

EFFICIENT KERNEL DENSITY ESTIMATION OF SHAPE AND INTENSITY PRIORS FOR LEVEL SET SEGMENTATION

Daniel Cremers

*Department of Computer Science,
University of Bonn
Bonn, Germany*

Mikael Rousson

*Department of Imaging and Visualization
Siemens Corporate Research,
Princeton, New Jersey, USA*

We propose a nonlinear statistical shape model for level set segmentation that can be efficiently implemented. Given a set of training shapes, we perform a kernel density estimation in the low-dimensional subspace spanned by the training shapes. In this way, we are able to combine an accurate model of the statistical shape distribution with efficient optimization in a finite-dimensional subspace. In a Bayesian inference framework, we integrate the nonlinear shape model with a nonparametric intensity model and a set of pose parameters that are estimated in a more direct data-driven manner than in previously proposed level set methods. Quantitative results show superior performance (regarding runtime and segmentation accuracy) of the proposed nonparametric shape prior over existing approaches.

1. INTRODUCTION

Originally proposed in [1, 2] as a means to propagate interfaces in time, the level set method has become increasingly popular as a framework for image

Address all correspondence to: Daniel Cremers, Department of Computer Science, University of Bonn, Bonn, Germany. Phone: (49)228-734380; Fax: (49)228-734382. dcremers@cs.uni-bonn.de.

segmentation. The key idea is to represent an interface $\Gamma \subset \Omega$ in the image domain $\Omega \subset \mathbb{R}^3$ implicitly as the zero level set of an embedding function $\phi : \mathbb{R}^3 \rightarrow \Omega$:

$$\Gamma = \{x \in \Omega \mid \phi(x) = 0\}, \quad (1)$$

and to evolve Γ by propagating the embedding function ϕ according to an appropriate partial differential equation. The first applications of this level set formalism for the purpose of image segmentation were proposed in [3, 4, 5]. Two key advantages over explicit interface propagation are the independence of a particular parameterization and the fact that the implicitly represented boundary Γ can undergo topological changes such as splitting or merging. This makes the framework well suited for the segmentation of several objects or multiply connected objects.

When segmenting medical images, one commonly has to deal with noise, and missing or misleading image information. For certain imaging modalities such as ultrasound or CT, the structures of interest do not differ much from their background in terms of their intensity distribution (see Figure 1). Therefore, they can no longer be accurately segmented based on the image information alone. In recent years, researchers have therefore proposed to enhance the level set method with statistical shape priors. Given a set of training shapes, one can impose information about which segmentations are *a priori* more or less likely. Such prior shape information was shown to drastically improve segmentation results in the presence of noise or occlusion [6, 7, 8, 9, 10, 11]. Most of these approaches are based on the assumption that the training shapes, encoded by their signed distance function, form a Gaussian distribution. This has two drawbacks: First, the space of signed distance functions is not a linear space; therefore, the mean shape and linear combinations of eigenmodes are typically no longer signed distance functions. Second, even if the space were a linear space, it is not clear why the given set of sample shapes should be distributed according to a Gaussian density. In fact, as we will demonstrate in this work, they are generally not Gaussian distributed. Recently, it was proposed to use nonparametric density estimation in the space of level set functions [8] in order to model nonlinear distributions of training shapes. (The term *nonlinear* refers to the fact that the manifold of permissible shapes is not merely a linear subspace.) While this resolves the above problems, one sacrifices the efficiency of working in a low-dimensional subspace (formed by the first few eigenmodes) to a problem of infinite-dimensional optimization.

In the present chapter, we propose a framework for knowledge-driven level set segmentation that integrates three contributions.¹ First, we propose a statistical shape prior that combines the efficiency of low-dimensional PCA-based methods with the accuracy of nonparametric statistical shape models. The key idea is to perform kernel density estimation in a linear subspace that is sufficiently large to embed all training data. Second, we propose to estimate pose and translation parameters in a more data-driven manner. Thirdly, we optimally exploit the intensity information in the image by using probabilistic intensity models given by kernel density estimates of previously observed intensity distributions.




