However, the discussion below does give the mathematical underpinnings of the proposed method and helps us understand why the results conform, qualitatively, to our expectations of the fourth-order flow.

In this discussion, we use vector notation for clarity, and we denote the energy \mathcal{G}_ϕ . We can rewrite this energy function by observing that the principal curvatures are functions of the derivatives of ϕ

$$\mathcal{G}_\phi = \int_U G(\phi) \|\nabla\phi\| dx. \quad (22)$$

Let $d\phi : \mathbb{R}^3 \rightarrow \mathbb{R}$ be a volume of incremental changes applied to $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}$. The change to \mathcal{G} induced by $d\phi$ can be expressed as the volume integral of the total derivative of the penalty function, $dG(\phi) \|\nabla\phi\|$, which is the product of $d\phi$ and the variation of the penalty function with respect to ϕ

$$d\mathcal{G}_\phi = \int_U \frac{d(G \|\nabla\phi\|)}{d\phi} d\phi dx. \quad (23)$$

Applying the product rule to $\frac{d(G \|\nabla\phi\|)}{d\phi}$, we obtain

$$d\mathcal{G}_\phi = \underbrace{\int_U \frac{dG}{d\phi} \|\nabla\phi\| d\phi dx}_{d\mathcal{G}_{\phi,1}} + \underbrace{\int_U G \frac{d\|\nabla\phi\|}{d\phi} d\phi dx}_{d\mathcal{G}_{\phi,2}}. \quad (24)$$

The total derivative $dG(\phi) \|\nabla\phi\|$ can be written in terms of the surface normals by using the equality

$$\frac{dG}{d\phi} d\phi = \frac{dG}{d\mathbf{N}} \cdot d\mathbf{N}, \quad (25)$$

given that the normal map is a function of ϕ . Then, the first term in (24) can be written as

$$d\mathcal{G}_{\phi,1} = \int_U \frac{dG}{d\mathbf{N}} \cdot d\mathbf{N} \|\nabla\phi\| dx. \quad (26)$$

To simplify (26), we derive $d\mathbf{N}$ as a function of \mathbf{N} and $d\nabla\phi$

$$d\mathbf{N} = (\mathbf{I} - \mathbf{N} \otimes \mathbf{N}) \frac{d\nabla\phi}{\|\nabla\phi\|}. \quad (27)$$

Substituting (27) for $d\mathbf{N}$ in (26) and using the commutivity of differentiation ($d\nabla\phi = \nabla d\phi$), we get

$$d\mathcal{G}_{\phi,1} = \int_U \left((\mathbf{I} - \mathbf{N} \otimes \mathbf{N}) \frac{d\mathbf{G}}{d\mathbf{N}} \right) \cdot \nabla(d\phi) dx. \quad (28)$$

We treat this energy minimization as an *adiabatic* problem, in which energy flow across the boundary of U is zero. Hence, using *Neumann* boundary conditions for U and integrating by parts, we obtain

$$d\mathcal{G}_{\phi,1} = \int_U \nabla \cdot \left((\mathbf{I} - \mathbf{N} \otimes \mathbf{N}) \frac{d\mathbf{G}}{d\mathbf{N}} \right) d\phi dx. \quad (29)$$

We now examine the second term in (24), $d\mathcal{G}_{\phi,2}$. As in Section 5.2, we treat G as a function of \mathbf{N} ; therefore, due to the decoupling between \mathbf{N} and ϕ , G can be considered independent of ϕ . Using this assumption, we can rewrite $d\mathcal{G}_{\phi,2}$ as

$$d\mathcal{G}_{\phi,2} = \int_U \frac{dG^{\mathbf{N}} \|\nabla\phi\|}{d\phi} d\phi dx, \quad (30)$$

where the superscript on $G^{\mathbf{N}}$ indicates that G is fixed with respect to ϕ . Taking the first variation of $dG^{\mathbf{N}} \|\nabla\phi\|$ yields

$$d\mathcal{G}_{\phi,2} = \int_U \nabla \cdot \left(G^{\mathbf{N}} \frac{\nabla\phi}{\|\nabla\phi\|} \right) d\phi dx, \quad (31)$$

since for a vector v , $\frac{d\|v\|}{dv} = \frac{v}{\|v\|}$. Finally, combining (23), (29), and (31), we can derive the desired relationship between the variations with respect to ϕ and \mathbf{N}

$$\frac{d(G \|\nabla\phi\|)}{d\phi} = \nabla \cdot \left((\mathbf{I} - \mathbf{N} \otimes \mathbf{N}) \frac{d\mathbf{G}}{d\mathbf{N}} + G^{\mathbf{N}} \frac{\nabla\phi}{\|\nabla\phi\|} \right). \quad (32)$$

