

Marcel Lüthi<sup>1</sup> · Anita Lerch<sup>1</sup> · Thomas Albrecht<sup>1</sup> · Zdzislaw Krol<sup>2</sup> ·  
Thomas Vetter<sup>1</sup>

# A hierarchical, multi-resolution approach for model-based skull-segmentation in MRI volumes.

**Abstract** We present a model-based approach for segmentation of the skull from T1 weighted MR images of the human head. Segmentation is performed by fitting a morphable skull model into a pre-segmented version of the image. This yields a segmentation result that is constrained to the normal skull anatomy and thus gives a statistically meaningful approximation of the skull, even in places where the bony structure cannot be distinguished from the surrounding tissue. We propose a multi-resolution approach to model-fitting that leads to an improved convergence rate. Further, we show how a hierarchy of statistical models can be used to increase the flexibility of the statistical model and thus to obtain more accurate segmentation results. To validate our method, we present experimental result using a statistical model based on 17 example skulls acquired from segmented CT-images. The results show that a good approximation to the skull structure can be found. Our experiments confirm that by incorporating strong prior knowledge even such difficult segmentation tasks as skull-segmentation from MR images become feasible.

**Keywords** Segmentation · Statistical Models · MRI · Multi-resolution · Model fitting

---

1.  
University of Basel  
Department of Computer Science  
Bernoullistrasse 16  
CH-4056 Basel  
E-mail: marcel.luethi@unibas.ch  
E-mail: anita.lerch@unibas.ch  
E-mail: thomas.albrecht@unibas.ch  
E-mail: thomas.vetter@unibas.ch

2.  
Hightech Research Center  
University of Basel  
Schanzenstrasse 46  
CH-4056 Basel  
E-mail: zkrol@uh-basel.ch

---

## 1 Introduction

Models of the human anatomy play an increasingly important role in understanding the human physiology and they start to become a prerequisite for many medical applications. For instance, cranio-facial interventions require the detailed planning of the surgery based on a model of the patient specific anatomy. Modern planning is usually performed using a combination of different diagnostic data-sets. While the different soft-tissues and vessels are relatively easy to segmented from MR images, the bony structure is hardly visible. Therefore, in practical surgery planning an additional CT image is acquired, which exposes the patient to additional harmful radiation. In this article we present a method for segmentation of the skull directly from T1 weighted MR images, using a morphable model (i.e. a statistical shape model) to estimate the skull-structure where it is not visible in the image.

Magnet resonance imaging is commonly applied in clinical practice to help diagnose or monitor a treatment. In contrast to computer tomography, the image acquisition poses almost no risk to the patient and should therefore be preferred over CT whenever possible. Unfortunately, the segmentation of bones in MR images is a difficult problem, as bony structure is hard to distinguish from the surrounding tissue and virtually impossible to distinguish from air in the images. The low resolution of a MR image and the occurrence of very fine bony structure in the human skull, makes skull-segmentation from MR images a particularly challenging problem. To overcome these difficulties, we exploit prior knowledge about the shape of normal skulls in the form of a statistical shape model. The statistical model is built from a number of skull-surfaces, segmented from CT data-sets. Hence, our method uses knowledge learned from CT data to restrict the solution space in the more difficult problem of segmenting MR images.

Our method proceeds in two steps. In a first step, we perform a pre-segmentation of the skull by applying simple thresholding and morphological operations to

the image. This yields a rough outline of the skull. In a second stage we fit a statistical shape model to the pre-segmented image. The fitting is done in a hierarchical way. First, a statistical model of the complete skull is fitted into the image. This yields a good initial estimate of the skull structure and allows to localize the individual anatomical structures, such as the cranium and the mandible. Then, a separate statistical model is fitted for each individual part, yielding a more accurate result. The fitted model serves as an approximation to the actual skull structure and provides complete information even in places where the bony structures are not visible in the MR image.

The contributions of this work are the formulation of a novel, multi-resolution fitting procedure that leads to improved convergence and allows to increase the flexibility by fitting a hierarchy of statistical models. Further, we present a new approach to skull-segmentation from MR images, by using a combination of a simple intensity based algorithm and model fitting.

*Related work and outline.* Automatic segmentation of medical images is an extremely well researched problem and the need for incorporating statistical information to constrain the segmentation process has clearly been recognized [4]. The use of statistical shape models for identifying structures in (2D) medical images has already been proposed by Cootes et al. [3] in 1993 and since used for segmentation of various structures from 3D CT and MR images [12,13,15,11,18]. All these methods have been evaluated on rather simple anatomical structures (such as the femur or pelvic bone). However, no results were shown for more complex structures such as the human skull. The literature on skull segmentation is very sparse in general. Results for the segmentation of the skull from CT images are shown by Kang et al. [10], where, in a multi-step approach, a sequence of standard segmentation techniques is applied to extract the bony structure. Dogas et al. [7] use techniques from mathematical morphology to segment the skull from MRI images. Hifai et al. [16] use a level-set segmentation technique to deform the contour of the scalp to fit the skull structure.