Figure 1. Segmentation challenges and estimated intensity distributions. The two curves on the right correspond to the empirical probability of intensities inside and outside the left ventricle (for the ultrasound image) and the prostate (for the CT image). The region-based segmentation of these structures is a challenging problem, because objects and background have similar histograms. Our segmentation scheme optimally exploits the estimated probabilistic intensity models. See attached CD for color version.

2. LEVEL SET SEGMENTATION AS BAYESIAN INFERENCE

The goal of level set segmentation can be formulated as the estimation of the optimal embedding function $\phi: \Omega \rightarrow \mathbb{R}$ given an image $I: \Omega \rightarrow \mathbb{R}$. In the Bayesian framework, this can be computed by maximizing the posterior distribution

$$\mathcal{P}(\phi | I) \propto \mathcal{P}(I | \phi) \mathcal{P}(\phi). \quad (2)$$

The maximization of (2) results in a problem of infinite-dimensional optimization. Given a set of training shapes encoded by their signed distance functions $\{\phi_i\}_{i=1 \dots N}$, Tsai et al. [7] proposed reducing the segmentation problem to one of finite-dimensional optimization by constraining the optimization problem to the finite-dimensional subspace spanned by the training shapes.

In this chapter we make use of this compact representation of the embedding function. Given the distance d on the space of signed distance functions defined by $d^2(\phi_1, \phi_2) = \int_{\Omega} (\phi_1(x) - \phi_2(x))^2 dx$, we align the set of training shapes with respect to translation and rotation. Subsequently, we constrain the level set function ϕ to a parametric representation of the form:

$$\phi_{\alpha, h, \theta}(x) = \phi_0(R_{\theta}x + h) + \sum_{i=1}^n \alpha_i \psi_i(R_{\theta}x + h), \quad (3)$$

where $\phi_0(x) = \frac{1}{N} \sum_{i=1}^N \phi_i(x)$ represents the mean shape, $\{\psi_i(x)\}_{i=1 \dots n}$ are the eigenmodes of the distribution, and $n < N$ is the dimension of the subspace spanned by the N training shapes. The parameter vector $\alpha = (\alpha_1, \dots, \alpha_n)$ models shape deformations, while the parameters $h \in \mathbb{R}^3$ and $\theta \in [0, 2\pi]^3$ model translation and rotation of the respective shape. In our applications, where the scale of objects is known, a generalization to larger transformations groups (e.g., similarity or affine) did not appear useful.

The infinite-dimensional Bayesian inference problem in Eq. (2) is therefore reduced to a finite-dimensional one where the conditional probability,

$$\mathcal{P}(\boldsymbol{\alpha}, h, \theta | I) \propto \mathcal{P}(I | \boldsymbol{\alpha}, h, \theta) \mathcal{P}(\boldsymbol{\alpha}, h, \theta), \quad (4)$$

is optimized with respect to the shape parameters $\boldsymbol{\alpha}$, and the transformation parameters h and θ . In the following, we will assume a uniform prior on these transformation parameters, i.e., $\mathcal{P}(\boldsymbol{\alpha}, h, \theta) = \mathcal{P}(\boldsymbol{\alpha})$. In the next section we will discuss three solutions to model this shape prior.

3. EFFICIENT NONPARAMETRIC STATISTICAL SHAPE MODEL

Given a set of aligned training shapes $\{\phi_i\}_{i=1\dots N}$, we can represent each of them by their corresponding shape vector $\{\boldsymbol{\alpha}_i\}_{i=1\dots N}$. In this notation, the goal of statistical shape learning is to infer a statistical distribution $\mathcal{P}(\boldsymbol{\alpha})$ from these sample shapes. Two solutions that have been proposed are based on the assumptions that the training shapes can be approximated by a **uniform distribution** [7, 9]: $\mathcal{P}(\boldsymbol{\alpha}) = \text{const.}$, or by a **Gaussian distribution** [6]:

$$\mathcal{P}(\boldsymbol{\alpha}) \propto \exp(-\boldsymbol{\alpha}^\top \Sigma^{-1} \boldsymbol{\alpha}), \quad \text{where } \Sigma = \frac{1}{N} \sum_i \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^\top. \quad (5)$$