Let us now consider the flow achieved by processing (17) and (19) back to back in one iteration of the main loop in Fig. 2 again. At the beginning of iteration n , the normals are computed from ϕ^n . If we evolve the normals for one step according to (17), instead of processing them multiple iterations, the new normals are

$$\mathbf{N}^{n+1} = \mathbf{N}^n - (\mathbf{I} - \mathbf{N} \otimes \mathbf{N}) \frac{d\mathbf{G}}{d\mathbf{N}}, \quad (33)$$

where we write $\frac{dG}{dN}$ instead of $\frac{dG}{d\mathbf{N}}$ because we are referring to the update for \mathbf{N} at a specific point in space. If we immediately apply (19) to fit ϕ to this new normal map, we get

$$\frac{dD}{d\phi} = H^{\phi^n} - \nabla \cdot \left(\mathbf{N}^n - (\mathbf{I} - \mathbf{N} \otimes \mathbf{N}) \frac{d\mathbf{G}}{d\mathbf{N}} \right), \quad (34)$$

where D is the local function defined in (18). Because \mathbf{N}^n is derived directly from ϕ^n , we have $\nabla \cdot \mathbf{N} = \mathbf{H}^{\phi^n}$, which gives the rule in our algorithm to make up this infinitesimal lag:

$$\frac{dD}{d\phi} = \nabla \cdot (\mathbf{I} - \mathbf{N} \otimes \mathbf{N}) \frac{d\mathbf{G}}{d\mathbf{N}}. \quad (35)$$

Comparing with (32), we find the rule to descend on the energy as a function of ϕ

$$\frac{d(G \|\nabla\phi\|)}{d\phi} = \frac{dD}{d\phi} + \nabla \cdot \left(G^{\mathbf{N}} \frac{\nabla\phi}{\|\nabla\phi\|} \right). \quad (36)$$

Thus, we see that we can mimic the fourth-order flow as a combination of second order terms; one that fits the surface to the normal field and the other that moves the surface between normal maps to find the set of normals that minimize the total curvature. In our experiments, we have found that the contribution of the second term is very small and does not change the results qualitatively. Therefore, we drop it for the sake of computational efficiency and implement only $\frac{dD}{d\phi}$ as described in Section 4. Because we discretize the equations in time in an uncoupled manner and recompute the normals asynchronously at the beginning of each iteration, the resulting numerical scheme is not *strictly* fourth order.

ACKNOWLEDGMENTS

This work was supported, in part, by the following research grants: ONR N00014-01-10033, NSF CCR0092065, and ARO DAAD19-01-1-0013.