The paper is organized as follows. In section 2 we present the mathematical basics of statistical models and the methods we applied for building our skull-model. Section 3 discusses the algorithms we use for pre-segmentation of the MR images. The fitting procedure is explained in section 4. We present experimental results in section 5 followed by a discussion of our method in section 6.

## 2 Statistical Skull Model

Most segmentation algorithms are defined in a generic way, in the sense that these algorithm would work on a grey-scale image of any anatomical structure. While



**Fig. 1** The mean skull and its first two modes of variation. Top to bottom: 1st Mode of variation ( $\mu \pm 2\sigma$ ) Left to right: 2nd model of variation ( $\mu \pm 2\sigma$ )

these algorithms are flexible, they do not exploit all the available information. For difficult segmentation problems, such as skull segmentation from MR images, we need to incorporate strong prior knowledge into our algorithm. A powerful and versatile approach for incorporating prior knowledge is by using a statistical shape model.

The main idea behind a statistical shape model is to span a space of shapes (3D Surfaces) from a set of normal examples by taking linear combinations of these examples [1]. A probability distribution on this space is defined, which quantifies the probability of observing a particular linear combination. Usually this is done by fitting a normal distribution to the example data. In many image processing tasks, and in particular in image segmentation, the use of a statistical model makes it possible to constrain the search space to normal instances of the given shape.

Figure 1 shows an example of the statistical shape model of the human skull, used throughout this paper. It shows the mean skull, as well as the first two principal modes of variations.

We now present the details of building the statistical shape model of the human skull. Let  $\{\Gamma_i | \Gamma_i \subset \mathbb{R}^3\}_{i=1}^n$  be  $n$  surfaces of normal skulls. Define an arbitrary surface, say  $\Gamma_1$ , as the reference surface, and assume that  $\hat{\Gamma}_1$  is a suitable discretization of  $N$  points (for instance,  $\hat{\Gamma}_1$  is represented as a triangle mesh). We represent the reference surface as a shape vector  $x_1 \in \mathbb{R}^{3N}$ ,

$$x_1 = (v_x^1, v_y^1, v_z^1, \dots, v_x^N, v_y^N, v_z^N)^T$$

where the vector  $v^i = (v_x^i, v_y^i, v_z^i)$  represents the  $x, y, z$  coordinates of the  $i$ -th vertex of  $\hat{\Gamma}_1$ . Using the non-rigid

registration algorithm proposed by Dedner et al. [5], we obtain a vector field  $\phi \subset \mathbb{R}^3 \rightarrow \mathbb{R}^3$  such that the surface  $\Gamma_i$  can be approximated by the shape vector

$$s_i = (v_x^1 + [\phi(v^1)]_x, v_y^1 + [\phi(v^1)]_y, v_z^1 + [\phi(v^1)]_z, \dots, v_x^N + [\phi(v^N)]_x, v_y^N + [\phi(v^N)]_y, v_z^N + [\phi(v^N)]_z)^T$$

such that corresponding components among all  $x_i$  represent corresponding points (e.g. the  $j$ -th component represents the top of the nasal bone in all the examples). The shape vectors  $\{s_i\}_{i=1}^n$  are observations of the (unknown) normal distribution  $\mathcal{N}(\mu, \Sigma)$ . The mean and covariance are estimated using the usual formulas

$$\mu = \frac{1}{n} \sum_{i=1}^n s_i \quad (1)$$

$$\Sigma = \frac{1}{n} \sum (s_i - \mu)(s_i - \mu)^T. \quad (2)$$

Note that in all practical situations  $n \ll N$  and hence  $\Sigma$  is of at most rank  $n$ . In these cases the probability density function does not exist, as to evaluate the density we need to be able to invert  $\Sigma$ . Using principal component analysis (PCA), we can, however, define a normal distribution of the subspace spanned by the  $n$  examples.

