Eurographics/ IEEE-VGTC Symposium on Visualization (2006)
Thomas Ertl, Ken Joy, and Beatriz Santos (Editors)

3D Soft Segmentation and Visualization of Medical Data Based on Nonlinear Diffusion and Distance Functions

B. Petersch, O. Serrano-Serrano and D. Hönigmann

Advanced Computer Vision GmbH, Vienna, Austria

Abstract

Visualization of medical 3D data is a complex problem, since the raw data is often unsuitable for standard techniques like Direct Volume Rendering. Some kind of pre-treatment is necessary, usually segmentation of the structures of interest, which in turn is a difficult task. Most segmentation techniques yield a model without indicating any uncertainty. Visualization then can be misleading, especially if the original data is of poor contrast.
We address this dilemma proposing a geometric approach based on distance on image manifolds and an alternative approach based on nonlinear diffusion. An effective algorithm solving Hamilton-Jacobi equations allows for computing a distance function for 2D and 3D manifolds at interactive rates. An efficient implementation of a semi-implicit operator splitting scheme accomplishes interactivity for the diffusion-based strategy. We establish a model which incorporates local information about its reliability and can be visualized with standard techniques. When interpreting the result of the segmentation in a diagnostic setting, this information is of utmost importance.

Categories and Subject Descriptors (according to ACM CCS): I.4.6 [IMAGE PROCESSING AND COMPUTER VISION]: Segmentation I.3.7 [COMPUTER GRAPHICS]: Three-Dimensional Graphics and Realism - Volume Rendering I.3.8 [COMPUTER GRAPHICS]: Applications

Keywords: uncertainty, visualization, segmentation, distance function, manifold, nonlinear diffusion

1. Introduction

With the advent of sophisticated 3D imaging devices for medical diagnosis and intervention planning, proper visualization of these data sets has become a major challenge. What is proper and not in this context is strongly dependent on the medical question at hand, but also on the way medical experts are used and trained to see and interpret the data. On the one hand 3D data processing and rendering techniques can be a valuable supplement to standard 2D imaging, as e.g. in breast cancer diagnosis [MSKB*04]. On the other hand, the possibility to render anatomical structures in 3D has opened the door to completely new medical areas like image-guided surgery [FSS*02]. Renderings of 3D data are frequently based on segmentation results. Segmentation deals with isolating objects from their background within digital data and usually yields "decided" models. Often such decided models do not represent the object adequately for several reasons: the structure under examination might not be sharply defined either due to poor data quality linked directly to acquisition problems or shortcomings of the device used for data acquisition, or due to the fact that the structure's perimeter itself is indistinct. Tumors in medical data sets are typical examples: In radiotherapy planning, defining the tumor's border and knowing the reliability of the border delineation is crucial for treatment planning and quality assurance [vHRRL00, MvHM02]. In mammography and breast ultrasound, lesion differentiation depends on correct assessment not only of the lesion's shape, but also of its boundary acutance [RFDA97, STR95, HRSB05].
Recently there is an increasing interest in the *visualization* of errors and uncertainty [PWL97, CR00, DKLP02, JS03, KLDP04, GR04, LP05]. While the scientific visualization community is well aware of the necessity to visualize errors, this still seems to be neglected in medical data visualization too often. We therefore want to demonstrate one possible way to achieve integration of segmentation and error visualization of anatomical structures from real world medical

data.

We present a two-step segmentation strategy. The first step is to compute an auxiliary n-dimensional function ϕ with low values inside the structures and high values outside. Two different possibilities to compute ϕ are proposed which can be used according to the data characteristics of the structures to be visualized. The second step then performs the actual segmentation on ϕ (see Fig.1).

The first approach settles in a geometric framework and will henceforth be called the *geometric approach*. It considers *n*-dimensional data, each point in *n*-space contributing *m* features (e.g., intensities, channels of an arbitrary color space, textural features, or the like). The basic concept of this approach is to treat image data as manifolds embedded in a higher dimensional space [Soc98, KMS00]. Distances are then calculated on this manifold with respect to (a) user specified input point(s) within the structure(s) as a measure of similarity. Distances will then be low inside the structure(s) and rise rapidly at locations where the data features change, i.e. the structure's border. This strategy is aimed at segmentation of a single (or few) rather compact structure(s) like tumors or cysts. It is very flexible since any number of features characterizing the structure(s) can be embedded to form a manifold for the calculation of the distance measure, which can be of advantage if signal strength alone is insufficient to characterize the structure of interest (for details on this method see [HPRP05]).

In order to overcome the drawback of having to choose a distance reference point and to further speed up the method, we developed a second approach which uses input data filtered by nonlinear diffusion [Wei98] as a basis for segmentation and is henceforth called the *diffusion approach*. This strategy is aimed at segmentation of a single structure or a multitude of structures which show(s) rather homogeneous signal σ. By remapping the intensities within the original data such that $\sigma \mapsto 0$, an auxiliary function ϕ can be obtained, which is similar to the distance function of the geometric approach. A diffusion filter then serves the purpose of edge-preserving smoothing to deal with small inhomogeneities of σ within the structure(s). In ultrasound (US) image data for example, blood vessels of different kind (liver vessels, brain ventricles, heart chambers,...), follicles and various cysts and tumors can all appear hypoechoic (low signal embedded in surrounding tissue of higher signal), in which case not even remapping of signals has to be performed but the diffusion approach can be applied directly.