In the present chapter we propose to make use of nonparametric density estimation [12] to approximate the shape distribution within the linear subspace. We model the shape distribution by the **kernel density estimate**:

$$\mathcal{P}(\boldsymbol{\alpha}) = \frac{1}{N\sigma^n} \sum_{i=1}^N K\left(\frac{\boldsymbol{\alpha} - \boldsymbol{\alpha}_i}{\sigma}\right), \quad \text{where } K(u) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{u^2}{2}\right). \quad (6)$$

There exist various methods to automatically estimate appropriate values for the width σ of the kernel function, ranging from the k th nearest neighbor estimates to cross-validation and bootstrapping. In this work, we simply set σ to be the average nearest neighbor distance: $\sigma^2 = \frac{1}{N} \sum_{i=1}^N \min_{j \neq i} |\boldsymbol{\alpha}_i - \boldsymbol{\alpha}_j|^2$.

In the context of level set-based image segmentation, the kernel density estimator (6) has two advantages over the uniform and Gaussian distributions:

- The assumptions of uniform distribution or Gaussian distribution are generally not fulfilled. In Figure 3, we demonstrate this for a set of silhouettes of sample shapes. The kernel density estimator, on the other hand, is known to approximate arbitrary distributions. Under mild assumptions, it was shown to converge to the true distribution in the limit of infinite sample size [13].





Figure 2. Schematic plots of different density estimates within a subspace. Darker shading indicates areas of high probability density for the respective models. The kernel density estimator adapts to the training data more flexibly since it does not rely on specific assumptions about the shape of the distribution.

- The space of signed distance functions is known to not be a linear space. Therefore, neither the mean shape ϕ_0 nor a linear combination of eigenmodes as in (3) will in general be a signed distance function. As a consequence, the functions $\phi(x)$ favored by the uniform or the Gaussian distribution cannot be expected to be signed distance functions. The kernel density estimator (6), on the other hand, favors shape vectors α , which are in the vicinity of the sample shape vectors α_i . By construction, these vectors correspond to signed distance functions. In fact, **in the limit of infinite sample size, the distribution inferred by the kernel density estimator (6) converges toward a distribution on the manifold of signed distance functions.**

Figure 2 shows schematic plots of the three methods for a set of sample data spanning a two-dimensional subspace in \mathbb{R}^3 . The kernel density estimator clearly captures the distribution most accurately. As we shall see in Section 5, constraining a level set-based segmentation process by this nonparametric shape prior will allow to compute accurate segmentations even for rather challenging image modalities.

Figure 3 shows a 3D projection of the estimated shape density computed for a set of silhouettes of a walking person. The bottom row shows shape morphing by sampling along geodesics of the uniform and the kernel density. These indicate that the kernel estimator captures the distribution of valid shapes more accurately.

In analogy to shape learning, we make use of kernel density estimation to learn the conditional probability for the intensity function I in (4) from examples. A similar precomputation of intensity distributions by means of mixture models was proposed in [14]. Given a set of presegmented training images, the kernel density estimate of the intensity distributions p_{in} and p_{out} of object and background are given by the corresponding smoothed intensity histograms. This has two advantages. First, the kernel density estimator does not rely on specific assumptions about the shape of the distribution. Figure 1 shows that the intensity distributions for ultrasound and CT images are not well approximated by Gaussian or Lapla-



Sampling along geodesic (uniform dens.) Sampling along geodesic (kernel dens.)

Figure 3. Linear versus nonlinear shape interpolation. The upper row shows 6 out of 49 training shapes and a 3D projection of the isosurface of the estimated (48-dimensional) shape distribution. The latter is clearly neither uniform nor Gaussian. The bottom row shows a morphing between two sample shapes along geodesics induced by a uniform or a kernel distribution. The uniform distribution induces a morphing where legs disappear and reappear and where the arm motion is not captured. The nonlinear sampling provides more realistic intermediate shapes. We chose human silhouettes because they exhibit more pronounced shape variability than most medical structures we analyzed.

cian models. Second, in contrast to the joint estimation of intensity distributions (cf. [15]), this simplifies the segmentation process, which no longer requires an updating of intensity models. Moreover, we found the segmentation process to be more robust to initialization in numerous experiments.

