An Evolutionary Strategy for Model–based Segmentation of Medical Data

Karin Engel and Klaus Toennies
Otto v. Guericke University Magdeburg, Germany

Abstract: Medical image segmentation often involves variants of deformable models to account for both the variability of object shapes and variation in image quality. Segmentation quality, however, highly depends on the initial estimate, and human guidance is often needed to guarantee acceptable results. For automating segmentation, our method employs a quality–of–fit function associated with a finite element model of shape in a search for the optimum parametrisation. A global search with an evolutionary strategy is employed to determine the set of optimum pose parameters for initialisation of the shape models. A local search subsequently optimises the non–rigid shape parameters by employing the deformable model paradigm. Experimental results are presented for different medical applications, which include object detection, localisation and segmentation, and show the good performance of our approach.

1 Introduction

The analysis of medical images demands efficient and precise automatic segmentation techniques, because manual segmentation is usually not straightforward and reproducibility is not satisfactory. Since medical data often come noisy and incomplete, while the anatomical structures of interest may vary dramatically from patient to patient and from scan to scan, segmentation can be seen as an ill–posed inverse problem. The general idea of solving ill–posed problems is to introduce some constraints to the solution [20], e.g. smoothness constraints. Hence, segmentation of medical data often involves deformable models [3, 5, 6, 10, 11, 13, 14, 19].

Under the deformable model framework, segmentation is considered as the problem of finding the contour $\mathcal{C} = \partial\Omega$ that represents the smooth boundary of a specific object Ω in an image I. This ill–posed problem is reformulated as finding the contour $\mathcal{C}$ that minimises the energy function

$$\epsilon(\mathcal{C}) = (1 - \zeta)\epsilon_{im}(\mathcal{C}|I) + \zeta\epsilon_{int}(\mathcal{C}), \tag{1}$$

where ϵ_{im} is an energy term that depends on the underlying image, ϵ_{int} is the regularisation term that depends on the deformation of the contour, and $\zeta \in [0, 1]$. This minimisation problem may be solved using gradient descend [5, 16], dynamic programming [1, 12], stochastic relaxation [21, 24] and the finite element method [4, 11], among others.

Success of deformable (shape) models, however, crucially depends on both reliable energies and the initial estimate, i.e. information about the object location, size and orientation of the desired object in an image. Directly estimating a globally optimum solution is usu-

ally impossible due to the high dimension of the configuration space. Also, the solution space is usually not convex, and multiple local optima may be found. Hence, currently human guidance in terms of initialisation [18] and generation of application–specific statistical shape models [5, 14] or atlases [9], is often needed to guarantee acceptable segmentation results. On the other hand, quality of the segmentations obtained by deformable models is rarely evaluated, but may be specifically useful for controlling the segmentation process.

This paper addresses these issues and presents a strategy for automating the deformable model–based segmentation of medical data sets. Our algorithm computes the most plausible explanation of the image content given a prior model of the desired object in terms of segmentation with maximum quality. In contrast to [19], where a highly application–specific segmentation plan is carried out by a single "self–aware" deformable agent, our algorithm initialises and optimises multiple model instances in parallel. Unlike the exhaustive search method used in [12] for global optimisation, our shape search uses an evolutionary strategy, which is inspired by the genetic algorithm Hill and Taylor employed for model–based segmentation of medical data [15]. Our approach also employs a population of individual solutions that undergo selection and "mutation" with the main difference that only the optimum pose parameters of the deformable shape model are estimated by the stochastic search. Further, no recombination (crossover) operators are used, and reproductive success does not vary with the relative "fitness" of the individuals. Instead, a rank–selection is implemented.

The search is organised by using a quality–of–fit (QOF) function, which not only accounts for support provided by the image features (as in [14, 15]), but uses a combined criterion according to equation 1. Compared with hybrid shape–based segmentation methods, e.g. [2, 3, 14], the QOF function allows evaluation of the model fit independent from the specific formulation of the prior model. Its general formulation allows us to focus on a-priori constrained prototypical models of shape variation (section 2.1), i.e. training regarding shape and appearance as in [2,3,5,8,12,14] is not required. Statistical information can, however, be incorporated in a straightforward manner [7]. We will demonstrate that the proposed objective function (section 2.2) allows determining success or failure of the deformable model fit in an evolutionary shape search. Our search strategy is described in section 2.3. Experimental results are provided in section 3 for different medical image segmentation applications, which allows us to demonstrate the use of our method as a means to detect, localise and segment the desired anatomical structures.

2 Method

By shape decomposition objects of similar types are represented using a combination of a set of basis functions with global support, where the basis can be either pre–defined in an a-priori manner [22], or obtained via training [5]. We assume that a prototypical representation of shape sufficiently captures variation of the desired object class, and valid instances of the object can be reconstructed from the set of a-priori constrained model

parameters. We therefore exemplarily employ the finite element decomposition of shape, which supports an efficient simulation of deformation.

The quality of a model instance projected into the image and deformed according to image–based forces, is evaluated using a quality–of–fit function, which estimates the probability that the image segmentation results are valid w.r.t. the prior model, and represent the true solution. This allows object localisation and segmentation to be integrated in a deformable shape search. The proposed evolutionary algorithm computes an optimal set of affine transformation parameters for automatic initialisation of the shape models. The optimum set of non–rigid shape parameters is found for each model instance by a local search that is solved using the finite element method. The best solutions are selected from a population of multiple deformed shape model instances and further evolved by initialising new shape instances with similar pose parameters.

2.1 Implementation of The Parametric Prototypical Deformable Model

A parametric deformable template $\mathcal{T}(\mathbf{p})$ represents the objects undeformed shape (rest shape or equilibrium shape) and a set of parameters $\mathbf{p}$ that define how it deforms under applied forces. The rest shape of a n–dimensional object can be understood as a continuous domain $\Omega \subset \mathbb{R}^n$, and its deformation is described by a boundary value partial differential equation. It is solved for the unknown displacement field $u(\mathbf{x})$, $\mathbf{x} \in \Omega$, using the finite element method yielding an algebraic function [4],

$$\mathbf{K}\mathbf{u}(t) = \mathbf{f}_b(t) + \mathbf{f}(t). \tag{2}$$

Equation 2 relates at each time step t of the numerical simulation the deformed positions $\mathbf{x}(t) = \mathbf{x}^0 + \mathbf{u}(t)$ of all N finite element nodes of the deformable template to the forces acting on it. Here, $\mathbf{u}(t)$ is the vector of nodal displacements and $\mathbf{x}^0$ denotes the nodal rest positions. $\mathbf{K}$ is a function of the elastic modulus E and Poisson ratio ν, and encapsulates the stiffness properties as well as the type of mesh and discretisation used. The vectors $\mathbf{f}_b(t)$ and $\mathbf{f}(t)$ denote the dynamic body forces and external model forces, respectively. The dynamic equilibrium equation has the form

$$\frac{\partial^2 \mathbf{u}}{\partial t^2}|_{t>0} = \mathbf{M}^{-1}(-\mathbf{C}\frac{\partial \mathbf{u}}{\partial t}|_{t>0} - \mathbf{K}\mathbf{u}(t) + \mathbf{f}(t)). \tag{3}$$

To simplify analysis, $\mathbf{C}$ and $\mathbf{M}$ represent velocity–dependent damping characteristics and a constant function of material density ρ, respectively [4, 22]. The deformed positions can be expressed in terms of a superposition of $m = m_2 - m_1$ displacement fields,

$$\mathbf{x}(t) = \mathbf{x}^0 + \sum_{k=m_1}^{m_2} \phi_k \mathbf{q}_k(t), \tag{4}$$

where $m_1 \geq 1$ and $m_2 \leq s$ for $s = nN$ degrees of freedom of the system. The modal vectors ϕ_k are solutions to the eigenproblem $(\mathbf{K} - \omega_k^2 \mathbf{M})\phi_k = 0$ and the $\mathbf{q}_k$ are the modal amplitudes, computed as described in [22]. Letting

$$\mathbf{p}^t = (\theta, \mathbf{q}(t)), \tag{5}$$

where θ denotes the affine transformation from the model coordinate frame to the image coordinate system, we can write equation 4 as $\mathbf{x}(t) = \theta(\mathbf{x}^0 + \sum_k \phi_k \mathbf{q}_k(t))$, and obtain a compact characterisation of the deformed shape $\mathcal{T}(\mathbf{p}^t) = \mathbf{x}(t)$.

The matching of model and data can be viewed as a *local optimisation*. In our case, a shape model deforms from an initial estimate $\mathcal{T}(\mathbf{p}^0)$ into an object instance $\mathcal{T}(\mathbf{p}^t)$ supported by features derived from the image I, such that solving equation 3 yields a balance of internal and external forces. External model forces $\mathbf{f}(t)$ are created by a sensor–based sparse sampling of a scalar potential field $\mathcal{W}$, whose local minima coincide with characteristic object features in the image, i.e.

$$\mathbf{f}(t) = -\nabla \mathcal{W}(\mathbf{x}(t)). \tag{6}$$

For example, contour nodes that typically represent the boundary of the desired object are usually associated with sensors that sample gradient magnitude maps of the form $|\nabla I|^2$ [16]. In contrast to the energy minimising formulation of deformable models according to equation 1, the force formulation also permits the use of more general types of external forces, which cannot be computed as the gradient of a scalar potential field $\mathcal{W}$, but e.g. using local texture information [11]. In the general case, a force field may be defined based on a map $\mathcal{D}(\mathbf{x})$ of the distance of each image pixel to the closest relevant point (that may or may not be defined based on a Gaussian potential field),

$$\mathbf{f}(t) = \kappa \mathcal{D}(\mathbf{x}(t)), \quad \kappa > 0. \tag{7}$$

As already mentioned, the set of basis functions (cf. equation 4) may alternatively be obtained by statistical analysis of training samples. The resulting Point Distribution Model or Active Shape Model (ASM) introduced by Cootes et al. [5] also estimates points of interest along the model boundary and iteratively deforms the model from its reference shape into similar shapes that match the target points. The optimum displacement of the model nodes may be based on the response to a linear filter or higher–order statistics represented by an Active Appearance Model (AAM) [6]. Consequently, the presented concepts for shape representation, fit and global search (cf. sections 2.2 and 2.3) also apply to statistical shape models. Incorporation of statistical information simply requires a basis transform [7] to determine deformations according to equation 4 by a set of constraints corresponding to variation modes instead of physically–based vibration modes ϕ_k.

2.2 The Quality–of–fit Function

Regularisation has close connections to Bayesian Maximum A Posteriori (MAP) estimation [20], i.e. each deformed model instance provides a locally optimum segmentation result[1]. For using the model in a search for the desired object, a high value of the associated quality–of–fit (QOF) function should indicate cases, where model and data match perfectly without significant deformation. Otherwise the QOF–value should be low. We therefore propose an objective function that combines a measure of deformation Q_{int} of

[1] Minimising the energy of the contour $\mathcal{C} = \partial\Omega$ defined in equation 1 is - with suitable assumptions - equivalent to finding from an initial estimate $\mathcal{C}^0$ the contour $\mathcal{C}$ that maximises a posterior probability $\epsilon(\mathcal{C}|I, \zeta, \mathcal{C}^0)$.

a model instance with its correlation with the data Q_{im} for estimating the total quality similar to equation 1,

$$\mathcal{Q}(\mathcal{T}) = (1-\zeta)Q_{im}(\mathcal{T}) + \zeta Q_{int}(\mathcal{T}),\ \zeta \in [0,1]. \tag{8}$$

We aim at formulating Q_{im} and Q_{int} such that their values are normalised to the interval $[0,1]$, where values close to 1 indicate high quality. The QOF–value $\mathcal{Q}(\mathcal{T})$ indicates the probability that the image segmentation result is valid w.r.t. the prior model $\mathcal{T}(\mathbf{p})$, and represents the true solution to the segmentation problem.

The quantity Q_{im} is usually determined by comparing the expected data with the candidate segmentation. Several approaches employ the evidence for edges along the model boundary [15], and evaluate the data fit as the average sensor response at each model node [2,9]. This simple approach may, however, not always be appropriate because object boundaries may or may not be indicated by high intensity gradients, while the intensity may vary in the object interior due to noise and pathological effects. Other methods are based on the local AAM search, which is driven by normalised texture differences between the model and the image [6]. For example, Heimann et al. employ the statistical evidence for the shape by using the posterior probabilities of local grey–value appearance [14]. We presented a similar approach that does not require training of textural information [11]. In all aforementioned cases the correspondence of the model with the expected data can be estimated using a force formulation according to equation 7. Hence in the ideal case, at the final step t of the simulation the nodal forces are assumed to vanish, i.e.

$$Q_{im}(\mathcal{T}) = \mathcal{F}\big(\mu(|\mathbf{f}_i(t)|^2)\big). \tag{9}$$

Here, μ computes the mean value over all nodes i and $\mathcal{F}$ maps the resulting values to the interval $[0,1]$, e.g. using the negative exponential.

As the space of possible solutions is constrained by the prior model, the obtained segmentations are assumed to be of similar shape, and the "degree of similarity" in terms of the deviation from the reference shape is rarely evaluated. Bergner et al. compute the ratio of the sum of all spring lengths over the sum of their rest lengths for a mass spring model of shape [2]. A drawback is the lack of specificity of their objective function that penalises global deformations, while the influence of local deformations may be neutralised due to averaging. Similar criteria were used in [9, 12], but apply only to specific representations of the prior model. The shape decomposition approach, on the other hand, allows a comparison of shape instances based on their coordinates in the shape space $\mathcal{S}(\mathcal{T})$ [17]. It is spanned by the set of non–affine deformation parameters associated with $\mathcal{T}$, such that similar shapes correspond to nearby locations. Hence, the quantity Q_{int} can be computed from the non–rigid deformation of the shape model instance in its un–rotated reference frame. In our case the strain energy, which is associated with the modal displacements, is adapted from [22], such that

$$Q_{int}(\mathcal{T}) = \mathcal{F}\big(\mu\big(\mathbf{q}_k^2(t)\omega_k^{-2}\big)\big). \tag{10}$$

In order to exclude the effect of global variations and reduce the influence of the high–frequency modes, which are sensitive to noise, we only consider the modal amplitudes

from a low–dimensional embedding $\left(\phi_{m_1} \dots \phi_{m_2}\right)$. We select the $m \ll s$ intermediate modes that explain a proportion β, e.g. $\beta = 0.25$, of the total variation (section 2.1).

Note that in contrast to other approaches, e.g. [9, 12, 14], the proposed quality–of–fit can be evaluated for any decomposition–based shape model, e.g. ASM [5]. In this case the $\mathbf{f}_i$ (in equation 9) may be evaluated as described in equation 7, and the $\mathbf{q}_k$ (equation 10) are coordinates in a statistically constrained shape space $\mathcal{S}(\mathcal{T})$.

2.3 The Deformable Shape Search

Solving the segmentation problem requires estimation of the optimal set of parameters $\mathbf{p}^t$ (equation 5), which represent the affine and non–rigid deformations of the shape model $\mathcal{T}$ that best matches the image. If we assume the model to be adequate w.r.t. the shape constraints and response to the crucial image feature detectors, the model is deformed from an initial estimate $\mathcal{T}(\mathbf{p}^0)$ into a similar shape by minimising an associated energy function (equation 1). In our case, the desired solution to equation 3 yields a balance of internal and external forces, and shall optimise the quality–of–fit. A solution $\mathcal{T}(\mathbf{p}^t)$ with a maximum QOF–value $Q(\mathcal{T})$ indicates the *globally optimal* segmentation iff the model is properly initialised. Hence, a global search is performed for possible initialisations.

The optimum set of affine transformation parameters is determined using an evolutionary strategy. Our deformable shape search initialises and optimises multiple model instances in parallel. Similar to genetic algorithms [15], our algorithm computes the most plausible explanation of image content given a prior model in terms of segmentation with maximum quality. A population of candidate solutions is evolved to one or several local optima within the search space by computing a ranked selection. We therefore evaluate the fitness of each instance using the QOF–function (8) introduced in section 2.2. According to the principles of stochastic optimisation [23], randomness is introduced in the direction of the global search. Since in our case the objective function to be maximised depends on the parametrisation of the deformable model, spontaneous variability is introduced similar to a random search by a "mutation" of the pose parameter values of selected individual solutions. Our implementation of the stochastic (Monte Carlo) sampling does not include cross–over mutations or recombination, but employs a zero–mean Gaussian mutation (cf. equation 12) to account for the non–convex nature of the objective function (Figure 2(a)).

Model instances are initialised by a transformation of the prototype $\mathcal{T}(\mathbf{p}^0)$ from the model coordinate frame to the image coordinate frame. Assuming the center of mass of the rest shape at the origin of the local frame of reference, the transformation θ defines the position $\mathbf{c}$, orientation ψ and scaling $\mathbf{s}$ of the model instance in the image. For the evolutionary shape search the pose parameters $\xi = \{\mathbf{c}, \mathbf{s}, \psi\}$, are considered as variates with a presumed

Gaussian[2] distribution $\xi \sim N(\mu, \varsigma)$, such that random, real–valued samples

$$x = \hat{\mu}(\xi) + r, \tag{11}$$

where

$$r = z\sqrt{\hat{\varsigma}(\xi)} \text{ and } z \sim N(0,1). \tag{12}$$

In contrast to [2,14,15], training data for estimating the parameters $\hat{\mu}$ and $\hat{\varsigma}$ of the probability density functions are in our case not available. We therefore specify an initial region of parameter values we are interested in. More specifically, we use pre–set tolerances $\hat{\varsigma}$ from the parameter values x' of a single model instance $\mathcal{T}^*$ generated from the representative manual segmentation, which serve as estimates for $\hat{\mu}$ (see sect. 3 for settings we used in our experiments). This kind of assumptions may be introduced by a human expert [9, 19], but can often also be derived from inspection of the images and example segmentations[3].

The search procedure is summarised in algorithm 1. In contrast to [14,15], a rank selection is used instead of employing probabilities for the selection of individual solutions. We organise the iterative search by employing a priority queue $\mathcal{R}$ of candidate solutions, where we use the quality–of–fit of the current model instances as the priority (step 4 of algorithm 1). Solutions with high quality $\mathcal{Q}(\mathcal{T}) > \tau_Q = q - \tau$, are selected (step 5), and further evolved until the overall quality of the current model instances,

$$q = \mu_b\big(\max(Q(\mathcal{T}) : \mathcal{T} \in b)\big), \tag{13}$$

converges. The threshold τ is predefined. For determining q, clusters of model instances with high energy are built by assigning them based on their center of mass to one of the bins b, which define a regular grid over the image[4]. This clustering allows to discard redundant shape instances. We identify close solutions for which the maximum distance of corresponding nodes is below a threshold δ_{user}, e.g. $\delta_{user} = 2mm$, and select the higher rated instance from each cluster of close solutions.

New shape generations are then generated in each iteration of the global search based on the parametrisation of the regionally (per cluster) best fitting shapes in terms of perturbations of the particular pose parameter values (step 6). Therefore, to the selected pose parameters θ we add a random vector created from a multi–variate zero mean Gaussian (step 7). More specifically, for each new instance we use in equation 11, $\hat{\mu}(\xi) = \upsilon'$, where υ' is the actual value for ξ from an instance with high quality, and in equation 12 we let $\hat{\varsigma}(\xi) = \bar{\varsigma} \in [0, 1]$. The values for $\bar{\varsigma}$ decrease over time, e.g. from $\bar{\varsigma}_{user}$ to 0.05. This strategy is related to annealing type optimisation methods, and helps to improve convergence.

[2]Gaussian sampling may be inadequate, for example for determining the optimum scaling (which depends on the selected field of view), or position of the desired objects. Alternatively, a uniform distribution may be initially assumed. Later this information can be refined by employing a match list of known solutions. Such importance sampling would also allow avoiding repetitive computations by discarding parametrisations in case of close proximity in the parameter space.

[3]The application–specific search ranges help to increase efficiency and robustness of the search. If no prior is known, a coarse exhaustive search can be performed by employing uniform sampling. In medical applications, however, the set of possible affine transformations can often be constrained a–priori on the basis of standardised acquisition parameters and general anatomical knowledge.

[4]Since the exact relation of the image size $\mathbf{s}_I$ and object is not in advance, we relate the size of a bin to the size of the template $\mathcal{T}^*$ constructed from the example segmentation, such that the bins are of size $0.5\hat{\mu}(\mathbf{s})$, where $\hat{\mu}(\mathbf{s}) = \mathbf{s}_I^{-1}\mathbf{s}_{\mathcal{T}^*}$. Alternatively, the bin size may vary dynamically with the size of the current best fitting shape.

Algorithm 1 Evolutionary shape search.

1: Initialise a population of shape model instances $\mathcal{P}(0) = \{\mathcal{T}(\mathbf{p}_j^0)\}, j = 1, \ldots, J(0)$, according to the pre–set distribution parameters in equations 11 and 12.
2: **repeat** $\{T = (T + 1)\}$
3: **Deformable Shape Fit:**
Fit each new individual model instance to the data by solving equation 3 to obtain a set of locally optimum solutions $\mathcal{P}(T) = \{\mathcal{T}(\mathbf{p}_j^t)\}, j = 1, \ldots, J(T)$.
4: **Evaluation**:
Evaluate the quality of fit $Q(\mathcal{T}(\mathbf{p}_j^t)), \forall j$ (equation 8) and generate a ranked list $\mathcal{R}(T)$ of the population $\mathcal{P}(T)$.
5: **Clustered Rank Selection**:
Select from the priority queue $\mathcal{R}(T)$ the $J_A \leq J(T)$ best fitting instances per bin g with $Q(\mathcal{T}(\mathbf{p}_j^t)) > \tau_Q$. Discard the $J_B < J_A$ redundant candidate solutions.
6: **Reproduction**:
Generate new instances based on the selected $\mathbf{p}_j^t, j = 1, \ldots J_A - J_B$, i.e.

$$\mathbf{p}_{new,j}^0 = \left(\theta_{new,j}, \mathbf{q}_j^0\right),$$

where $\theta_{new,j} = \theta(\xi_j)$ (and $r = 0$ in equation 11), and $\mathbf{q}_j^0 = 0$.
7: **Gaussian Randomisation**:
Introduce variability in the values $\mathbf{p}_{new,j}^0, j = 1, \ldots J_A$, by adding a small random vector to $\theta_{new,j}$ (equation 12).
8: **Initiate Blind Search:**
Replace the $J(T) - (J_A - J_B)$ instances by a new population according to step 1.
9: **until** convergence.
10: **return** the list of ranked solutions $\mathcal{R}(T)$.

Shape model instances with low quality (and redundant shapes, resp.) are replaced in step 8 by new instances. Misleading shape searches due to an insufficient parametrisation as well as an exponential increase in the number of solutions can thus be avoided, while the additional new trials keep the search independent of known solutions.

The search converges if no improvement in overall quality q (and quality of the best fitting model instance provided at the top of the priority queue) can be achieved with a pre–defined minimum and/or maximum number of trials. Depending on the desired precision, the search may be terminated when a solution $\mathcal{T}(\mathbf{p}_j^t) \in \mathcal{R}$ can be reported with a QOF–value above a pre–defined threshold τ_{user}. On the other hand, if the QOF–value $Q(\mathcal{T}(\mathbf{p}_j^t))$ of the best candidate solution is below the pre–defined threshold, the desired object was not detected in the given image.

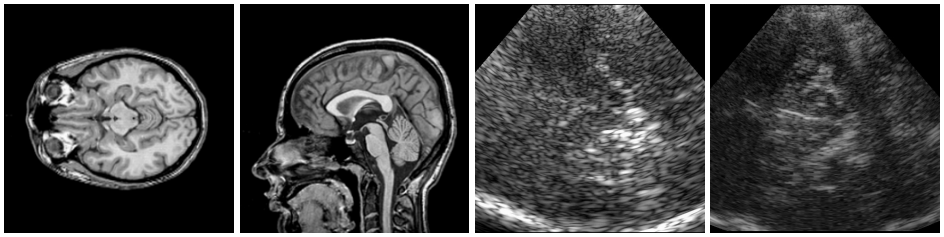

Figure 1: Example images used in our experiments. From left to right: transversal MRI slice (A), sagittal MRI slice (C), TCS image (B), TCS image acquired using Tissue Harmonic Imaging (THI).

3 Experimental Evaluation

We chose three representative medical applications to validate our approach to deformable shape search:

A Segmentation of the mesencephalic brain stem (midbrain) in transversal slices from cerebral MRI data sets,

B Segmentation of the midbrain in transcranial sonographies (TCS),

C Segmentation of the corpus callosum in sagittal slices from MRI data sets.

30 T1–weighted MRI data sets were available (from 30 subjects, scanned at 3 Tesla) with a resolution of 256^3 isotropic pixels of size $1mm^3$. The 20 TCS data sets from 9 patients and 1 healthy subject were acquired with varying imaging protocols in clinical routine. The subjects were scanned using conventional B–mode ultrasound (and using Tissue Harmonic Imaging (THI) in 14 of the 20 data sets). The TCS data sets had an in–plane resolution of 666×888 isotropic pixels with a variable spacing between $0.13mm$ and $0.27mm$. Example data sets are provided by Figure 1.

All data sets were available in 8 bit format, but varied w.r.t. the signal to noise ratio, as well as magnetic field inhomogeneities (in the MRI data) and echomorphology of the midbrain (in the TCS data sets), respectively. An automatic segmentation without prior shape information and initialisation close to the desired solution was prone to fail due to inter–subject anatomical variability, as well as the low contrast to neighbouring structures. The segmentation of the midbrain from the TCS images was particularly challenging due to the highly variable quality of the temporal bone window of the scanned subjects, the non–deterministic characteristic of ultrasound data, artefacts (e.g., due to reflection and absorption of the ultrasound by the calvaria), and pathological changes in the echotexture.

The segmentation was performed in all cases in 2D allowing an evaluation of our method using qualitative and quantitative analyses based on ground truth segmentations provided by neuroscientists. We constructed 2D–shape models, which prototypically represent the rest shape of the corpus callosum and midbrain, from a randomly selected expert segmentation in one example data set per application. These segmentations were available in terms of manually delineated contours. The finite element models were build on the basis

of a triangulation of the segmentations, and implemented using linear shape functions for interpolation of the field variables over the finite element mesh.

3.1 Parameter Settings

We conducted three different experiments, where we first evaluated localisation and segmentation accuracy under the assumption that the desired anatomical structure was contained in the input images (exactly once). In the second experiment (C) we selected from each MRI data set 2 sagittal slices, which did / did not contain the desired structure, respectively, in order to explore the ability of the proposed method to also detect the presence of the desired object in the input image.

The segmentation consisted of two steps, finding the optimum pose parameters using the presented evolutionary strategy and searching for the locally optimum non–rigid shape parameters by employing the finite element method for optimisation.

In our experiments the local shape searches were parametrised using $E = 0.9, \nu = 0.25, \rho = 1$ and $\kappa = 100$. We computed application–specific external model forces. The distance–based potential functions $\mathcal{D}$ (equation 7) were pre–computed for the MRI data sets on the basis of intensity gradient magnitude maps. In application B texture–based forces were estimated at runtime on the basis of first and second–order statistics over the actual segmentations [11]. In order to improve the runtime of the local searches, we pre–computed the first–order statistics over the TCS images. As convergence criterium we used a minimum change in nodal displacements and a maximum number of iterations ($t < 1000$) of the numerical simulation, respectively.

For the deformable shape search we let in all experiments $\beta = 0.25$, i.e. $m = 11$ intermediate modes were used in equation 10 for application A, and $m = 6$ for applications B and C. We further let $\zeta = 0.5$, $\tau = 0.1$, and $\tau_{user} = 0.7$ was the required minimum QOF–value of the final solution. For mapping the quality–of–fit values according to equation 8 to the interval $[0, 1]$, we used the negative exponential function $\mathcal{F}(Q) = \exp(-\alpha Q^2)$ with $\alpha = 0.1$. We set $\delta_{user} = 2mm$ and $\bar{\varsigma} = 0.4$.

The initial constraints on the search space were $\hat{\varsigma}(\mathbf{c}) = 0.1\mathbf{s}_I$, with $\mathbf{s}_I$ being the image width, $\hat{\varsigma}(\psi) = 2^\circ$ and $\hat{\varsigma}(\mathbf{s}) = 0.1\hat{\mu}(\mathbf{s})$, except for (C) where $\hat{\varsigma}(\mathbf{c}) = 0.2\mathbf{s}_I$. The global search was run for a maximum number of $T = 20$ iterations with a population of initially $J(0) = 40$ instances in the localisation/segmentation experiments. We chose the particular number of instances according to the application–specific restrictions on the search space, i.e. a higher number of $J(0) = 60$ initial individuals was chosen for experiment C, because here a larger region of possible locations in the image coordinate frame was explored.

Iteration of the global search stopped if the overall quality did not improve after $T > 5$ iterations and the quality of the best rated solution in the rank list $\mathcal{R}$ was above τ_{user}.

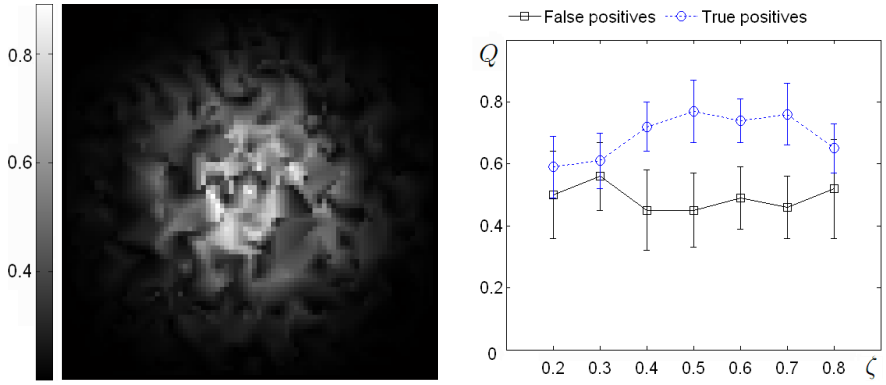

(a) (b)

Figure 2: Figure (a) shows the quality–of–fit evaluated for a small region of 81×81 positions around the true location within an example image (application C), using a population of shape model instances with fixed optimum size and orientation. Multiple local maxima can be identified in this subregion of the search space. Figure (b) shows the influence of the regularisation parameter ζ (equation 8) on the significance of the QOF–values. For the default value $\zeta = 0.5$, true and false positive segmentation results can be identified based on a threshold $\tau_{user} = 0.7$.

3.2 Experimental Results

Since our method is based on the assumption that valid instances of the object can be reconstructed from the a-priori constrained model parameters (cf. section 2.3), we first evaluated the constructed shape models w.r.t. target error. We found that segmentation was always successful given a proper manual initialisation close to the object of interest in the data (note that in this experiment we manually selected from the MRI data sets (C) the sagittal slices containing the corpus callosum). In all cases, the resulting segmentations showed a high overlap with the ground truth segmentations of $o = 95 \pm 2\%$ ($\mu \pm \sigma$). Moreover, the QOF–values increased during the deformation process at average from $\bar{Q}(\mathcal{T}(\mathbf{p}^0)) = 0.42$ to $\bar{Q}(\mathcal{T}(\mathbf{p}^t)) = 0.75$. We also found from this experiment with user initialisation that slight *deterministic* changes in the parametrisations not necessarily resulted in further improvement, i.e. not in all cases the segmentation quality was inversely proportional to the distance from the reference segmentation. The non–convexity of the quality–of–fit function over the search space is illustrated in Figure 2(a).

Our results from the second experiment (automatic localisation) indicated that the stochastic search algorithm 1 can, however, cope with this situation, as it may less likely get stuck in local optima compared with deterministic methods. Table 1 shows that in all cases low values indicated a misled shape search, while the best fitting instance represented a valid solution (type 1–error was $p < 0.05$) close to the true solution. Typical results of our method are displayed in Figure 3.

Both the adequacy of the deformation and significance of the QOF–function (8) depend, however, on the actual choice of material parameters. All our experiments have been

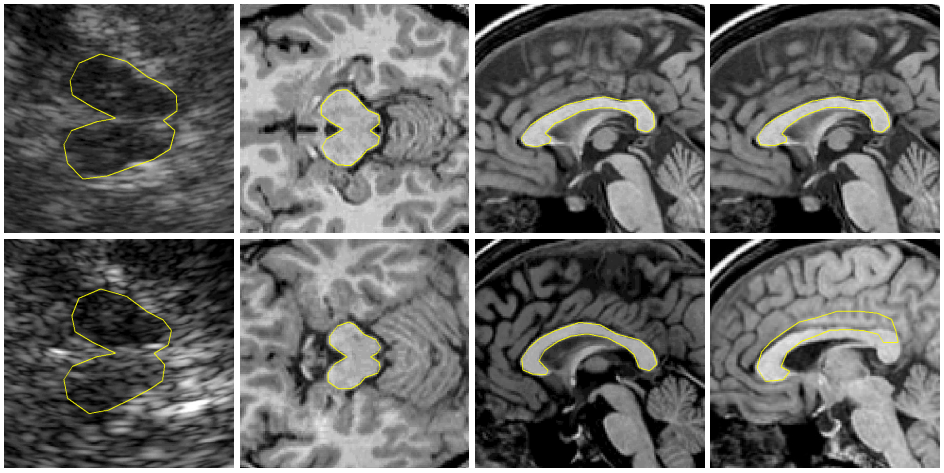

Figure 3: Typical segmentation results with a QOF $Q(\mathcal{T}) > \tau_{user} = 0.7$ for the example applications studied (The images were cropped only for displaying the regions of interest.). The solution in the top right figure was favoured by the expert but assigned a smaller QOF–value $Q = 0.78$ (rank 9 of the priority queue) compared with the solution in the top middle right figure (best rated solution with $Q = 0.81$). Such cases happened rarely, but need further evaluation. The bottom right figure shows a false positive solution reported by our algorithm in the detection task (C) on the basis of a QOF–value $Q > \tau_{user}$. The overlap with the ground truth segmentation was below 90%.

done with the same set of empirically chosen parameter values (listed in section 3.1) that yielded plausible deformations. From our experience, the method is robust to changes in the parameter values.

A rigid template matching, i.e. an exhaustive search for the optimum affine set of parameters using a correlation measure as objective function, also provided reliable initial estimates for applications A and C. For the non–deterministic TCS data (B), however, false positive solutions could be avoided in 9 of 10 cases only by combining texture–based forces [11] and the proposed deformable shape search. This optimisation strategy is computationally more expensive. However, from the third experiment on object detection (case study C) we can conclude that an initially less constrained search space (e.g., assuming an uniform distribution of the object locations) was efficiently reduced by employing match lists of known solutions (i.e. importance sampling). By avoiding repetitive computations the overall running time was reduced by about 40%.

The best rated segmentations provided at the top of the priority queue always compared well with the manual segmentations (see table 1). Average overlap was $94 \pm 1\%$ and boundary errors were at the order of $1mm$ for experiments 2 and 3, where we evaluated our method to the automatic localisation/detetction and segmentation of specific shapes. These results compare with the results from experiment 1, were the shape model instances were manually initialised. This clearly shows that the combination of a prototypical shape template – and reliable, i.e. image modality dependent, model forces – with the proposed search strategy allows for both localising and segmenting the desired objects.

As expected, the detection rate for application C could be improved using a more liberal threshold τ_{user} only at the cost of an increasing false positive rate (increasing e.g. for $\tau_{user} = 0.6$ from 0.03 to 0.2). The proposed QOF–function, however, was more significant compared with an purely image–based objective function (realised by letting $\zeta = 0$ in equation 8), as used in [15].

Figure 2(b) provides statistics over the quality of segmentations resulting from successful and misleading shape searches, which were computed by applying to all data sets the deformable shape search with different parametrisations for the regularisation factor $\zeta \in [0, 1]$. For an analysis, we classified the candidate solutions from the final populations $\mathcal{P}(T)$ according to their overlap with the manual segmentations as true and false positives, respectively. A false positive was defined by an area overlap of $o < 90\%$. Our results indicate the importance for evaluating segmentation quality of including the amount of non–rigid deformation necessary to fit the data. Most notably, in the detection experiment no false negative solution was reported at the top of the priority queue $\mathcal{R}(T)$ for $\zeta \in [0.4, 0.6]$. In all cases QOF–values above 0.7 could be interpreted as a high probability for the true solution, while no preferred value $\zeta \in [0.4, 0.6]$ for the influence of the strain energy became evident from our results.

In contrast to the method proposed by Hill and Taylor [15], where the search space is spanned by the set of affine *and* non–rigid parameters, our global search computes the optimum pose parameters only. By incorporating a population of deformable model instances, however, hidden dimensions are added to the problem making the presented search also more efficient than regular sampling [23]. Our approach, however, cannot guarantee to find the globally optimal parameter set. This would require a recursive subdivision of the parameter space and analysis of all possible matchings for transformations with parameters contained in the respective region, possibly resulting in an exhaustive search.

We finally applied the shape search for all data sets with an increased initial population size $J(0) = 100$. Several strategies (e.g., a-priori constraining the search space, importance sampling, clustered ranking, pre–computation of image–based force fields, etc.) allowed rendering the search efficient. The computational load, however, was at the order of 80%

	instances			TP			FP		time
	$J(0)$	overall	$Q > \tau_{user}$	rate	$\bar{Q}$	δ	rate	$\bar{Q}$	
(A)	40	406	9	1.0	0.74	0.8 ± 0.2	0	-	1004
(B)	40	418	12	1.0	0.81	1.2 ± 0.3	0	-	2733
(C)	60	559	14	0.97	0.85	0.6 ± 0.1	0.03	0.82	1128

Table 1: Segmentation results. The table provides the number of instances initialised, the size of the population $\mathcal{P}(T)$ after convergence, and the number of instances from $\mathcal{R}(T)$ with a QOF–value above τ_{user}, the fraction (%) of true positive (TP) and false positive (FP) solutions and their average QOF–value $\bar{Q}$, the average squared distance δ (in mm) to the manual segmentations, and the average running time (given in seconds) for a complete search (on a 2.4GHz P4, Matlab/C). Note that the running time in (B) is largely dominated by on–line computation of the texture–based external model forces according to [11].

higher compared with the original settings (sect. 3.1), while no improvement over the results presented in table 1 was achieved. This effect may be due to the noise and ambiguity in the data as well as the possible lack of specificity and/or sensitivity of the prior model, and needs further evaluation.

4 Conclusion and Outlook

This paper presented an approach to automatic medical image interpretation by employing a quality–of–fit function in a deformable shape search. Given our assumptions on the decomposition–based representation of specific shapes are met, the proposed evolutionary strategy controls the segmentation process and returns multiple model–based interpretations of the image along with a confidence measure from the interval $[0, 1]$. A major difference to existing shape search strategies [2, 5, 9, 12, 14, 15, 19] is that the presented quality–of–fit function allows evaluation independent from the specific formulation of the parametric model regarding its basis and external model forces. Experimental results showed its adequacy for deciding upon the success of the segmentation. In all cases where small deformations were required to bring the prototypical model into correspondence with the data, the resulting segmentations were indicated by high QOF–values, while poorly fitted model instances yielded low values. This effect was independent from the application/modality–specific external model–forces and image quality.

As the proposed evolutionary strategy was shown to work successful with little, prototypical prior information, our method may become an option especially for the segmentation of higher dimensional data. We believe that our approach provides a means to replace current strategies that rely on purely image–based criterion functions [14, 15] and/or require statistical a-priori information [3,5,6,9,12,14], which may not always be available. Future research will focus on the detailed evaluation of the influence of the various parameters on possible random and systematic errors of our approach. In particular, we aim at finding an optimum parametrisation on the basis of exemplar–based segmentations in order to improve the a-priori information involved in both the local and global search.

References

[1] Amini A., et al.: Using Dynamic Programming for Solving Variational Problems in Vision. IEEE Trans Patt Anal Mach Intell **12**(9) (1990) 855–867

[2] