Image Segmentation with a Statistical Appearance Model and a Generic Mumford-Shah Inspired Outside Model

Thomas Albrecht and Thomas Vetter

University of Basel

Abstract. We present a novel statistical-model-based segmentation algorithm that addresses a recurrent problem in appearance model fitting and model-based segmentation: the "shrinking problem". When statistical appearance models are fitted to an image in order to segment an object, they have the tendency not to cover the full object, leaving a gap between the real and the detected boundary. This is due to the fact that the cost function for fitting the model is evaluated only on the inside of the object and the gap at the boundary is not detected. The state-of-the-art approach to overcome this shrinking problem is to detect the object edges in the image and force the model to adhere to these edges. Here, we introduce a region-based approach motivated by the Mumford-Shah functional that does not require the detection of edges. In addition to the appearance model, we define a generic model estimated from the input image for the outside of the appearance model. Shrinking is prevented because a misaligned boundary would create a large discrepancy between the image and the inside/outside model. The method is independent of the dimensionality of the image. We apply it to 3-dimensional CT images.

1 Introduction

We present a novel statistical-model-based segmentation algorithm that addresses a recurrent problem in appearance model fitting and model-based segmentation: the "shrinking problem", see the "Examples of Failure" in [1] for instance. When statistical appearance models are fitted to an image in order to segment an object, they have the tendency not to cover the full object, leaving a gap between the real and the detected boundary. The model seems to shrink inside the real object, a typical example can be seen in Figure 3b. This is due to the fact that the cost function for fitting the model is evaluated only on the inside of the object and the gap at the boundary is not detected. The state-of-the-art approach to overcome this shrinking problem is to detect the object edges in the image and force the model to adhere to these edges [2]. While this can in fact prevent shrinking, it requires the accurate detection of the object boundary. But in many applications the boundary detection can be almost as difficult as the original segmentation task.

Here, we introduce a region-based method that aims at solving the shrinking problem without the need to explicitly detect edges. The idea is borrowed from the Mumford-Shah functional for image segmentation [3], and is widely used in the field of level set segmentation [4]: Instead of detecting edges, we try to partition the image into different regions. The edges are only implicitly determined as the boundary between these regions, see [5] for an illustrative explanation of this principle. While the original Mumford-Shah method seeks any regions that offer an optimal piecewise approximation of the image, in our case of appearance model fitting, we have a very strong preconception into what regions we wish to partition the image: The foreground object, which is an instance of our statistical appearance model, and the background, the area around the model. In this way, we combine some of the advantages of Mumford-Shah based level set segmentation and appearance model fitting.

Statistical appearance models are built from example data sets and model the shape and appearance of a specific object class. Typical object classes in the literature are faces, organs or bones. In some cases, the example data sets may offer representative data not only for the inside of the object but also for the background. In these cases, the background can be modeled in a similar way to the foreground, see [6], or the model can simply be enlarged to include some of the background information. This problem can be regarded as solved.

We on the other hand focus on cases where the example data sets do not provide representative data for the background. Even though it would be desirable to develop a complete model of an object and all possible backgrounds and adjacent objects, it is easier and often the only realistic option to focus on one object of interest at a time. Our main motivation is a femur bone model we developed from CT scans of isolated bones. In the scans, the bones are surrounded only by air, but in most practical applications, bones will be surrounded by soft tissue, adjacent bones, etc., see Figure 1. A similar situation is found in face modeling, where a model has to be fitted to faces without any prior knowledge about the background.

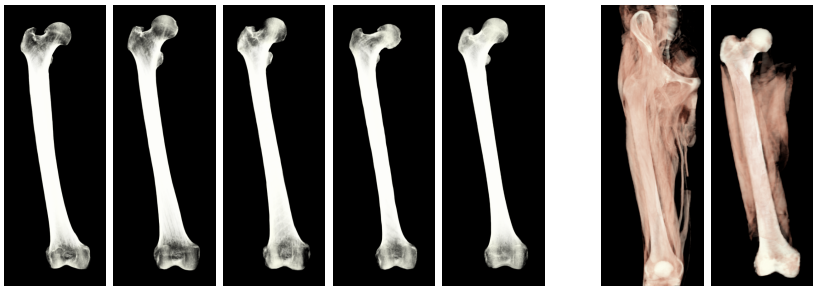

Fig. 1: On the left: a subset of the bones used to build the appearance model. The model is built from isolated bones. No useful information is available for the outside of the bones that could be used in real segmentation tasks such as segmenting the femur in the images on the right.

With no usable example data on the outside of the object, we look again to the Mumford-Shah functional [3], which is generic in the sense that it does not model regions based on examples. Instead, the appearance for the different regions is estimated form the input image itself. The original formulation of the functional proposes two possible ways to estimate the appearance for a given region from the intensity values: the mean intensity or a smoothed version of the image on that region. More sophisticated models have later been developed, see [4]. However, introducing and evaluating the different modeling techniques for different applications is beyond the scope of this paper and we introduce our method based on the simple models from [3].

Prior Work Image segmentation remains one of the most important challenges in image analysis. When the objects of interest like for instance bones, organs or faces, cannot be identified by simple intensity thresholding, the most successful and intensively researched approach is to include prior knowledge in the form of a statistical shape model into the segmentation algorithms and allow only shapes than can be represented by this model as segmentation results. We can distinguish between algorithms with strict and those with relaxed shape constraints. The algorithm we propose here enforces a strict shape constraint, which means that the segmentation results have to strictly lie in the space of shapes modeled by the statistical model. Essentially, such a segmentation can also be regarded as a model fitting algorithm, as it optimizes only the pose and model parameters. In principle the strict shape constraint could be relaxed by one of the methods proposed in [6], but we have not yet implemented these methods.

There are two main segmentation frameworks that are commonly used as the basis for shape-model-constrained segmentation and our method combines features from both of them. The first is based on the active shape and appearance models as proposed by Cootes, Taylor et al. [1] or Blanz and Vetter [7]. The second framework is that of level set segmentation and is mostly based on the Mumford-Shah functional [3] and its level set formulation [5]. Instead of trying to summarize the extensive research performed in this area, we refer to the comprehensive reviews by Heimann et al. [6] and Cremers et al. [4]. Both frameworks share the distinction between edge- and region-based segmentation. Region-based methods have proven to be more robust and successful as they do not only consider local edge information, but rather complete regions like the inside and the outside of the segmented object, see [4, 6].

Shape and Region Modelling Conceptually, the main difference between the level set and active appearance model frameworks is the representation of the shapes. Active shape and appearance models represent shapes by discrete point sets or grids, while level set methods represent shapes by implicit (= level set) functions. Consequently, for including prior knowledge about an object class of shapes into the segmentation algorithm, appropriate statistical models have been proposed for each framework. Active shape and appearance models represent the class of shapes by deformations of a reference. New shapes are generated by linear

combinations of example deformations. In order to determine these example deformations, the example data sets have to be brought into correspondence with a non-rigid registration algorithm. In level set methods, a class of shapes is modeled by linear combination of example level set functions. For this, the shapes do not need to be in correspondence but only rigidly aligned.

Further, the different shape model representations determine a different treatment of the image regions within the model. With correspondence information available, it is easy to transfer the image intensity or "appearance" of each example to the reference and build a separate linear model of appearances. For level set based models without correspondence information, such a straight-forward appearance modeling is not possible. Therefore, level set methods typically use intensity models estimated from the input image or histogram-based statistic that do not require correspondence information.

If we wish to use the statistical appearance information from the example data sets for the inside of the model and a generic model estimated from the input image on the outside, we have two possibilities: 1. Find a way to integrate a correspondence-based intensity model in the level set framework. 2. Integrate a generic input-image-based outside model into the active appearance model fitting. While the first possibility is an interesting research topic and may be the subject of a future paper, we take the second approach here. In this way the appearance model can be used in its original form and only needs to be complemented by an outside model.

2 Segmentation Method

In this section, we give a more detailed description of the models and show how they can be combined. Then, we show the feasibility of our approach on a few qualitative results. A thorough comparison with state of the art segmentation methods would require the implementation of many edge- and region-based segmentation methods, which is beyond the scope of this paper.