4. ENERGY FORMULATION AND MINIMIZATION

Maximizing the posterior probability in (2), or equivalently minimizing its negative logarithm, will generate the most probable segmentation of a given image. With the nonparametric models for shape and intensity introduced above, this leads to an energy of the form

$$E(\boldsymbol{\alpha}, h, \theta) = -\log \mathcal{P}(I|\boldsymbol{\alpha}, h, \theta) - \log \mathcal{P}(\boldsymbol{\alpha}). \quad (7)$$

The nonparametric intensity model permits to express the first term, and equation (6) gives exactly the second one. With the Heaviside step function H and the short hand $H_\phi = H(\phi_{\boldsymbol{\alpha}, h, \theta}(x))$, we end up with

$$E(\boldsymbol{\alpha}, h, \theta) = -\int_{\Omega} H_\phi \log p_{\text{in}}(I) + (1-H_\phi) \log p_{\text{out}}(I) dx - \log \left(\frac{1}{N\sigma} \sum_{i=1}^N K \left(\frac{\boldsymbol{\alpha} - \boldsymbol{\alpha}_i}{\sigma} \right) \right).$$

With $e(x) = \left[\log \frac{p_{\text{out}}(I(x))}{p_{\text{in}}(I(x))} \right]$, $K_i = K\left(\frac{\alpha - \alpha_i}{\sigma}\right)$, and $\psi = (\psi_1, \dots, \psi_n)$, we obtain the following system of coupled gradient descent equations:

$$\begin{cases} \frac{d\alpha}{dt} = \int_{\Omega} \delta(\phi_{\alpha, h, \theta}(x)) \psi(R_{\theta}x + h) e(x) dx + \frac{1}{\sigma^2} \frac{\sum_{i=1}^N (\alpha_i - \alpha) K_i}{\sum_{i=1}^N K_i}, \\ \frac{dh}{dt} = \int_{\Omega} \delta(\phi_{\alpha, h, \theta}(x)) \nabla \phi_{\alpha, h, \theta}(x) e(x) dx, \\ \frac{d\theta}{dt} = \int_{\Omega} \delta(\phi_{\alpha, h, \theta}(x)) (\nabla \phi_{\alpha, h, \theta}(x) \cdot \nabla_{\theta} R x) e(x) dx. \end{cases} \quad (8)$$

In all equations, the Dirac delta function δ appears as a factor inside the integrals over the image domain Ω . This allows to restrict all computations to a narrow band around the zero crossing of ϕ . While the evolution of translation and pose parameters h and θ are merely driven by the data term $e(x)$, the shape vector α is additionally drawn toward each training shape with a strength that decays exponentially with the distance to the respective shape.

5. EXPERIMENTAL RESULTS AND VALIDATION

5.1. Heart Segmentation from Ultrasound Images

Figures 4–6 show experimental results obtained for the segmentation of the left ventricle in 2D cardiac ultrasound sequences, using shape priors constructed from a set of 21 manually segmented training images.

The segmentation in Figure 4 was obtained by merely imposing a small constraint on the length of the segmenting boundary. As a consequence, the segmentation process leaks into all darker areas of the image. The segmentation of the left ventricle based on image intensities and purely geometric regularity constraints clearly fails.

The segmentation in Figure 5 was obtained by constraining the shape optimization to the linear subspace spanned by the eigenmodes of the embedding function of the training set. This improves the segmentation, providing additional regularization and reducing the degrees of freedom for the segmentation process. Nevertheless, even within this subspace there is some leakage into darker image areas.