In both scenarios, the segmentation result shall provide, for every data point, information about the location of the object(s) as well as an indication of the reliability of the segmentation. Finally, we aim for a solution which allows to segment 3D data sets of moderate size at interactive rates. In the present context, we consider visualization the ultimate goal of the segmentation task. I.e., our work aims at providing an integrated framework for "soft" segmentation and visualization of the uncertainty. Despite the capability of both approaches to operate on *n*-dimensional data, we content ourselves to standard 3D visualization techniques such as direct volume rendering. To illustrate the concept as well as the results one can achieve, we will mainly use 3D ultrasound data sets, this modality being particularly prone to mediocre image quality through various imaging artifacts [SSG95] while requiring very fast segmentation techniques due to interactive data acquisition.

This paper is organized as follows. In section 2 we briefly summarize relevant work related to the subject. In section 3 we discuss the calculation of distance on manifolds (3.1) and the semi-implicit operator splitting scheme for nonlinear diffusion (3.2), while section 3.3 outlines the actual soft segmentation technique. Finally we discuss results in section 4 and conclude our work with section 5.

2. Related work

Nonlinear Diffusion. Nonlinear or anisotropic diffusion filtering [PM90, CLMC92, Gil02] can be used for effective removal of different kinds of noise. The main disadvantage of such filters is usually their slow calculation speed, which is especially crucial in 3D data processing. For our diffusion approach we therefore use the fast operator splitting scheme proposed in [Wei98]. Data denoising problems can also be formulated within a geometric framework.

The Beltrami Framework. In [Soc98] and [KMS00], Sochen, Kimmel, and Malladi derive a formulation unifying some classical flow-based image denoising algorithms using concepts of high energy physics. They coin the term *Beltrami framework*. The underlying idea is to treat image processing as an evolution of an image manifold embedded in a higher dimensional Riemannian space towards a minimal surface. Their framework is capable of edge-preserving denoising of vector-valued images of arbitrary dimension. Among the results emanating from this work are, e.g., orientation diffusion filtering [KS02] and Beltrami flow on implicit surfaces [SDP03].

Geometric Segmentation. In addition to geometrically motivated filtering techniques, there is a number of *segmentation* methods inspired by geometry. These include level set propagation [OF02, Set99], geodesic active contours [CKS97, Kim04], and others [Kim04, Sap01]. Our first approach settles in a geometric segmentation framework and employs level sets of distance functions on manifolds.

Distance on manifolds. Distance functions on manifolds are a central concept in our first approach. For the computation of such functions, two basic methods have been proposed: fast marching [MS01, SK04] and fast sweeping [TCOZ03]. Applications concern for example morphology on surfaces [PSH*04]. Geodesics already appear in geodesic active contours, but possess a number of remarkable further applications in image processing and computer vision [Kim04]. Instead of geodesics, we use appropriately modified level sets

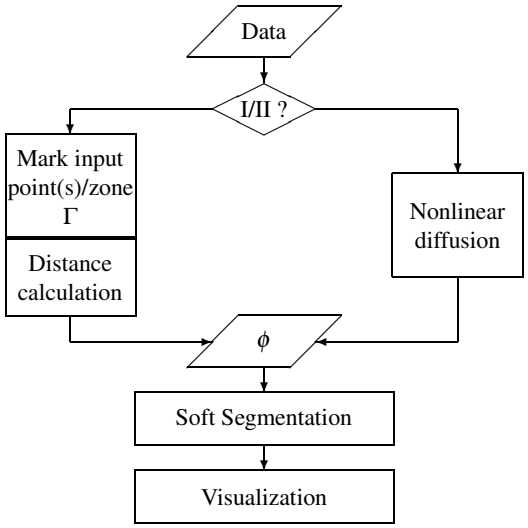

Figure 1: *Flowchart of the complete process using the geometric (I) or the diffusion approach (II).*

of distance functions on image manifolds for soft segmentation.

Fuzzy segmentation and statistical approaches. A popular approach for linking segmentation and uncertainty visualization dates back to the work of Udupa [US96] who introduced the concept of *fuzzy connectedness* for image segmentation. Fuzzy segmentation approaches usually suffer from drawbacks such as the result strongly depending on (i) the choice of the functions defining the pixel affinity, (ii) the selection of an appropriate threshold for the fuzzy affinity map. Implementations of fuzzy segmentation algorithms are often based on dynamic programming, resulting in moderate speed. It exists a lot of work dealing with probabilistic segmentation (e.g. [PT01, KvU05]), being quite different from our approach which does not require any statistical prior model of data features and noise.

3. Materials and methods

In order to segment anatomical structures from 3D medical data we are first aiming at generating an auxiliary 3D function ϕ which contains low values inside the object(s) to be segmented and visualized and high values outside. The actual soft segmentation procedure (see section 3.3) is then performed on ϕ rather than on the original data. We present two different approaches to generate ϕ, which are aimed at solving different segmentation scenarios. Figure 1 shows a flowchart of the overall process.