*Principal Component Analysis* Although PCA is a well known method, we will give a rather detailed derivation, in order to clarify notation and introduce the concepts used later in section 4. We start by defining the mean free data matrix  $X \in \mathbb{R}^{3N \times (n)}$  as

$$X = [s_1 - \mu \dots s_n - \mu],$$

with columns  $x_i := s_i - \mu$ . Let  $\tilde{n} = \text{rank}(X)$  be the column rank of  $X$ . Using the singular value decomposition (see e.g. Demmel [6]), we can write

$$X = UDV^T. \quad (3)$$

where  $U = [u_1 \dots u_{\tilde{n}}] \in \mathbb{R}^{3N \times \tilde{n}}$  is a matrix satisfying  $U^T U = I$ ,  $D = \text{diag}(\sigma_1, \dots, \sigma_{\tilde{n}}) \in \mathbb{R}^{\tilde{n} \times \tilde{n}}$  with  $\tilde{n}$  strictly positive diagonal entries  $\sigma_i$  and  $V^T = [v_1 \dots v_n]^T \in \mathbb{R}^{\tilde{n} \times n}$  is a matrix, satisfying  $V^T V = I$ . By choosing a  $3N$ -by- $(3N - \tilde{n})$  matrix  $\tilde{U}$  so that  $[U \tilde{U}]$  is square and orthogonal (for example by using the Gram-Schmidt process) we can write

$$\begin{aligned} \Sigma &= \frac{1}{n} \sum_{i=1}^n (s_i - \mu)(s_i - \mu)^T = \frac{1}{n} \sum_{i=1}^n x_i x_i^T = \frac{1}{n} X X^T \\ &= \frac{1}{n} U D V^T V D U^T = [U \tilde{U}] \begin{bmatrix} \frac{1}{n} D^2 & 0 \\ 0 & 0 \end{bmatrix} [U \tilde{U}]^T. \end{aligned} \quad (4)$$

Equation (4) is an eigendecomposition of  $\Sigma$  and hence the columns of  $[U \tilde{U}]$  are the eigenvectors of  $\Sigma$ . From

equation 4 we see that, expressed in this basis of eigenvectors,  $\Sigma$  becomes a diagonal matrix. It further follows that the diagonal entry  $\sigma_i^2$  of  $D^2$  is the variance that is captured by the associated eigenvector  $u_i$ . We refer to the eigenvector  $u_i$  as the  $i$ -th *principal component*.

Defining the probability density is now straightforward. As all the variance in our data is captured by the  $\tilde{n}$  principal components, we represent our data with respect to this basis, and assume that the observations follow the normal distribution  $\mathcal{N}(U^T \mu, \frac{1}{n} D^2)$ .<sup>1</sup> The probability of observing the shape vector  $s \in \mathbb{R}^N$  is given by

$$p(s) = z \exp\left(-\frac{1}{2} U^T (s - \mu) \frac{1}{n} D^{2-1} U^T (s - \mu)^T\right). \quad (5)$$

where  $z = \frac{1}{(2\pi)^{(n-1)/2} |D|^{1/2}}$  is the normalization factor. We usually express the data as a linear combination of the  $u_i$ :

$$s = \mu + U\alpha$$

and call the coefficients  $\alpha$  the *PCA-coefficients*. In this representation the probability density has a particularly simple form, namely

$$p(s) = z \exp\left(\alpha \frac{1}{n} D^{2-1} \alpha^T\right) = z \exp\left(\frac{1}{n} \|D^{2-1} \alpha\|^2\right).$$

This expression states that the probability of observing a given shape vector  $s$  is proportional to the weighted norm of its pca-coefficients - a fact turns out to be useful in section 4.

It is important to stress at this point that any linear combination of the principal components  $u_i$  can be expressed as a linear combination of the example data and vice-versa. To see this, we right-multiply (3) by  $V D^{-1}$  to obtain an expression for the matrix of principal components  $U$  in terms of the data-matrix  $X$ :

$$U = X V D^{-1}. \quad (6)$$

Let  $\alpha = (\alpha_1, \dots, \alpha_{\tilde{n}})^T$ ,  $\beta = (\beta_1, \dots, \beta_n)$  be coefficient vectors. Then

$$\sum_{i=1}^{\tilde{n}} \alpha_i u_i = U\alpha = X V D^{-1} \alpha = X\beta = \sum_{i=1}^n \beta_i x_i.$$

Hence we have the relations


$$\beta = V D^{-1} \alpha \quad (7)$$

and

$$\alpha = D V^T \beta. \quad (8)$$

This fact allows us to keep the same linear combination of example vectors, while freely changing among different basis representations.

<sup>1</sup> This is in fact a rather strong assumption. It states that the shape-space is completely spanned by the  $n$  given examples, i.e. all the possible shapes can be represented as a simple linear combination of the given data. For small  $n$ , this is clearly not the case. This problem can be addressed by regularization (see e.g. Schäfer et al. [17]). This is, however, out of scope for this paper.