The segmentation in Figure 5 was obtained by additionally imposing a non-parametric statistical shape prior within this linear subspace. While the subspace allows for efficient optimization (along a small number of eigenmodes), the non-parametric prior allows to accurately constrain the segmentation process to a sub-manifold of familiar shapes (see also Figure 2). This prevents any leakage of the boundary and enables the segmentation of the left ventricle despite very limited and partially misleading intensity information.

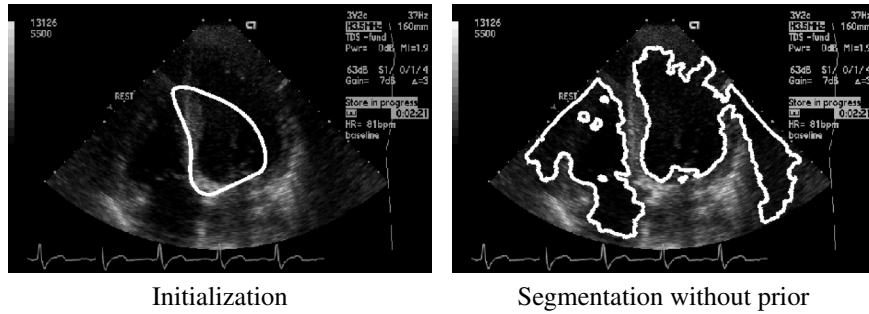


Figure 4. Segmentation without prior. Since there is no shape constraint imposed upon the contour — other than a small length constraint present in the Chan-Vese model — the boundary gradually separates brighter from darker areas. This indicates that intensity information is insufficient to induce the desired segmentation.




Figure 5. Boundary evolution for an ultrasound segmentation with uniform shape prior. By constraining the level set evolution to the linear subspace spanned by the first few eigenmodes computed from a set of training shapes, one can improve the segmentation of the given image (see, e.g., [7]). Nevertheless, in our application, the uniform shape prior does not sufficiently constrain the segmentation process, permitting the boundary to leak into darker image areas.




Figure 6. Boundary evolution for an ultrasound segmentation with nonparametric shape prior. Imposing a non-parametric shape prior within the eigenspace spanned by the training shapes leads to a segmentation process that is sufficiently constrained to enable an accurate segmentation of the left ventricle. In contrast to the uniform prior (see Figure 5), the nonparametric one does suppress leaking of the boundary, because it constrains the level set function to a well-defined submanifold around the training shapes (see also Figure 2).

As a quantitative evaluation we computed the percentage of correctly classified object pixels and that of misclassified ones. During energy minimization, the percentage of correctly classified pixels increases from 56 to 90%, while the percentage of false positives decreases from 27 to 2.7% by using the kernel prior. Using the uniform prior, we attain 92% correctly classified, yet the percentage of




Figure 7. Prostate segmentation for two patients with the same shape model. Each row shows axial slices of the same segmentation for one patient. The manual segmentation is in black and the automatic one white.

false positives increases to 42%. Merely constraining the boundary evolution to the linear subspace spanned by the training shapes is insufficient to provide for accurate segmentation results.

5.2. Prostate Segmentation from 3D CT Images

5.2.1. A single statistical shape model for different patients?

Segmentation of the prostate from CT images is an important and challenging problem in radiotherapy. It may help to avoid the exposure to radiation of vital organs that are not infected by the cancer. In this image modality, the prostate appears with an intensity level very close to the one of adjacent organs like the bladder. The key assumption of our work is that the shape of the prostate in a given segmentation task is statistically similar to prostate shapes observed in a training set. Most related works on prostate segmentation are indeed model-based [7, 10, 11]. In contrast to existing works, we will show that a single (sufficiently sophisticated) statistical shape model can be applied to the segmentation of different patients.

To this end, we built a nonparametric 3D shape model of the prostate using 12 manually extracted prostates (with seminal vesicles) collected from two different patients.