3.1. Geometric Approach

Metric and distance functions. In the geometric approach, we regard n-dimensional image data comprising m-dimensional feature vectors as an embedding $X : \Sigma \to M$ of an n-dimensional manifold Σ with coordinates $\sigma^1, \sigma^2, \cdots, \sigma^n$ in an $(n+m)$-dimensional hybrid space M of mixed spatial coordinates and feature coordinates $X^1, X^2, \cdots, X^{n+m}$.

Here, we assume that the embedding space is Riemannian, i.e., in M there is locally an inner product defined by the metric $\mathbf{H}$ which is a symmetric positive definite (spd) matrix. Then the induced metric $\mathbf{G} = (g_{\mu\nu})$ of the Riemannian space $(\Sigma, \mathbf{G})$ is explicitly given by the spd matrix

$$g_{\mu\nu} = \left(\frac{\partial \mathbf{X}}{\partial \sigma^\mu}\right)^T \cdot \mathbf{H} \cdot \left(\frac{\partial \mathbf{X}}{\partial \sigma^\nu}\right). \tag{1}$$

In our examples on 3D medical intensity data $I(x,y,z)$, we chose an embedding $(X^1, X^2, X^3, X^4) = (x,y,z,I(x,y,z))$. The metric of the embedding space $\mathbb{R}^4$ has been chosen to be the canonical Euclidean one, i.e., $\mathbf{H} = (h_{ij}) = \delta_{ij}$. Thus, we obtain with equation (1) the induced metric

$$\mathbf{G} = \begin{pmatrix} 1+I_x^2 & I_xI_y & I_xI_z \\ I_xI_y & 1+I_y^2 & I_yI_z \\ I_xI_z & I_yI_z & 1+I_z^2 \end{pmatrix}. \tag{2}$$

Let ϕ be a function on Σ, i.e., a function in Σ's parameter space. The function describes a distance function on Σ iff its gradient with respect to the given metric is normalized. It is well known that this is precisely the case if the ordinary gradient $\nabla\phi$ satisfies the Hamilton-Jacobi equation

$$\nabla\phi^T \cdot \mathbf{G}^{-1} \cdot \nabla\phi = 1. \tag{3}$$

In our applications, we will be given a reference zone Γ in parameter space $D \subset \mathbb{R}^n$, then compute its image on the manifold Σ, and the distance function to that. The representation of that distance function in D will simply be called the distance function ϕ of Γ.

Fast sweeping. Tsai et al. [TCOZ03] propose a fast scheme for the numerical solution of a general class of convex Hamilton-Jacobi equations. The authors derive their scheme exemplarily for a distance function ϕ to a set Γ on a 2-dimensional manifold, i.e., solving Hamilton-Jacobi equations of the form (3) with a spd 2×2 matrix $\mathbf{G}^{-1}$. The scheme finds the distance transform in O(N) time, where N is the number of pixels. The generalization to n-dimensional manifolds can be found in [Zha05]. We implemented Tsai et al.'s scheme for 2- and 3-dimensional manifolds, i.e., to the solution of Hamilton-Jacobi equations of the form (3) with a spd 3×3 matrix $\mathbf{G}^{-1}$. The interested reader is referred to [HPRP05] for more details.

3.2. Diffusion Approach

The main disadvantage of the geometric approach is that at least one input point (Γ) inside each object has to be chosen. This can be impractical if a multitude of objects have to be segmented from one data set. Also, elongated structures are hard to characterize by a distance function which is based on

a single reference point since the Euclidean part of the distance tends to produce isodistance surfaces which resemble spheres (in 3D). As mentioned before, data $I \in [0, I_{max}]$ containing anatomical structures which are already well characterized by their rather homogeneous signal σ can be transformed to resemble the distance function ϕ of the geometric approach by linearly mapping $0 \mapsto \sigma, \sigma \mapsto 0, I_{max} \mapsto I_{max} - \sigma$. The main problem then consists in getting rid of small signal inhomogeneities like noise. We propose the use of fast nonlinear diffusion for noise removal based on the semi-implicit operator splitting scheme described by Weickert et al. [Wei98], which shall be briefly outlined here. The so-called CLMC filter calculates filtered data $\phi(\mathbf{x}, t)$ of $I(\mathbf{x})$ as a solution of the diffusion equation

$$\partial_t \phi = div[g(|\nabla \phi_\sigma|^2)\nabla\phi] \tag{4}$$

with the original data as initial state $\phi(\mathbf{x}, 0) = I(\mathbf{x})$ and $\nabla\phi_\sigma$ the gradient of a smoothed version of ϕ. We use ϕ to denote the filtered data here, since it will play the same role in the ensuing soft segmentation process as the distance function ϕ described in the previous section. Weickert proposes a discretization of (4) which leads to a (linear-implicit) semi-implicit iteration scheme

$$[\mathbb{1} - \tau \mathbf{A}(\phi^k)]\phi^{k+1} = \phi^k \tag{5}$$

$$a_{ij}(\phi^k) := \begin{cases} \frac{g_i^k + g_j^k}{2h^2} & [j \in N(i)], \\ -\sum_{n \in N(i)} \frac{g_i^k + g_n^k}{2h^2} & (j = i), \\ 0 & (\text{else}), \end{cases} \tag{6}$$

with $\mathbf{A}(\phi^k) = (a_{ij}(\phi^k))$, $\mathbb{1} \in \mathbb{R}^N$ the unit matrix, g the diffusivity, h the data grid size, $N(i)$ the set of two neighbours of pixel i and τ the diffusion time. This requires to solve a linear system, where the system matrix is tridiagonal and diagonally dominant. It can be efficiently solved using the Thomas algorithm. The semi-implicit iteration scheme for the m-dimensional CLMC equation leads to a linear system which for $m \geq 2$ can no longer be solved using the efficient Thomas algorithm as in the 1-D case. The proposed solution is a modification called additive operator splitting (AOS)

$$\phi^{k+1} = \frac{1}{m} \sum_{l=1}^{m} [\mathbb{1} - m\tau \mathbf{A}_l(\phi^k)]^{-1} \phi^k. \tag{7}$$

It treats all coordinate axes in exactly the same manner (the matrix $\mathbf{A}_l$ corresponds to derivatives along the lth coordinate axis), in contrast to alternative multiplicative operator splitting (MOS) schemes.

We have implemented both the AOS and MOS scheme for processing of 2D and 3D data in an efficient way using functions from Intel's IPP library as well as the Intel C++ compiler with its capability to generate code optimized for Intel Pentium IV. As diffusivity function g we are using $g(\nabla I) = e^{-||\nabla I||}$ (see e.g. [PM90]), where the gradient magnitude $||\nabla I||$ is calculated on a Gaussian smoothed version of the input data.

3.3. Soft Segmentation

We will assume that the function ϕ has been generated using one of the two methods described above and henceforth refer to it as 'distance'. Isodistance hypersurfaces (level sets) of ϕ are crucial in our algorithm, but not taken directly as the segmentation result. Instead, we base our algorithm on the following ideas: (i) Data points belonging to the object will have similar low distance values. (ii) The distance will increase rapidly in areas where the data (object) features change most, i.e., at the object border. (iii) There is one isodistance hypersurface $S(d^*)$ which can serve as best approximation of the actual object border. (iv) Based on d^*, distance values can be mapped to probabilities using an appropriate scheme. Therefore, the rough structure of the algorithm for soft segmentation is as follows:

1. Determine the level set $S(d^*)$ of ϕ, which best approximates the object's boundary; its distance value is d^*.
2. Based on d^*, map distance values to probabilities $p_{1,2}$.
3. If visualization of the object's border with its uncertainty is desired, compute a density function ρ to be used for display.

Finding the best approximating isodistance hypersurface. In [HPRP05] we have described an algorithm for detection of the distance value d^* marking the isodistance hypersurface $S(d^*)$ that approximates the actual object border best for the geometric approach. If diffusion filtered data is used as 'distance', a simpler strategy for detection of $S(d^*)$ based on the cumulative histogram of ϕ has proven fruitful.

Mapping distances to probabilities. Based on the previous calculation of d^*, one can set up a *probability distribution on the distance scale* in order to assess the reliability of the segmentation. For the sake of simplicity we have chosen a Gaussian probability distribution centered at d^*, $G(\phi) = e^{-(\phi - d^*)^2/2\sigma^2}$. It has to be emphasized, that this distribution is *not* aimed at mimicking the statistics that led to the fuzzy appearance of the object boundary, which are a priori unknown! The probability is defined on the distance space and not within the spatial domain. The distance function itself takes care of proper spreading of boundary probabilities across the spatial domain via its varying gradient, i.e., the spatial density of isodistance hypersurfaces, see Fig.2.

Two different probability values arise naturally: (i) the probability p_1 for each data point $\mathbf{x}$ to be part of the object (for an example see Fig.3(b)) and (ii) the probability p_2 for each data point to be part of the object border. For the first one, we assume that all data points $\mathbf{x}$ with $\phi(\mathbf{x}) \leq d^*$ are for sure within the object, whereas probability decreases if $\phi(\mathbf{x}) > d^*$, i.e.,

$$p_1(\mathbf{x}) = \begin{cases} 1 & \text{if } \phi(\mathbf{x}) \leq d^*, \\ G(\phi(\mathbf{x})) & \text{if } \phi(\mathbf{x}) > d^*. \end{cases} \tag{8}$$

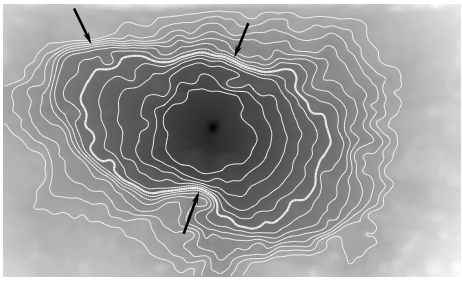

Figure 2: *Cross section of a 3D distance volume (geometric approach) with some equidistant isodistance surfaces. The bold white line marks the isodistance surface $S(d^*)$.*