**Fig. 2** The pre-segmentation procedure. The original image (a) is masked with the brain and scalp mask (b) to obtain after thresholding the skull-mask (c).

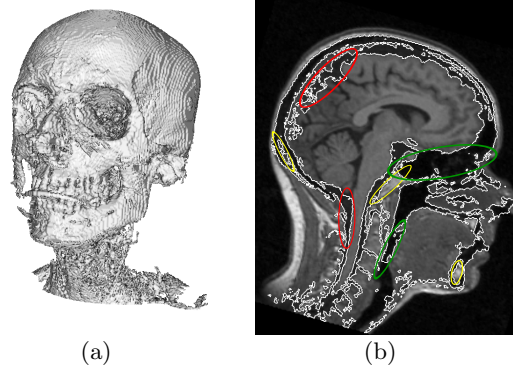
### 3 Pre-segmentation

The pre-segmentation step provides an automated procedure for finding a rough approximation of the skull structure, by means of simple image-processing operations. While segmenting bony structure directly from MR images is difficult, the different soft-tissue types can easily be distinguished. Fortunately, the brain and scalp, which are easy to segment, constrain the position of the skull rather well. We make use of this fact by segmenting the scalp and the brain and use them to mask the MRI-image. On the remaining region, we perform a threshold operation using the Otsu-threshold (the threshold value which maximizes the within class variance) and select the largest connected component. The different segmentation results are shown in figure 3.

*Scalp-segmentation* For scalp segmentation we are using the method proposed by Dodgas et al. [7]. The method works by using a combination of thresholding operations, to segment the soft-tissue from the bone, and mathematical morphology to close the auditorial channels and the nasal cavity. Performing a hole filling, leads to the scalp mask.

*Brain-segmentation* For the brain segmentation, we follow the method proposed by Géraud et al [8,16]. As for the scalp-segmentation, the method works by a combination of thresholding and mathematical morphology operation. Morphology is used to separate the brain from the surrounding structures. Selecting the largest component yields the final brain segmentation.

*Pre-segmentation results* The result from the brain and scalp segmentation, as well as the resulting mask for the bony parts, are shown in figure 3. It can be seen that the structure of the skull is roughly approximated by this method. However, the pre-segmentation fails at several places. For example in the area around the sinuses, thin bones and air are located side by side. The algorithm fails to distinguish them, as air and bones give almost the same signal. Also, CSF cannot be distinguished from bone by this simple method, as the intensity value of CSF in a T1 weighted image lies close to the one of bone. A



**Fig. 3** (a) shows a 3D contour of the pre-segmentation result. (b) While the outline of the skull is roughly correct, there are many places where the pre-segmentation (indicated by the white line) is wrong. The ellipses mark these wrong segmentation results. The green ellipses show areas where air is classified as bone, the yellow show the same for bone-marrow and the red ellipses for CSF.

different problem occurs with the bone marrow, which gives a distinct signal from bone and is therefore not excluded in the thresholding step.

### 4 Fitting

As discussed in section 2, the main assumption of our statistical model is, that given a large enough number  $n$  of examples, all shapes of the same class can be obtained by taking linear combinations of the example shapes. The shapes are most conveniently represented as linear combinations of the principal components  $u$  (cf. section 2):

$$s = \mu + \sum_{i=1}^n \alpha_i u_i.$$

Recall that the probability of observing  $s$  depends solely on the size of  $\|D^{2^{-1}} \alpha\|^2$ . These properties leads naturally to an approach for improving the result from the pre-segmentation step: We try to find a linear combination of the examples, such that the shape matches the one given from the pre-segmentation step well, but discourages unlikely combinations by penalizing large  $\alpha$ .

We proceed by stating this idea more formally. Let  $\tilde{I} : \Omega \subset \mathbb{R}^3 \rightarrow \{0, 1\}$  be the resulting binary image from the pre-segmentation step and  $\tilde{\Gamma} \subset \mathbb{R}^3$  be the surface that induces the partition in  $\tilde{I}$ . We define the distance image as

$$I(x) = \begin{cases} \text{dist}(x, \tilde{\Gamma}) & x \in \text{outside}(\tilde{\Gamma}) \\ 0 & x \in \tilde{\Gamma} \\ -\text{dist}(x, \tilde{\Gamma}) & x \in \text{inside}(\tilde{\Gamma}). \end{cases} \quad (9)$$