We employed a leave-one-out strategy by removing the image of interest from the training phase. Figure 7 shows 2D cuts of a few results obtained using this strategy. With a one-click initialization inside the organ, the algorithm led to a steady-state solution in less than 10 seconds. We obtained 86% successfully classified organ voxels and 11% misclassified organ voxels. This compares favorably to the intra-patient results reported in [11]. One should note that these quantitative evaluations underestimate the quality of our results since the “ground-truth” seg-




Figure 8. Comparison of the segmentations obtained with the kernel prior (white) and with alternative approaches (black).

segmentations are in general not perfect. Figure 7 provides qualitative comparisons to the manual segmentation, as well as to the segmentations obtained with uniform and Gaussian approximations of the shape distribution.

5.2.2. Quantitative analysis of segmentation accuracy

To further quantify the segmentation accuracy, we consider three different criteria: the Dice coefficient, the average surface distance, and the centroid distance. The Dice coefficient is defined as

$$DSC = \frac{2|S_{\text{manual}} \cap S_{\text{auto}}|}{|S_{\text{manual}}| + |S_{\text{auto}}|}, \quad (9)$$

where $|S_{\text{manual}}|$ and $|S_{\text{auto}}|$ are the volumes of the manual and automatic segmentations, and $|S_{\text{manual}} \cap S_{\text{auto}}|$ is the volume of their intersection. This coefficient can be expressed directly with the level set representations:

$$DSC = \frac{2 \int_{\Omega} H(\phi_{\text{manual}}) H(\phi_{\text{auto}}) dx}{\int_{\Omega} H(\phi_{\text{manual}}) dx + \int_{\Omega} H(\phi_{\text{auto}}) dx}. \quad (10)$$

In general, a value of DSC superior to 0.7 is considered a good agreement. The other two criteria can also be expressed in similar manner. The average surface distance is given by

$$D_{\text{surface}} = \frac{1}{2} \left(\frac{\int_{\Omega} |\nabla H(\phi_{\text{manual}})| |\phi_{\text{auto}}| dx}{\int_{\Omega} |\nabla H(\phi_{\text{manual}})| dx} + \frac{\int_{\Omega} |\nabla H(\phi_{\text{auto}})| |\phi_{\text{manual}}| dx}{\int_{\Omega} |\nabla H(\phi_{\text{auto}})| dx} \right). \quad (11)$$

Essentially, this quantity amounts to averaging the distance of each contour point on one contour to the nearest contour point on the other contour (and vice versa).

The centroid distance is the distance between the centers of mass:

$$c_{\text{manual/automatic}} = \frac{\int_{\Omega} x H(\phi_{\text{manual/automatic}}) dx}{\int_{\Omega} H(\phi_{\text{manual/automatic}}) dx}.$$

Table 1. Quantitative Validation on 26 CT Images

	DSC	$D_{surface}$ (mm)	$D_{centroid}$ (mm)
Average	0.8172	3.38	2.86
Standard deviation	0.0807	1.11	1.63
Minimum	0.6573	2.07	0.50
Maximum	0.9327	5.42	5.75

One should point out, however, that the centroid distance has only a very limited capacity to quantify shape differences. Obviously, it cannot distinguish between any two segmentations that share the same centroid.

Table 1 gives the average value of all three criteria computed for the entire dataset in the leave-one-out strategy mentioned above. In addition, we displayed the standard deviation, minimum, and maximum value of each criterion. Overall, these values show that our segmentations typically agree well with the manual ground truth.

5.2.3. Robustness to initialization

The level set method for image segmentation and also its implementation with nonparametric statistical shape priors are *local* optimization methods. As a consequence, experimental results will depend on the initialization. This aspect is a common source of criticism, it is generally believed that *local* indicates that segmentations can only be obtained if contours are initialized in the vicinity of the desired segmentation. Yet, this is not the case for the region-based segmentation schemes like the one developed in this work. The segmentation without shape prior in Figure 4 shows a drastic difference between initial and final boundary: clearly contours can propagate over large spatial distances from the initialization to the “nearest” local minimum.

In order to quantify the robustness of our method to initialization, we translated the initialization by a certain distance in opposite directions and subsequently computed the accuracy of the resulting segmentation process with nonparametric shape prior. Table 2 shows that the accuracy is quite robust with respect to displacements of the initialization up to 10mm in each direction.