2.1 Inside Shape and Appearance Model

To build the inside appearance model, we need to acquire a representative collection of example data sets of the organ we wish to model. In our experiments, we acquired $n = 47$ CT data sets of isolated human femur bones. After a rigid pre-alignment, these are brought into correspondence with a non-rigid image registration method [8]. We single out one of the data sets as the reference and register all n data sets to this reference. This introduces a bias in the model towards the reference, but the interested reader can find strategies to remove or reduce this bias in [6]. Once the data sets are registered, a statistical shape and appearance model can be built along the lines of those proposed by Blanz and Vetter or Cootes and Taylor [7, 1]. The registration algorithm produces n deformation fields $u_i : \Omega \to \mathbb{R}^d$ defined on the reference's image domain $\Omega \subset \mathbb{R}^d$, where the dimensionality d is of course typically 2 or 3. When we denote by

$\Gamma \subset \Omega$ the points on the inside of the reference shape, the inside of the target shapes is represented as

$$\Gamma_i = \{x + u_i(x) \,|\, x \in \Gamma\}, \tag{1}$$

i.e. Γ_i is a forward warp of Γ with $x+u_i(x)$. On the other hand, we can backward-warp the CT intensities ct_i of the targets to the reference as:

$$\widetilde{ct}_i(x) = ct_i(x + u(x)) \text{ for } x \in \Gamma_i. \tag{2}$$

In this way, the shape and the intensity information of all examples is available on the reference. In practice Γ is a finite set with $m := |\Gamma|$ elements, and each shape can be represented by a dm-dimensional vector $\mathbf{s}_i$ of coordinates and each appearance by an m-dimensional vector $\mathbf{t}_i$ of intensity values. From these, we can calculate mean vectors $\bar{\mathbf{s}}$ and $\bar{\mathbf{t}}$ and covariance matrices $\Sigma_s = \frac{1}{n}\mathbf{X}_s\mathbf{X}_s^T$, $\Sigma_t = \frac{1}{n}\mathbf{X}_t\mathbf{X}_t^T$, where $\mathbf{X}_s, \mathbf{X}_t$ are the mean free data matrices with columns $\mathbf{s}_i - \bar{\mathbf{s}}$ resp. $\mathbf{t}_i - \bar{\mathbf{t}}$. The actual statistical modeling consists of assuming multivariate normal distributions $\mathcal{N}(\bar{\mathbf{s}}, \Sigma_s)$, $\mathcal{N}(\bar{\mathbf{t}}, \Sigma_t)$ for the shape and intensity data. After a singular value decomposition of the data matrices:

$$\tfrac{1}{\sqrt{n}}\mathbf{X}_s = \mathbf{U}_s\mathbf{W}_s\mathbf{V}_s^T \text{ resp. } \tfrac{1}{\sqrt{n}}\mathbf{X}_t = \mathbf{U}_t\mathbf{W}_t\mathbf{V}_t^T, \tag{3}$$

we can represent the shapes and intensities of the statistical model as:

$$\mathbf{s}(\alpha) = \bar{\mathbf{s}} + \mathbf{U}_s\mathbf{W}_s\,\alpha =: \bar{\mathbf{s}} + \mathbf{Q}_s\,\alpha, \text{ and } \mathbf{t}(\beta) = \bar{\mathbf{t}} + \mathbf{U}_t\mathbf{W}_t\,\beta =: \bar{\mathbf{t}} + \mathbf{Q}_t\,\beta. \tag{4}$$

where α and β are coefficient vectors. Under the assumption of the above normal distributions, α and β are distributed according to $\mathcal{N}(0, \mathbf{I})$. While the 3D Morphable Model [7] was originally only defined *on* and not on the inside of the modeled 3D object and the active appearance model [1, 6] was originally only introduced for 2D shapes, this is a straight-forward extension of these models. Such models are usually called PCA models as Equations (3) and (4) constitute a principal component analysis of the data matrices.