Finding a likely shape that fits the partition boundary of  $\tilde{I}$  amounts to finding transformation- and model parameters, such that the resulting surface comes to lie on

the zero-level set of  $I$ . This is satisfied by the solution of the optimization problem

$$\min_{s,t,R,\alpha} \sum_{i=1}^N w_i I(sR[\delta_i(\mu + \sum_{j=1}^n \alpha_j u_j)] + t)^2 + \lambda \|D^{2^{-1}} \alpha\|^2 \quad (10)$$


where  $s \in \mathbb{R}$  is a scaling factor,  $t \in \mathbb{R}^3$  a translation,  $R \in \mathbb{R}^{3 \times 3}$  a rotation matrix and  $\lambda$  a weighting coefficient. The function  $\delta_i(x) : \mathbb{R}^{3N} \rightarrow \mathbb{R}^3$  extracts the coordinates belonging to the  $i$ -th vertex from the shape vector. Finally,  $w_i \in \mathbb{R}$  is a weight for point  $i$ . For the current discussion, we just let all  $w_i = 1$ . Note that the first term is zero, if and only if the shape defined by the optimization parameters lies exactly on the zero-level set that defines the partition boundary  $\bar{I}$ . The second term in (10) serves as a regularizer and penalizes unlikely shapes. The problem (10) can be minimized by using any gradient based optimization scheme. Unfortunately, the dependence on  $I$  makes it highly non-convex and we are likely to get stuck in a local minimum.

*Multi-resolution* In image registration, a common technique to escape local minima and, at the same time to speed up the optimization procedure, is to use a multi-resolution scheme. We show how the same idea can be adapted for fitting a statistical model to an image. In a first step, we build an image pyramid of the distance image  $I$ . Let  $I^{(0)} = I$  be the original image. The image at level  $l$  is obtained by a convolution with a Gaussian kernel  $G(\sigma)$  of variance 1:

$$I^{(l)} = I^{(l-1)} \star G(1).$$

This results in smoothed version of the image and, consequently, in a smoother objective function (10) with fewer local minima. Additionally, the resolution of the  $l$ -th image can be reduced by the factor  $2^l$  in each dimension, which speeds up the optimization procedure considerably. Note however, that this simple approach does not work, as the zero-level set in the the images  $I_i, i > 0$ , does not represent the original surface anymore and we can not expect to find a good fit of the statistical model. Therefore, we need to apply the same smoothing to the reference surface  $\Gamma_0$ . We form the distance image of  $\Gamma_0$  and create the image pyramid, as described above. Using a contour extraction algorithm, such as for example marching cube [14], we extract the zero level set to obtain the reference surface  $\Gamma_0^{(l)}$  for level  $l$ . Figure 4 shows the four levels of the reference surfaces for our skull-model. Note, that only the most dominant features of the skulls are still represented in the surfaces. Repeating the procedure described in section 2 with  $\Gamma_0^{(l)}$  as the reference surface, we build a separate statistical model for each level.

As we built the statistical model separately for every resolution level, the PCA components may be different



**Fig. 4** The first four resolution levels of the reference skull. In the last resolution (right) the details are smoothed out and only dominant areas are still present in the surface.

---

#### Algorithm 1: The fitting algorithm

---

**Input:** Number-of-levels  $L$   
Reference-Mesh  $\Gamma_0$   
Model  $(\mu^{(l)}, U^{(l)}, V^{(l)}, D^{(l)}), l \in \{0, \dots, L\}$   
Distance-Image  $I^{(l)}, l \in \{0, \dots, L\}$   
**Output:**  $(s, R, t, \alpha)$

- 1  $(s^{(L+1)}, R^{(L+1)}, t^{(L+1)}, \alpha^{(L+1)}) := (1, \mathcal{I}, 0^3, 0^n)$
- 2  $D^{(L+1)} = V^{(L+1)} := \mathcal{I}$
- 3 **foreach**  $l \in [L, \dots, 0]$  **do**
- 4      $\bar{\alpha} := D^{(l+1)} V^{T^{(l+1)}} V^{(l)} D^{(l)-1} \alpha^{(l+1)}$ .
- 5     **SetInitialParameters**  $(s^{(l+1)}, R^{(l+1)}, t^{(l+1)}, \bar{\alpha})$
- 6      $(s^{(l)}, R^{(l)}, t^{(l)}, \alpha^{(l)}) :=$   

$$\min \sum_{i=1}^N w_i I^{(l)}(sR[\delta_i(\mu^{(l)} + \sum_{j=1}^n \alpha_j u_j^{(l)})] + t)^2 + \lambda \|D^{(l)2^{-1}} \alpha\|^2$$
- 7 **return**  $(s^{(0)}, R^{(0)}, t^{(0)}, \alpha^{(0)})$