5.2.4. Robustness to noise

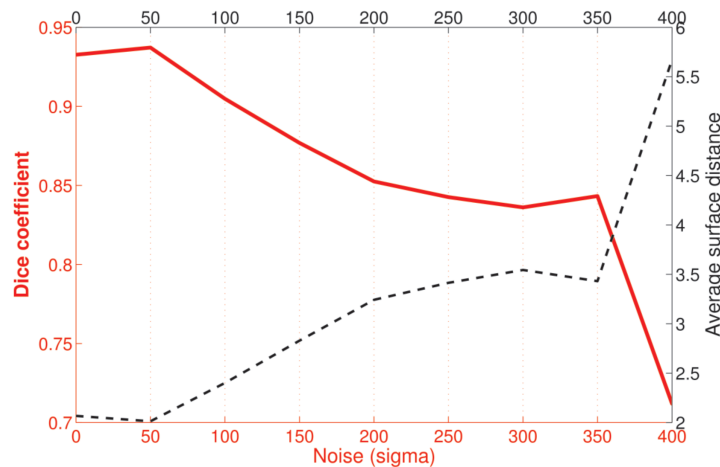
The prostate CT images are in themselves rather challenging, since prostate and surrounding tissue have fairly similar intensities (see Figure 1, right side). The combination of statistically learned nonparametric models of both the intensity

Table 2. Robustness to Initialization

X Translation (mm)	-10	-5	0	5	10
DSC	0.9287	0.9297	0.9327	0.9289	0.9300
$D_{surface}$ (mm)	2.1358	2.0910	2.0673	2.1080	2.1105
$D_{centroid}$ (mm)	1.3942	1.4657	1.4826	1.4685	1.5481

distribution and the distribution of the shape embedding functions nevertheless allows to compute the desired segmentation. Yet, one may ask where the limitations of our model are. At what point does segmentation accuracy break down?

To investigate this, we artificially added noise to the images, computing at each time the segmentation accuracy. Figure 9 shows both the Dice coefficient defined in (10) and the average surface distance defined in (11) of the final segmentation as a function of the noise. While the segmentation is rather good over a large range of noise values, it does decay at very large values of noise.

**Figure 9.** Robustness to noise. See attached CD for color version.

5.2.5. Efficiency versus accuracy: How many eigenmodes are needed?

The efficiency of our implementations arises because we solve the level set computation in the low-dimensional linear subspace spanned by the training shapes. Given N training shapes, this will typically amount to an optimization of $N - 1$ parameters.

Table 3. Segmentation Accuracy for Different Numbers of Modes

	DSC	$D_{surface}$ (mm)	$D_{centroid}$ (mm)
3 modes	0.8015	3.55	3.32
10 modes	0.8172	3.38	2.86
25 modes	0.8173	3.46	2.95

While there exist many ways to parameterize this subspace, the representation in terms of principal components (eigenshapes of the embedding function) has the additional advantage that the principal components associated with the largest eigenvalues by definition capture the largest variation of the embedding function. Hence, one could further reduce the dimensionality of the problem, by using merely the first few eigenmodes.

To quantify the loss in segmentation accuracy when using fewer eigenmodes in the optimization, we show in Table 3 the values of the Dice coefficient, the surface distance, and the centroid distance obtained when using 3, 10, and 25 eigenmodes. The reported quantities are averages computed for each of the 25 test images. As expected, the higher-order eigenmodes contain very little additional shape information, so that the accuracy increases only by a little amount when going from 10 to 25 eigenmodes, while the computation time scales linearly with the number of eigenmodes considered.

6. CONCLUSION

We proposed herein an efficient and accurate statistical shape prior for level set segmentation that is based on nonparametric density estimation in the linear subspace spanned by the level set surfaces of a set of training shapes. In addition, our segmentation scheme integrates nonparametric estimates of intensity distributions and efficient optimization of pose and translation parameters.