Density function for visualization. If modeling of the object's border reflecting its uncertainty is the primary goal, we propose a two step procedure. First, probability values p_2 are assigned to distances over the complete range covered by the Gaussian,

$$p_2(\mathbf{x}) = G(\phi(\mathbf{x})). \tag{9}$$

This produces maximum probability at $S(d^*)$ with probabilities fading to both sides. In the final step, the probability values are weighted point-wise with the gradient magnitude of the distance function $\|\nabla\phi\|$ to produce a final density ρ which can be used for display,

$$\rho(\mathbf{x}) = p_2(\mathbf{x}) \cdot \|\nabla\phi(\mathbf{x})\|. \tag{10}$$

Weighting with $\|\nabla\phi\|$ compensates for the fact that $S(d^*)$ is just a guess of the location of the object's border at locations of low distance gradient. An example can be seen in Fig.3(c) and 4(d). The variance of the Gaussian can be used to modulate the *spatial* extent of the zone of high border probability within the data set. Thus, if $S(d^*)$ does not correspond well to the real object's border at some locations, the border probability distribution can still reach into the area of the actual object border and be correctly emphasized by the distance gradient.

Figure 3 demonstrates steps of the soft segmentation process on 2D cross sections through a 3D ultrasound scan of a breast lesion (input point and distance function based approach; rendering see figure 5(d)) while figure 4 demonstrates the process using a diffusion filtered ultrasound scan of liver vessels (rendering see figure 5(e)).

4. Results

One of the main concerns in medical 3D segmentation and visualization is calculation speed, above all when it comes to treating data provided by imaging devices which are operated interactively like ultrasound. The central parts of the segmentation method presented here are the distance function calculation via fast sweeping, respectively the nonlinear diffusion filtering using the semi-implicit operator splitting scheme, for both of which we devised highly efficient implementations.

In contrast to fast marching methods, the sweeping characteristic qualifies for a straightforward parallelized implementation using, e.g., Intel's Streaming SIMD Extensions (SSE) and thus operating on four single-precision floats simultaneously. Our implementation allows for computing ϕ for a 150×252×120 8 bit data set within 1,95 seconds per iteration (eight sweeps) on a 3.0GHz Pentium IV with 2GB of memory. One to two iterations have proven sufficient for this application.

For our implementation of the AOS diffusion we have used functions from Intel's IPP library as well as the Intel C++ compiler with its capability to generate code optimized for Intel Pentium IV. This allows for computing the diffusion process for a 150×252×120 8 bit data set within 700 milliseconds per iteration on the same hardware. One iteration is sufficient since the degree of smoothing can be steered with the diffusion time τ alone when using the diffusion scheme described, the diffusion process being stable, which is one of this algorithm's main advantages.

For further speedup, downsampling of the input data can be an eligible option. Downsampling by a factor 2 in every dimension produces speedup by a factor 2^3 for 3D data if algorithms of $O(N)$ are used. Employing this kind of downsampling, volume data of above mentioned size takes approximately 590ms to segment if the fast sweeping based distance function is used (2 iterations), and 146ms for the diffusion based approach. For all the diffusion based segmentation results presented here, 1 iteration with $\tau = 250$ has been used.

In figure 5 (bottom row) 3D renderings of soft segmentations of three different ultrasound data sets are shown: a breast lesion (see also figure 3), liver vessels (see also figure 4) and a fetal heart, all rendered via DVR with gradient shading using SGI Volumizer [Sil05]. For comparison, results of visualization attempts without prior segmentation are presented as well (top row). E.g., rendering of the distance function gradient magnitude alone still shows too much confusing clutter since the distance function's gradient per se has no "knowledge" of what is noise and what belongs to the object border. Windowing (remapping of signal intensities) alone, which is the easiest solution for structures of homogeneous signal, suffers from similar problems. It is clearly visible how the soft segmentation based visualization not only improves the overall appearance of the renderings but also discriminates between well defined object regions (shiny, surface like appearance) and ill defined object borders (fuzzy, more transparent appearance), like some rims of the breast lesion, the lower rim of the big liver vessel or part of the surface of the upper left heart chamber. Figure 6 shows more renderings of 3D soft segmentation results. Fig 6(a) shows a brain stem segmented from an MR data set, 6(b) shows a rendering of a cerebral ventricle system segmented from US data and 6(d) shows another vessel tree from an US liver scan. The fuzzy look of one of the ventricle's ends and some of the liver vessels indicate data regions of higher uncertainty. The shell-like result of the soft segmentation process makes it possible to cut through objects and actually render views

(a) Original data, reference point indicated by +

(b) Probability p_1 of voxels belonging to the structure under examination

(c) Density volume ρ, intended for direct volume rendering of the object's border

Figure 3: *Cross sections of a 150×252×120 8 bit medical ultrasound data set of a breast lesion (Fig.3(a)) and its soft segmentation. The lesion is sharply defined mostly in the lower central region. Its border is indistinct in the right and very left region. Soft segmentation appropriately handles the fuzziness of the data, see Fig.3(b) and 3(c).*

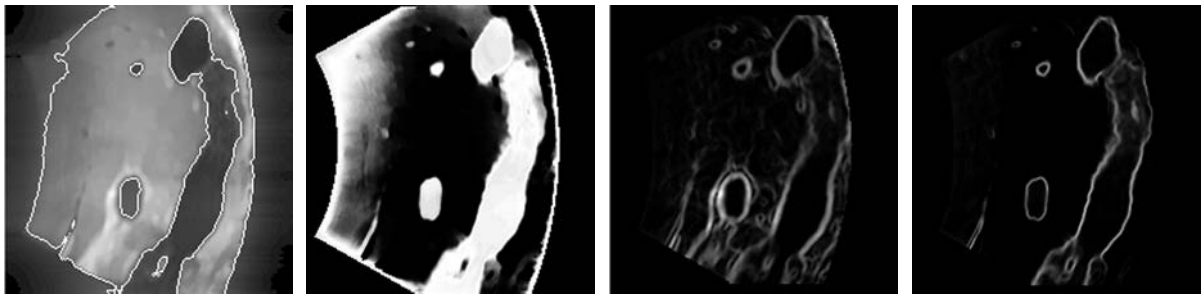

(a) Liver vessels (cross section of diffusion with $S(d^*)$ overlaid)

(b) Unweighted soft segmentation result (probability p_2)

(c) Gradient of diffusion filtered data

(d) Density volume ρ (weighting of 4(b) with 4(c)). DVR see 5(e).

Figure 4: *Demonstration of soft segmentation process using diffusion filtered US data set.*

of their interior structure, which has been attempted in figure 6(c) with the fetal heart from figure 5(f). The chamber walls as well as a connection between two chambers appear clearly.

The visualization of uncertainty in our approach is mainly based on transparency, in contrast to methods which use "artificial" elements like adding glyphs, color encoding or modification of the geometry of the object [PWL97]. In medical visualization, these elements can rather impair the quality of the rendering, while transparency indicates uncertainty or missing information quite naturally.

5. Conclusions

We have presented a technique for "soft" segmentation of data sets of arbitrary dimension n. Two approaches have been presented which are based on generation of distance functions, i.e. n-dimensional auxiliary functions ϕ which show low values inside the object(s) to be segmented and high values outside. The first one uses embedding of the data into a high-dimensional Riemannian space and treats it as a manifold on which distances with respect to a reference point or zone within the object can be calculated. The second one uses nonlinear diffusion filtering. These auxiliary functions are used for segmentation and are mapped to object border probabilities and densities which in turn can serve as input for visualization. Our method thus automatically yields uncertainty information without the need to know the underlying nature of the statistics, object properties, or data acquisition imperfections. Due to fast implementations of both a sweeping algorithm for calculation of a distance function and a nonlinear diffusion filter, our method allows for real-time segmentation of 2D and interactivity for 3D data.

Acknowledgment The authors would like to acknowledge the contributions of Armin Schoisswohl of GE Healthcare Kretztechnik in providing the 3DUS data sets and valuable discussions. This work has been carried out within the K*plus* Competence Center ADVANCED COMPUTER VISION. This work was funded by the K*plus* program.

References

[CKS97] CASELLES V., KIMMEL R., SAPIRO G.: Geodesic active contours. *Intl. J. Computer Vision 22*, 1 (1997), 61–79.

[CLMC92] CATTÉ F., LIONS P. L., MOREL J. M., COLL T.: Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. on Num. Anal. 29* (1992), 182–193.

[CR00] CEDILNIK A., RHEINGANS P.: Procedural annotation of uncertain information. In *Proc. Visualization* (2000), pp. 77–84, 542.

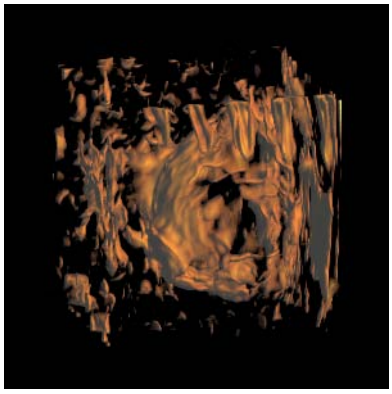

(a) Breast lesion (rendering of distance gradient magnitude)

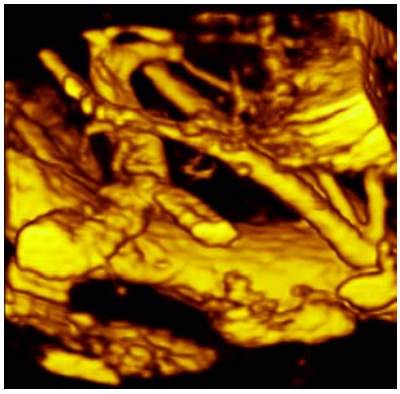

(b) Liver vessels (windowed rendering)

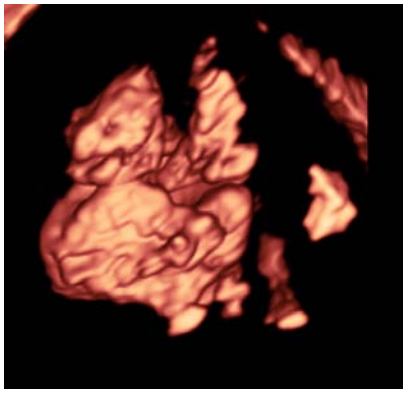

(c) Fetal heart (windowed rendering)

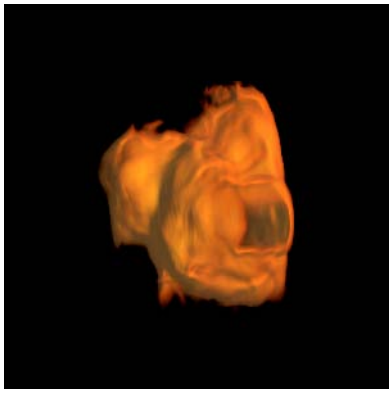

(d) Breast lesion (soft segmentation)

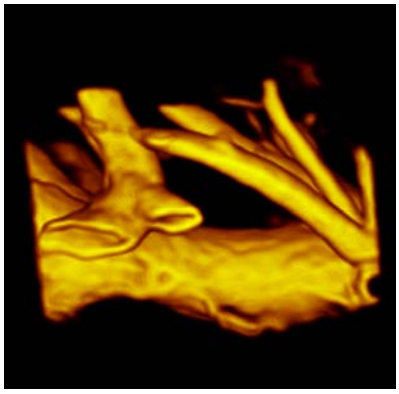

(e) Liver vessels (soft segmentation)

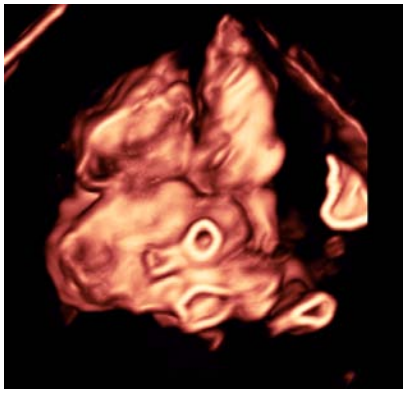

(f) Fetal heart (soft segmentation)

Figure 5: *Renderings of soft segmentation results (bottom row) on different ultrasound data sets and comparison with simpler visualization techniques (top row).*

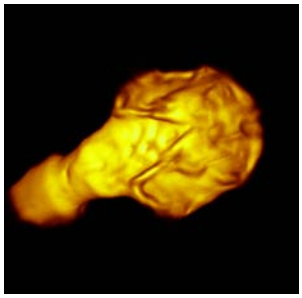

(a) Brain stem segmented from magnetic resonance data

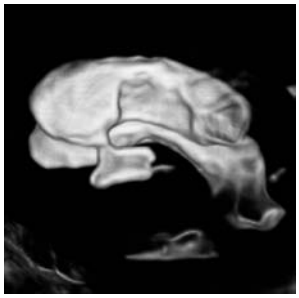

(b) Brain ventricle (US)

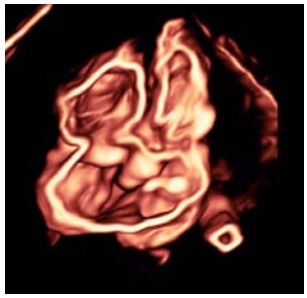

(c) Fetal heart cut open (US; same as in Fig.5(f))

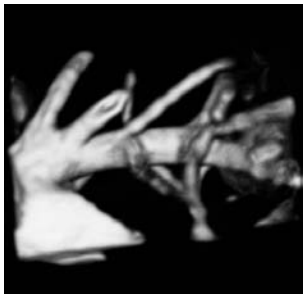

(d) Liver vessels (US)

Figure 6: *Further example renderings of soft segmentation results on different medical data sets.*

[DKLP02] DJURCILOV S., KWANSIK K., LERMUSIAUX P., PANG A.: Visualizing scalar volumetric data with uncertainty. *Computers & Graphics 26*, 2 (2002), 239–248.

[FSS*02] FENSTER A., SURRY K., SMITH W., ET AL.: 3d ultrasound imaging: applications in image-guided therapy and biopsy. *Computers and Graphics 26*, 4 (2002), 557–568.

[Gil02] GILBOA G.: Forward-and-backward diffusion processes for adaptive image enhancement and denoising. *IEEE Trans. Image Process. 11*, 7 (2002), 689–703.

[GR04] GRIGORYAN G., RHEINGANS P.: Point-based probabilistic surfaces to show surface uncertainty. *IEEE Trans. on Visualization and Computer Graphics 10*, 5 (2004), 564–573.

[HPRP05] HÖNIGMANN D., PETERSCH B., RUISZ J., POTTMANN H.: Manifold-based soft segmentation. *submitted* (2005).

[HRSB05] HONG A. S., ROSEN E. L., SOO M. S., BAKER J. A.: Bi-rads for sonography. *American Journal of Roentgenology 184* (2005), 1260–1265.

[JS03] JOHNSON C. R., SANDERSON A. R.: A next step: Visualizing errors and uncertainty. *IEEE Computer Graphics and Applications 23*, 5 (2003), 6–10.

[Kim04] KIMMEL R.: *Numerical Geometry of Images*. Springer, 2004.

[KLDP04] KAO D. L., LOVE A., DUNGAN J. L., PANG A.: Picturing data with uncertainty. In *Proc. SIGGRAPH* (2004).

[KMS00] KIMMEL R., MALLADI R., SOCHEN N.: Images as embedded maps and minimal surfaces: movies, color, texture, and volumetric medical images. *Int. J. of Computer Vision 39*, 2 (2000), 111–129.

[KS02] KIMMEL R., SOCHEN N.: Orientation diffusion or how to comb a porcupine. *J. of Visual Communication and Image Representation 13*, 1-2 (2002), 238–248.

[KvU05] KNISS J., VAN UITERT R.: Statistically quantitative volume visualization. In *VIS'05, IEEE* (2005), pp. 287–294.

[LP05] LOVE A. L., PANG A.: Visualizing spatial multivalue data. *IEEE Computer Graphics and Applications 25*, 3 (2005), 69–79.

[MS01] MEMOLI F., SAPIRO G.: Fast computation of weighted distance functions and geodesics on implicit hyper-surfaces. *J. Comput. Physics 173*, 2 (2001), 730–764.

[MSKB*04] MEYBERG-SOLOMAYER G. C., KRAEMER B., ET AL.: Does 3-d sonography bring any advantage to noninvasive breast diagnostics? *Ultrasound in Medicine & Biology 30*, 5 (2004), 583–589.

[MvHM02] MCKENZIE A., VANHERK M., MIJNHEER B.: Margins for geometric uncertainty around organs at risk in radiotherapy. *Radiotherapy and Oncology 62*, 3 (2002), 299–307.

[OF02] OSHER J. S., FEDKIW R. P.: *Level Set Methods and Dynamic Implicit Surfaces*, 1 ed. Springer, 2002.

[PM90] PERONA P., MALIK J.: Scale-space and edge detection using anisotropic diffusion. *IEEE TPAMI 12*, 7 (1990), 629–39.

[PSH*04] POTTMANN H., STEINER T., HOFER M., ET AL.: The isophotic metric and its application to feature sensitive morphology on surfaces. In *ECCV04* (2004), vol. 3024, pp. 560–572.

[PT01] POKRIC M., THACKER A.: Multi-dimensional medical image segmentation with partial voluming. In *MIUA 01* (2001), University of Birmingham.

[PWL97] PANG A. T., WITTENBRINK C. M., LODHA S. K.: Approaches to uncertainty visualization. *The Visual Computer 13*, 8 (1997), 370–390.

[RFDA97] RANGAYYAN R., FARAMAWY N., ET AL.: Measures of acutance and shape for classification of breast tumors. *IEEE T. Med. Imaging 16*, 6 (1997), 799–810.

[Sap01] SAPIRO G.: *Geometric Partial Differential Equations and Image Analysis*. Cambridge Univ. Press, 2001.

[SDP03] SOCHEN N., DERICHE R., PEREZ L. L.: The beltrami flow over implicit manifolds. In *Computer Vision, Proc. 9th IEEE Int. Conf. on* (2003), pp. 832–839.

[Set99] SETHIAN J. A.: *Level set methods and fast marching methods*. Cambridge monographs on applied and computational mathematics. Cambridge Univ. Press, 1999.

[Sil05] SILICONGRAPHICS: http://www.sgi.com/products/software/volumizer/.

[SK04] SPIRA A., KIMMEL R.: An efficient solution to the eikonal equation on parametric manifolds. *Interfaces and Free Boundaries 6*, 3 (2004), 315–327.

[Soc98] SOCHEN N.: A general framework for low level vision. *IEEE Trans. Image Process. 7*, 3 (1998), 310–318.

[SSG95] SAKAS G., SCHREYER L.-A., GRIMM M.: Preprocessing and volume rendering of 3d ultrasonic data. *IEEE Computer Graphics and Applications 15*, 4 (1995), 47–54.

[STR95] STAVROS A. T., THICKMAN D., RAPP C. L.: Solid breast nodules: use of sonography to distinguish between benign and malignant lesions. *Radiology 196* (1995), 123–134.

[TCOZ03] TSAI Y. R., CHENG L. T., OSHER S., ZHAO H. K.: Fast sweeping algorithms for a class of hamilton-jacobi equations. *SIAM J. on Num. Anal. 41*, 2 (2003), 673–694.

[US96] UDUPA J. K., SAMARASEKERA S.: Fuzzy connectedness and object definition: theory, algorithms, and applications in image processing. *CVGIP: Graphical Models and Image Processing 58*, 3 (1996), 246–261.

[vHRRL00] VAN HERK M., REMEIJER P., ET AL.: The probability of correct target dosage: dose-population histograms for deriving treatment margins in radiotherapy. *Int. J. of Rad.Oncol.Biol.Phys. 47*, 4 (2000), 1121–1135.

[Wei98] WEICKERT J.: Efficient and reliable schemes for nonlinear diffusion filtering. *IEEE Trans. Image Process. 7*, 3 (1998), 398–410.

[Zha05] ZHAO H. K.: A fast sweeping method for eikonal equations. *Mathematics of Computation 74*, 250 (2005), 603–627.