---

and the coefficients  $\alpha$  do not represent the same shape anymore. We need to perform a change of basis from to the new principal components. Let  $\alpha^{(l)}$  be the final PCA coefficients vector obtained from optimizing (10) for level  $l$ . Recall from section 2, equation (7) and (8) the relation between the PCA-coefficients  $\alpha$  and the example-coefficients  $\beta$ . This relationship allows us to compute the coefficients for the next level from the previous solution. First, we compute from the PCA-coefficients the corresponding coefficients  $\beta$  that determine the linear combination of the examples  $x_1, \dots, x_n$

$$\beta^{(l)} = V^{(l)} D^{-1(l)} \alpha^{(l)}.$$

Then we use the same linear combination of the examples given by  $\beta^{(l)}$  as the initial solution for the next level. For this, we have to represent these coefficients in terms of the PCA-components  $U^{(l-1)}$  for level  $l-1$ . The coefficients for level  $l-1$  become

$$\alpha^{(l-1)} = D^{(l-1)} V^{T^{(l-1)}} \beta^{(l)} = D^{(l-1)} V^{T^{(l-1)}} V^{(l)} D^{-1(l)} \alpha^{(l)}.$$

We proceed iteratively until we reach the finest resolution. Algorithm 1 shows the detailed steps of the multi-resolution fitting.

**Algorithm 2:** Multi-part fitting algorithm

---

**Input:** Number of parts  $P$   
Reference-Mesh  $\Gamma_0$   
Model  $(\mu^{(p)}, U^{(p)}, V^{(p)}, D^{(p)}), p \in \{0, \dots, P\}$   
Distance-Image  $I$

**Output:**  $(s^{(p)}, R^{(p)}, t^{(p)}, \alpha^{(p)}), p \in \{0, \dots, P\}$

- 1  $(s^{(0)}, R^{(0)}, t^{(0)}, \alpha^{(0)}) = \text{Fit}(\Gamma_0)$
- 2 **foreach** part  $\Gamma_0^{(p)} \subset \Gamma_0$  **do**
- 3      $\bar{\alpha} := D^{(p)} V^{T^{(p)}} V^{(0)} D^{(0)-1} \alpha^{(0)}$
- 4     **SetInitialParameters**  $(s^{(0)}, R^{(0)}, t^{(0)}, \bar{\alpha})$
- 5      $(s^{(p)}, R^{(p)}, t^{(p)}, \alpha^{(p)}) = \text{Fit}(\Gamma_0^{(p)})$
- 6 **return**  $(s^{(p)}, R^{(p)}, t^{(p)}, \alpha^{(p)}), p \in \{0, \dots, P\}$

---

*Multi-parts fitting* When working with real data, especially with such complex structures as the skull, obtaining a large database of normal examples is difficult. Hence, our central assumption that any new skull we see can be accurately represented by our model is often violated. We therefore try to increase the flexibility of the model for the case when we are given only a small number of examples. This can be done by segmenting the skull into different parts and fit each part separately. A natural segmentation for the skull would, for instance, be to separate the mandible and the cranium. This separation offers several advantages. Given the same number of examples, it provides a greater flexibility, as we can obtain different optimal parameters for each part. Furthermore, as the structures to be fitted are simpler, fewer examples are needed to capture all the variation of this structure. Lastly, working with separate models allows us to use a different number of examples for each structure.

Our approach is as follows: In a first step, we fit a statistical model of the whole skull. This yields a good approximation of the overall structure and localizes the different parts. Using this result as an initial solution, we fit a separate statistical model for each part into the image. From the technical point of view, this hierarchical fitting does not pose additional challenges. We can employ the same method for changing among the different bases as described for the multi-resolution approach.<sup>2</sup> Algorithm 2 illustrates this procedure in more detail.

*Discretization of the reference surface* The fitting method (10) described above depends strongly on the discretization of the reference surface used to build the statistical model. The cost-function (10) is only evaluated at a discrete number of points on the surface. It is clear, that areas that are more densely sampled will have a larger influence on the overall error. The algorithm will strive to minimize the distance at these points, at the expense of a larger error at more sparsely sampled region. This

<sup>2</sup> For this procedure to work, we need to make sure that the example are processed in the same order. In case that the parts-model contains more basis vectors than the global model, we just set the corresponding coefficients to zero.

**Table 1** Parameter settings used in the experiments

Parameter	Value
Number of Levels	5
Number of Iterations	2000
Regularization parameter $\lambda$	0.1
Number of skulls in models	17

dependence on the discretization is almost always undesirable. Therefore, we represent the reference surface as a triangle mesh and weight the points according to the size of its neighborhood. More precisely, at the vertex  $v_i$ , we compute the area of the triangles that  $v_i$  belongs to. The weight  $w_i$  in (10) is then set proportional to this area. This greatly reduces the dependence on the sampling.