REFERENCES

- [1] D. Mumford and J. Shah, "Boundary Detection by Minimizing Functionals," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1985.
- [2] D. Mumford and J. Shah, "Optimal Approximations by Piecewise Smooth Functions and Associated Variational Problems," *Comm. Pure and Applied Math.*, vol. 42, pp. 577-685, 1989.
- [3] N. Nordstrom, "Biased Anisotropic Diffusion—A Unified Regularization and Diffusion Approach to Edge Detection," *Image and Vision Computing*, vol. 8, no. 4, pp. 318-327, 1990.
- [4] J. Shah, "Segmentation by Nonlinear Diffusion," *Proc. Conf. Computer Vision and Pattern Recognition*, pp. 202-207, 1991.
- [5] M.J. Black, G. Sapiro, and D.H. Marimont, "Robust Anisotropic Diffusion," *IEEE Trans. Image Processing*, vol. 7, pp. 421-432, Mar. 1998.
- [6] T. Tasdizen, R. Whitaker, P. Burchard, and S. Osher, "Geometric Surface Smoothing via Anisotropic Diffusion of Normals," *Proc. IEEE Visualization*, pp. 125-132, Oct. 2002.
- [7] R.M. Bolle and D.B. Cooper, "On Optimally Combining Pieces of Information, with Application to Estimating 3D Complex-Object Position from Range Data," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 8, no. 5, pp. 619-638, Sept. 1986.
- [8] R. Bajcsy and F. Solina, "Three Dimensional Object Representation Revisited," *Proc. First Int'l Conf. Computer Vision*, pp. 231-240, June 1987.
- [9] A.P. Pentland, "Recognition by Parts," *Proc. First Int'l Conf. Computer Vision*, pp. 612-620, June 1987.
- [10] G. Taubin, "Estimation of Planar Curves, Surfaces, and Nonplanar Space Curves Defined by Implicit Equations with Applications to Edge And Range Image Segmentation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 11, pp. 1115-1138, Nov. 1991.
- [11] D. DeCarlo and D. Metaxas, "Adaptive Shape Evolution Using Blending," *Proc. Fifth Int'l Conf. Computer Vision*, pp. 834-839, 1995.
- [12] G. Turk and M. Levoy, "Zippered Polygon Meshes from Range Images," *Proc. SIGGRAPH*, pp. 311-318, July 1994.
- [13] Y. Chen and G. Medioni, "Fitting a Surface to 3D Points Using an Inflating Ballon Model," *Proc. Second CAD-Based Vision Workshop*, A. Kak and K. Ikeuchi, eds., vol. 13, pp. 266-273, 1994.
- [14] B. Curless and M. Levoy, "A Volumetric Method for Building Complex Models from Range Images," *Proc. SIGGRAPH (Computer Graphics)*, July 1996.
- [15] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface Reconstruction from Unorganized Points," *Computer Graphics*, vol. 26, no. 2, pp. 71-78, 1992.
- [16] A. Hilton, A.J. Stoddart, J. Illingworth, and T. Windeatt, "Reliable Surface Reconstruction from Multiple Range Images," *Proc. European Conf. Computer Vision*, 1996.
- [17] R.T. Whitaker, "A Level-Set Approach to 3D Reconstruction from Range Data," *Int'l J. Computer Vision*, vol. 29, no. 3, pp. 203-231, 1998.
- [18] R. Whitaker and J. Gregor, "A Maximum Likelihood Surface Estimator for Dense Range Data," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, Oct. 2002.
- [19] J. Gregor and R.T. Whitaker, "Indoor Scene Reconstruction from Sets of Noisy Range Images," *Graphical Models*, vol. 63, pp. 304-332, 2002.
- [20] P. Perona and J. Malik, "Scale Space and Edge Detection Using Anisotropic Diffusion," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629-639, July 1990.
- [21] D. Geiger and A. Yuille, "A Common Framework for Image Segmentation," *Int'l J. Computer Vision*, vol. 6, no. 3, pp. 227-243, 1991.
- [22] W. Snyder, Y.-S. Han, G. Bilbro, R. Whitaker, and S. Pizer, "Image Relaxation: Restoration and Feature Extraction," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, no. 6, pp. 620-624, June 1995.
- [23] R.T. Whitaker, "Volumetric Deformable Models: Active Blobs," *Visualization in Biomedical Computing*, R.A. Robb, ed., SPIE, 1994.
- [24] I. 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 Structures," *Proc. Computer Vision and Pattern Recognition*, 2000.
- [25] G. Taubin, "A Signal Processing Approach to Fair Surface Design," *Proc. SIGGRAPH*, pp. 351-358, 1995.
- [26] U. Clarenz, U. Diewald, and M. Rumpf, "Anisotropic Geometric Diffusion in Surface Processing," *Proc. IEEE Visualization*, pp. 397-405, 2000.
- [27] Y. Ohtake, A.G. Belyaev, and I.A. Bogaevski, "Polyhedral Surface Smoothing with Simultaneous Mesh Regularization," *Geometric Modeling and Processing*, 2000.
- [28] G. Taubin, "Linear Anisotropic Mesh Filtering," Technical Report RC22213, IBM Research Division, Oct. 2001.
- [29] D.L. Chopp and J.A. Sethian, "Motion by Intrinsic Laplacian of Curvature," *Interfaces and Free Boundaries*, vol. 1, pp. 1-18, 1999.
- [30] W. Welch and A. Witkin, "Free-Form Shape Design Using Triangulated Surfaces," *Proc. SIGGRAPH '94*, pp. 247-256, 1994.
- [31] R. Schneider and L. Kobbelt, "Generating Fair Meshes with g^1 Boundary Conditions," *Proc. Geometric Modeling and Processing*, pp. 251-261, 2000.
- [32] R. Whitaker and V. Elangovan, "A Direct Approach to Estimating Surfaces in Tomographic Data," *J. Medical Image Analysis*, vol. 6, no. 3, pp. 235-249, 2002.
- [33] G. Sapiro, *Geometric Partial Differential Equations and Image Analysis*. Cambridge Univ. Press, 2001.
- [34] D. Geiger and A. Yuille, "A Common Framework for Image Segmentation," *Int'l J. Computer Vision*, vol. 6, no. 3, pp. 227-243, 1991.
- [35] A. Polden, "Compact Surfaces of Least Total Curvature," technical report, Univ. of Tübingen, Germany, 1997.
- [36] L. Ambrosio and S. Masnou, "A Direct Variational Approach to a Problem Arising in Image Reconstruction," *Interfaces and Free Boundaries*, vol. 5, no. 1, pp. 63-81, Jan. 2003.