Segmentation with this statistical model is now performed by finding those coefficients α, β for which the difference between the shape and appearance associated with the vectors $\mathbf{s}(\alpha), \mathbf{t}(\beta)$ and the target object in the input image $\mathcal{I}(x)$ is minimal. In addition to the shape and appearance, we also need to estimate the pose of the object in the image, which can be represented by a rigid or similarity transform T_ρ with parameters ρ. The segmentation can be formulated as a minimization problem. For better readability we treat the vectors $\mathbf{s}, \mathbf{t}$ as continuous functions and write the problem as an integral:

$$E(\alpha, \beta, \rho) = \int_\Gamma \Big(\underbrace{(\bar{\mathbf{t}} + \mathbf{U}_t\beta)(x)}_{\text{model intensity}} - \underbrace{\mathcal{I}\big(T_\rho\,(\bar{\mathbf{s}} + \mathbf{U}_s\alpha)(x)\big)}_{\text{image intensity at model point}} \Big)^2 dx + \eta_s\|\alpha\|^2 + \eta_t\|\beta\|^2. \tag{5}$$

The norms of α and β with weighting terms η_t, η_s act as regularization terms motivated by the normal distributions $\mathcal{N}(0, \mathbf{I})$ of α and β. The optimal parameters α, β, ρ can be sought with any standard optimization algorithm, and in this

fashion the shape, appearance and position of any object that can be represented by the statistical model can be identified. However, only points on the inside of the model are considered and all points on the outside of the model are ignored, which can lead to the adverse effect of "shrinking" described in the introduction.

2.2 Mumford Shah Model

In their landmark paper [3], Mumford and Shah introduced what is now known as the Mumford-Shah functional for image segmentation, which seeks to simultaneously find an edge set C and a piecewise smooth approximation $\mathcal{J}$ of an input image $\mathcal{I} : \Omega \to \mathbb{R}$. In [5] Chan and Vese proposed a simplified version of this functional for the case that C is a closed contour (represented by a level set function) that separates the image domain Ω into an inside and an outside, $\text{in}(C)$ and $\text{out}(C)$ of C. In this case, the Mumford-Shah functional can be written as:

$$F(C, \mathcal{J}) = \lambda \int_{\text{in}(C)} (\mathcal{J}_{\text{in}} - \mathcal{I})^2 + \lambda \int_{\text{out}(C)} (\mathcal{J}_{\text{out}} - \mathcal{I})^2 + \mu \ \text{length}(C) + \nu \int_{\Omega \setminus C} |\nabla \mathcal{J}|^2 \ , \quad (6)$$

where $\text{length}(C)$ denotes the length of the segment boundary C and acts as a regularization term. Typically, the functional is minimized with an interlaced algorithm. In every other iteration the boundary C is kept fixed and the image approximation $\mathcal{J}$ is optimized and in the next iteration $\mathcal{J}$ is fixed and C optimized. Mumford and Shah showed that if C is fixed, $\mathcal{J}$ optimizes the functional if and only if it satisfies the following elliptic boundary value problem with zero Neumann boundary conditions on each of the segments, here written out only for $\text{out}(C)$:

$$-\Delta \, \mathcal{J}_{\text{out}} = \frac{\lambda}{\nu} \left(\mathcal{I} - \mathcal{J}_{\text{out}} \right) \text{ on } \text{out}(C) \qquad \frac{\partial \mathcal{J}_{\text{out}}}{\partial n} = 0 \text{ on } \partial(\text{out}(C)). \quad (7)$$

This means that $\mathcal{J}$ has to be a smoothed version of $\mathcal{I}$ with sharp edges on the boundary C, which is why the functional is minimal when C coincides with edges in the image, while on the segments the image can be approximated well by smooth functions. The great advantage over methods based on actual edge detection is that when no sharp edges are present in the image, the minimizing edge set C will still separate the different regions in the image in an optimal way when $F(C, \mathcal{J})$ is minimized. If $\frac{\lambda}{\nu} \to 0$, the optimal approximation $\mathcal{J}$ is a piecewise constant function which takes on the mean value of the function $\mathcal{I}$ on each of the segments, i.e. $\mathcal{J}_{\text{out}} \equiv c_{\text{out}} = \frac{1}{|\text{out}(C)|} \int_{\text{out}(C)} \mathcal{I}$. More sophisticated approximation strategies for $\mathcal{J}_{\text{in,out}}$, e.g. based on texture can be found in [4].