We reported quantitative evaluation of segmentation accuracy and speed for cardiac ultrasound images and for 3D CT images of the prostate. In particular, we quantitatively validated that the proposed segmentation scheme is robust to the initialization and robust to noise. Furthermore, we demonstrated that one can increase efficiency by reducing the number of eigenmodes considered in the optimization while losing a little accuracy of the average segmentation results. These results indicate that the proposed nonparametric shape prior outperforms previously proposed shape priors for level set segmentation.

7. ACKNOWLEDGMENTS

We thank Christophe Ched'hotel for fruitful discussions. We thank Marie-Pierre Jolly for providing us with image and training data for the ultrasound sequences.

8. NOTES

1. A preliminary version of this work was presented in [16]

9. REFERENCES

1. Dervieux A, Thomasset F. 1979. A finite element method for the simulation of Rayleigh-Taylor instability. In *Approximation methods for Navier-Stokes problems*, pp. 145–158. Ed R Rautmann. Berlin: Springer.
2. Osher SJ, Sethian JA. 1988. Front propagation with curvature dependent speed: algorithms based on Hamilton-Jacobi formulations. *J Comput Phys* **79**:12–49.
3. Caselles V, Catté F, Coll T, Dibos F. 1993. A geometric model for active contours in image processing. *Num Math* **66**:1–31.
4. Malladi R, Sethian JA, Vemuri BC. 1994. A topology independent shape modeling scheme. In *Proceedings of the SPIE conference on geometric methods in computer vision*, Vol. 2031, pp. 246–258. Bellingham, WA: SPIE.
5. Kichenassamy S, Kumar A, Olver PJ, Tannenbaum A, Yezzi AJ. 1995. Gradient flows and geometric active contour models. In *Proceedings of the fifth international conference computer vision (ICCV'95)*, pp. 810–815. Washington, DC: IEEE Computer Society.
6. Leventon M, Grimson W, Faugeras O. 2000. Statistical shape influence in geodesic active contours. In *Proceedings of the IEEE international conference on computer vision and pattern recognition (CVPR)*, Vol. 1, pp. 316–323. Washington, DC: IEEE Computer Society.
7. Tsai A, Yezzi AJ, Willsky AS. 2003. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Trans Med Imaging*, **22**(2):137–154.
8. Cremers D, Osher SJ, Soatto S. 2006. Kernel density estimation and intrinsic alignment for shape priors in level set segmentation. *Int J Comput Vision*. **69**(3):335–351.
9. Rousson M, Paragios N, Deriche R. 2004. Implicit active shape models for 3d segmentation in MRI imaging. In *Proceedings of the international conference on medical image computing and computer-assisted intervention (MICCAI 2000)*. *Lecture notes in computer science*, Vol. 2217, pp. 209–216. New York: Springer.
10. Dam EB, Fletcher PT, Pizer S, Tracton G, Rosenman J. 2004. Prostate shape modeling based on principal geodesic analysis bootstrapping. In *Proceedings of the international conference on medical image computing and computer-assisted intervention (MICCAI 2003)*. *Lecture notes in computer science*, Vol. 2217, pp. 1008–1016. New York: Springer.
11. Freedman D, Radke RJ, Zhang T, Jeong Y, Lovelock DM, Chen GT. 2005. Model-based segmentation of medical imagery by matching distributions. *IEEE Trans Med Imaging* **24**(3):281–292.
12. Rosenblatt F. 1956. Remarks on some nonparametric estimates of a density function. *Ann Math Stat* **27**:832–837.
13. Silverman BW. 1992. *Density estimation for statistics and data analysis*. London: Chapman and Hall.
14. Paragios N, Deriche R. 2002. Geodesic active regions and level set methods for supervised texture segmentation. *Int J Comput Vision* **46**(3):223–247.
15. Chan LA Vese TF. 2001. Active contours without edges. *IEEE Trans Med Imaging*, **10**(2):266–277.
16. Rousson, M., Cremers, D., 2005. Efficient Kernel Density Estimation of Shape and Intensity Priors for Level Set Segmentation, *International conference on medical image computing and computer-assisted intervention (MICCAI 2005)*, **2**: 757–764.