- [37] C. Ballester, M. Bertalmio, V. Caselles, G. Sapiro, and J. Verdera, "Filling-In by Joint Interpolation of Vector Fields and Gray Levels," *IEEE Trans. Image Processing*, vol. 10, pp. 1200-1211, Aug. 2001.
- [38] M.P. DoCarmo, *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.
- [39] M. Bertalmio, L.-T. Cheng, S. Osher, and G. Sapiro, "Variational Methods and Partial Differential Equations on Implicit Surfaces," *J. Computational Physics*, vol. 174, pp. 759-80, 2001.
- [40] B. Tang, G. Sapiro, and V. Caselles, "Diffusion of General Data on Non-Flat Manifolds via Harmonic Maps Theory: The Direction Diffusion Case," *Int'l J. Computer Vision*, vol. 36, no. 2, pp. 149-161, 2000.
- [41] A. Lefohn, J.M. Kniss, C. Hansen, and R. Whitaker, "Interactive Deformation and Visualization of Level Set Surfaces Using Graphics Hardware," *Proc. IEEE Visualization*, pp. 75-82, 2003.
- [42] A. Lefohn, J. Cates, and R. Whitaker, "Interactive, GPU-Based Level Sets for 3D Brain Tumor Segmentation," *Medical Image Computing and Computer Assisted Intervention*, pp. 564-572, 2003.
- [43] M. Rumpf and R. Strzodka, "Level Set Segmentation in Graphics Hardware," *Proc. Int'l Conf. Image Processing*, pp. 1103-1106, 2001.



Tolga Tasdizen received the BS degree in electrical engineering from Bogazici University, Istanbul in 1995. He received the MS and PhD degrees in engineering from Brown University in 1997 and 2001, respectively. Since 2001, he has been at the University of Utah where he is a postdoctoral research associate with the School of Computing. He is also a member of the Scientific Computing and Imaging Institute at the University of Utah. He is a member of the IEEE and the IEEE Computer Society.



Ross Whitaker received the BS degree in electrical engineering and computer science from Princeton University in 1986, earning Summa Cum Laude. From 1986 to 1988, he worked for the Boston Consulting Group, entering the University of North Carolina (UNC) at Chapel Hill in 1989. At UNC, he received the Alumni Scholarship Award, and completed the PhD in computer science in 1994. From 1994 to 1996, he worked at the European Computer-Industry Research Centre in Munich Germany as a research scientist in the User Interaction and Visualization Group. From 1996 to 2000, he was an assistant professor in the Department of Electrical Engineering at the University of Tennessee. Since 2000, he has been at the University of Utah where he is an associate professor in the College of Computing and a faculty member of the Scientific Computing and Imaging Institute. He is a member of the IEEE.

► **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**