This segmentation method separates those two regions which can be best approximated by mean intensities or smooth approximations. However, it is by no means guaranteed that these coincide with the organs we want to segment in the image.

2.3 Combining the Models

We now present a way to combine the prior knowledge of the statistical shape and appearance model and the generic ad-hoc modeling technique of the Mumford-Shah segmentation method. For the inside of the object, we use the appearance model exactly as described in Section 2.1. The outside Mumford-Shah model is derived from (6) with a few adjustments. First of all, we only use the terms concerning the outside region. The length term can be omitted as the regularization properties of the statistical model provide a superior regularization method.

The terms in Equation (6) are defined on a part of the input image domain, whereas Equation (5) is defined on a part of the reference domain. To seamlessly integrate the outside terms into the appearance model segmentation, we need to transform them to the reference domain. In Equation (5), the spatial transformation from the reference model to the image is given by $\Phi_{\alpha,\rho}(x) := T_\rho\left(\bar{\mathbf{s}}+\mathbf{Q}_s\alpha\right)(x)$. and thus the transformation ("change of variables") of the outside terms by:

$$\begin{aligned}
&\lambda \int_{\mathrm{out}(C)} (\mathcal{J}_{\mathrm{out}}-\mathcal{I})^2+\nu \int_{\mathrm{out}(C)} |\nabla \mathcal{J}_{\mathrm{out}}|^2 = \lambda \int_{\Phi_{\alpha,\rho}(\Gamma_{\mathrm{out}})} (\mathcal{J}_{\mathrm{out}}-\mathcal{I})^2+\nu \int_{\Phi_{\alpha,\rho}(\Gamma_{\mathrm{out}})} |\nabla \mathcal{J}_{\mathrm{out}}|^2 \\
&= \lambda \int_{\Gamma_{\mathrm{out}}} \left(\mathcal{J}_{\mathrm{out}} \circ \Phi_{\alpha,\rho} - \mathcal{I} \circ \Phi_{\alpha,\rho}\right)^2 |\det D\Phi_{\alpha,\rho}| + \nu \int_{\Gamma_{\mathrm{out}}} |\nabla \mathcal{J}_{\mathrm{out}} \circ \Phi_{\alpha,\rho}|^2 \, |\det D\Phi_{\alpha,\rho}|,
\end{aligned} \tag{8}$$

where Γ_{out} is the outside of the model in the reference domain. In principal, Γ_{out} should be chosen so that $\Phi_{\alpha,\rho}(\Gamma_{\mathrm{out}}) = \mathrm{out}(C)$, but in practice, any neighborhood of Γ can be used. Then, contrary to the original integral from the Mumford-Shah functional, the transformed integral does not depend on the function or parameters we wish to optimize, which greatly simplifies the minimization. The only dependence remains in the determinant term from the transformation formula $|\det D\Phi_{\alpha,\rho}|$. However, this is where we introduce a simplifying approximation and assume $|\det D\Phi_{\alpha,\rho}| \equiv 1$, as it would be very time-consuming to compute the derivative of the deformations caused by the matrix $\mathbf{Q}_s$. Secondly, this term measures the volume change caused by $\Phi_{\alpha,\rho}$ and would allow the minimization of the functional simply by contracting the model, which is not desirable. Our proposed combined cost function is then given as:

$$\begin{aligned}
G(\alpha,\beta,\rho) = &\int_{\Gamma} \left((\bar{\mathbf{t}} + \mathbf{Q}_t\beta)(x) - \mathcal{I} \circ \Phi_{\alpha,\rho}(x)\right)^2 dx + \eta_s\|\alpha\|^2 + \eta_t\|\beta\|^2 \\
&+ \lambda \int_{\Gamma_{\mathrm{out}}} \left(\mathcal{J}_{\mathrm{out}} \circ \Phi_{\alpha,\rho} - \mathcal{I} \circ \Phi_{\alpha,\rho}\right)^2 + \nu \int_{\Gamma_{\mathrm{out}}} |\nabla \mathcal{J}_{\mathrm{out}} \circ \Phi_{\alpha,\rho}|^2.
\end{aligned} \tag{9}$$