---


## 5 Results

In this section we present results that show the feasibility and benefits of our model based approach. In particular, we show that the skull structure is generally well approximated and good results are achieved even in areas where methods, which are based only on the intensity-values, must fail. For several reasons, we do not perform a comparison with other segmentation methods or try to give precise quantitative results at this stage. First, the skull model we are using at the moment is built from only 17 data-sets, some of which are of rather poor quality. Hence, any precise evaluation would show the limitations of the model rather than those of the presented method. Furthermore, at the current stage, where only small first steps towards skull-segmentation from MRI-data have been performed, we feel that a comparison of the different algorithm would not give much more insight. However, a quantitative evaluation and comparison is planned as future work.

*Experimental setup* We implemented our method using the Insight Registration and Segmentation Toolkit (itk) [9]. Optimizing the cost function (10) is performed using ITK's *RegularStepGradientDescentOptimizer*. The rotation is represented as a versor rotation. Before the skulls are fitted, they are roughly aligned with the reference surface using landmark based rigid registration [20]. Our statistical model has been built from 16 segmented skulls from diverse sources, out of which 7 stem from the Bosma collection [2]. The parameter settings used for all our experiments is given in table 1.

*Experiments and Results* In the first experiment, we tested if the multi-part fitting leads to improvements over the fitting result obtained by performing only a fit of the whole skull. For this, we segmented the skull into two parts, namely the mandible and the cranium. Figure 5 shows how fitting the parts separately increases the flexibility of the model. While for the cranium the improvement is small (the difference is mainly in the area of the






**Fig. 5** The multi-part fitting (c) clearly improves the result compared to the global fitting (b). At the places where the pre-segmentation (a) fails, the multi-parts fitting (c) is able to segment the structures correctly.

sinuses), the result for the mandible is clearly much better. This is due to the fact that the accumulated error for the cranium dominates the cost-function and hence the algorithm strives to decrease the error for the cranium. Due to the limited flexibility of the model, this comes at the cost of a larger error for the mandible. Abolishing this coupling leads to a better fit for both parts. It can also be clearly seen in figure 5 that at the places where the pre-segmentation algorithm fails, good results can be achieved by using model-based fitting. The missing information in the MR image is compensated by the prior knowledge about the skull structure encoded in the model.

Figure 5 shows the fitting result in six slices through a different MRI volume. Here as well, the model approximates the shape of the skull generally well. However, it can also be seen that in some places, for instance in the area of the frontal sinuses, the fitting deviates from the true skull contour. We attribute this to the limited number of examples in the model. To achieve a better fit, the model has to be enlarged significantly, or more



**Fig. 6** Slices through an MRI of the skull, showing the fitting result for the cranium (red) and mandible (orange).

independent parts would have to be considered. Another problem that becomes apparent in Figure 5 is that only structures that are in the model can be segmented. As the model has only 28 teeth, the 32 teeth in the image can never be correctly fitted. This suggests to extend the method, such that different statistical models can be chosen for each part, depending on the patient's anatomy.

All our experiments have been performed using the multi-resolution scheme with 5 resolution levels. Our tests showed that for this particular segmentation task, similar results could be achieved using only a single resolution level. The multi-resolution scheme led, however, to a speedup of factor 2 to 5, depending on the parameter settings. The reason that the single resolution method does perform surprisingly well, lies on the one hand in the use of the distance image, which yields consistent gradient directions, and on the other hand the relatively good pre-alignment of the model. Other tests that we performed on CT images directly, confirmed that the multi-resolution does not only lead to faster convergence, but is a requirement for this fitting procedure to be more generally applicable.

## 6 Discussion

In this paper we have presented an approach for the segmentation of the human skull by means of fitting a statistical shape model to a pre-segmented MR image of the head. To efficiently solve the resulting optimization problem, we describe a multi-resolution approach to model fitting. Our method allows only solution that lie in the span of the statistical skull-models, i.e. linear combinations of the examples used to build the model. This strong restriction of the solution space allows the method to yield meaningful segmentation results, even in places where the bone cannot be distinguished from the surrounding tissue in the image. At the same time, it is this strong bias that limits the flexibility of the model. It would require a large number of example skulls in order to span the space of all normal skull shapes. We therefore proposed the use of a hierarchy of statistical models to increase the flexibility of the model, for the same number of examples. Our experiments showed that this greatly improves the accuracy of the segmentation.

The statistical shape model has been built from example skulls that were semi-automatically segmented from CT-Data. Hence our approach can be seen as a means of transferring shape information from CT-data to MR images. At some point this might reduce the need for acquiring a separate CT image of the patient and make it possible to obtain all information that is needed for the medical application from a single MR image.

The full potential of the hierarchical model has not been exploited in this paper. The fitting procedure described here is by no means limited to skull segmentation, but can be applied for any segmentation task, where a good enough pre-segmentation can be performed. A direct application is for the segmentation of bone from CT images, where the pre-segmentation can easily be performed by a simple thresholding operation. Moreover, one could consider more complex hierarchies and use separate model for every part. For example, a separate statistical model for each tooth could be built. Combined with an intelligent algorithm to automatically detect missing parts in data (using a scheme like the one described by Toews et al. [19]) it would be possible to fit images where no one-to-one correspondence to the reference-parts exists. The development of such a method, together with a more thorough evaluation of the segmentation algorithm, will be subject of future work.

**Acknowledgements** This work was funded by the Swiss National Science Foundation in the scope of the NCCR CO-ME project 5005-66380 (<http://co-me.ch>).

## References

1. 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, pp. 187–194. ACM Press (1999). DOI <http://doi.acm.org/10.1145/311535.311556>
2. Board, A.: Brief Communication: A Sample of Pediatric Skulls Available for Study. *AMERICAN JOURNAL OF PHYSICAL ANTHROPOLOGY* **103**, 415–416 (1997)
3. Cootes, T., Hill, A., Taylor, C., Haslam, J.: The Use of Active Shape Models for Locating Structures in Medical Images. *LECTURE NOTES IN COMPUTER SCIENCE* pp. 33–33 (1993)
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), 195–215 (2007)
5. Dedner, A., Lüthi, M., Albrecht, T., Vetter, T.: Curvature guided level set registration using adaptive finite elements. In: *Pattern Recognition*, pp. 527–536 (2007)
6. Demmel, J.W.: *Applied Numerical Linear Algebra*. SIAM (1997)
7. Dogdas, B., Shattuck, D., Leahy, R.: Segmentation of Skull and Scalp in 3-D Human MRI Using Mathematical Morphology. *HUMAN BRAIN MAPPING* **26**(4), 273 (2005)
8. Geraud, T., Mangin, J., Bloch, I., Maitre, H.: Segmenting internal structures in 3D MR images of the brain by Markovian relaxation on a watershed based adjacency graph. *Proc. of IEEE International Conference on Image Processing* **3**, 548–551
9. Ibanez, L., Schroeder, W., Ng, L., Cates, J.: *The ITK Software Guide*. Kitware, Inc. (2005)
10. Kang, Y., Engelke, K., Kalender, W.A.: A new accurate and precise 3d segmentation method for skeletal structures in volumetric ct data. *IEEE Trans. Med. Imaging* **22**(5), 586–598 (2003)
11. Kelemen, A., Szekely, G., Gerig, G.: Elastic model-based segmentation of 3-D neuroradiological data sets. *Medical Imaging, IEEE Transactions on* **18**(10), 828–839 (1999)
12. Lamecker, H., Seebaß, M., Hans-Christian, Hege, Deuffhard, P.: A 3d statistical shape model of the pelvic bone for segmentation. *Medical Imaging, IEEE Transactions on* **22**(2), 137–154 (2003)
13. Leventon, M., Grimson, W., Faugeras, O.: Statistical shape influence in geodesic active contours. *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on* **1** (2000)
14. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH Comput. Graph.* **21**(4), 163–169 (1987)
15. Lorigo, L., Faugeras, O., Grimson, W., Keriven, R., Kikinis, R.: Segmentation of Bone in Clinical Knee MRI Using Texture-Based Geodesic Active Contours. *LECTURE NOTES IN COMPUTER SCIENCE* pp. 1195–1204 (1998)
16. Rifai, H., Bloch, I., Hutchinson, S., Wiart, J., Garnero, L.: Segmentation of the skull in MRI volumes using deformable model and taking the partial volume effect into account. *Medical Image Analysis* **4**(3), 219–233 (2000)
17. Schäfer, J., Strimmer, K.: A Shrinkage Approach to Large-Scale Covariance Matrix Estimation and Implications for Functional Genomics. *Statistical Applications in Genetics and Molecular Biology* **4**(1), 32 (2005)
18. Schmid, J., Magnenat-Thalmann, N.: Mri bone segmentation using deformable models and shape priors. In: *MICCAI '08*. Springer (2008)
19. Toews, M., Arbel, T.: A Statistical Parts-Based Model of Anatomical Variability. *Medical Imaging, IEEE Transactions on* **26**(4), 497–508 (2007)
20. Umeyama, S.: Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**, 376–380 (1991)