The principal component matrix $\mathbf{Q}_s$ and mean vector $\bar{\mathbf{s}}$ used in $\Phi_{\alpha,\rho}$ have been defined only for the inside model Γ in Section 2.1. They need to be extended to the outside in order to calculate the outside terms of Equation (9). If the deformation fields u_i from which the model is calculated are defined on the entire

image domain of the reference, which is the case for the registration algorithm we used, this extension can be performed in a straight-forward manner. The mean vector $\bar{\mathbf{s}}$ is naturally extended by the mean of the fields on Γ_{out}. The matrix $\mathbf{Q}_s$ can be extended by employing the same linear combination of the original deformation fields on the outside as on the inside. The linear combinations are stored in the matrix $\mathbf{V}_s$ from the singular value decomposition Equation (3), and we can compute the extension of $\mathbf{Q}_s$ to Γ_{out} as $\mathbf{Q}_s = \frac{1}{\sqrt{n}}\mathbf{X}_s\mathbf{V}_s$.

2.4 Implementation

The minimization of the functional G defined in Equation (9) is handled in an interlaced algorithm similar to that described in Section 2.2: We alternately calculate the ad-hoc model $\mathcal{J}_{\text{out}}$ for the current parameters α, β, ρ, and find the parameters for the next iteration step with a standard optimization algorithm; we use the LBFGS optimizer [9]. $\mathcal{J}_{\text{out}}$ needs to be calculated from the image intensities as the mean or according to the elliptic equation (7) (on Γ_{out} instead of $\text{out}(C)$). Like the inside Γ, we represent Γ_{out} by an unstructured grid, and implemented a Gaussian smoothing with Neumann zero boundary conditions on this grid to approximate the solution of Equation (7). A more exact solution could be achieved by computing a finite element solution on this grid.

2.5 Results

Fig. 2: On the left: A CT slice with its approximation by the inside and outside model. The inside is an instance of the statistical model, while the outside is modeled as a smoothed version of the image intensities. On the right, the outside is modeled as the mean value of the outside intensities, which works best for uniform outside intensity.

We conclude by showing a few examples of bone segmentations that show the feasibility of segmentation with our proposed combined method and its advantages over the individual methods of level set and appearance model segmentation. As we are using a strict shape constraint, none of the segmentation results are perfect. They are only the best approximation within the space of the shape model that the optimization algorithm was able to find. We used the LBFGS algorithm with a landmark-based rigid alignment of the mean model as initialization. The method is not very sensitive to the parameters. For all experiments, we have chosen $\lambda = 1$, $\eta_s = 100$, $\eta_t = 10$. For the Gaussian smoothing of the outside model, we have used a variance that corresponds to $\nu = 300$.

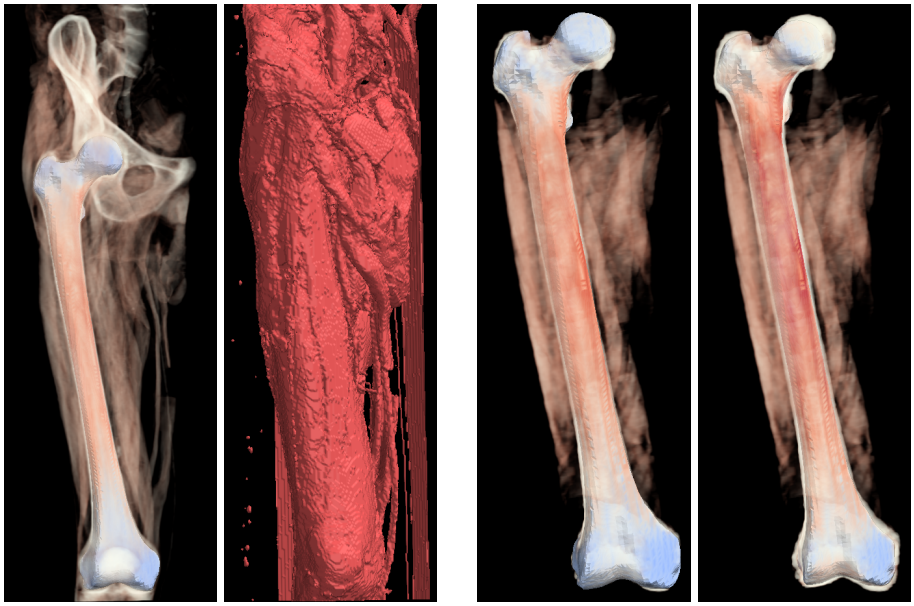

(a) (b)

Fig. 3: Comparison of the our segmentation method with each of the original methods. The input images are shown in Figure 1. On the left in (a), the proposed method identifies to femur bone, whereas the original Mumford-Shah level set segmentation on the right separates air from not-air, and the segmentation boundary shows the muscle tissue and not the femur bone.
On the left in (b), the proposed method identifies the femur as well as the model permits, while on the right the appearance model without outside model "shrinks" and leaves a small gap between the model and the real bone surface.

Figure 2 illustrates the two proposed method for outside models. Only when the outside of the bone is very uniform as for instance in the case of isolated bones is the constant approximation by the mean preferable over the smooth approximation. Note that the aim of the outside model is not the perfect representation of the input image, that would of course be given by the unsmoothed image itself. The aim is to give a homogeneous representation of the outside which encourages the correct placement of the model boundary because any other placement would incur a higher cost in the functional from Equation (9).

In Figure 3a we see how our method can identify the femur in a CT image with soft tissue and other bones. In contrast, the Mumford-Shah level set segmentation finds the most prominent segment boundary in the image, that between air and everything that is not air. While this is the optimal boundary from the point of view of this segmentation method, it is not the boundary we are interested in if we wish to segment femur bones. In Figure 3b, we see another successful segmentation with our combined model, contrasted with a

result of using only the inside appearance model. As expected, in this case, the segmentation leaves a narrow gap around the boundary of the model.

3 Discussion

We have showed that including an outside model term motivated by the Mumford-Shah functional can help reduce the effect of "shrinking" in active appearance model fitting/segmentation. Without the need of explicit edge detection, the outside model discourages the incorrect placement of the boundary. Obviously, this works best in regions where the foreground and background have distinct intensity values, but the correct separation of fore- and background should always be at least as good as without the outside model. We did not yet perform a quantitative comparison with state-of-the-art segmentation or model fitting algorithms. In fact, we think that a fine-tuned edge-based method may perform equally well or even better if the correct edges can be found. But we have shown that a region-based combined model can improve model-based segmentation while completely circumventing the difficult and often unstable problem of edge detection.

Possible future work includes a thorough quantitative comparison with state-of-the-art methods, the evaluation of more advanced outside region models, and ways to relax the strict shape constraint.

References

1. Cootes, T., Taylor, C.: Statistical models of appearance for medical image analysis and computer vision. In: Proc. SPIE Medical Imaging. Volume 4322. (2001) 236–248
2. Romdhani, S., Vetter, T.: Estimating 3d shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2005. CVPR 2005. Volume 2. (2005)
3. Mumford, D., Shah, J.: Optimal Approximations by Piecewise Smooth Functions and Associated Variational Problems. Center for Intelligent Control Systems (1988)
4. Cremers, D., Rousson, M., Deriche, R.: A Review of Statistical Approaches to Level Set Segmentation: Integrating Color, Texture, Motion and Shape. International Journal of Computer Vision **72**(2) (2007) 195–215
5. Chan, T.F., Vese, L.A.: Active contours without edges. IEEE Trans. Image Process. **10**(2) (2001) 266–277
6. Heimann, T., Meinzer, H.: Statistical shape models for 3D medical image segmentation: A review. Medical Image Analysis (2009)
7. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques, ACM Press (1999) 187–194
8. Dedner, A., Lüthi, M., Albrecht, T., Vetter, T.: Curvature guided level set registration using adaptive finite elements. In: Pattern Recognition. (2007) 527–536
9. Zhu, C., Byrd, R., Lu, P., Nocedal, J.: L-BFGS-B: Fortran subroutines for large-scale bound constrained optimization. ACM Transactions on Mathetmatical Software **23**(4) (1997) 550–560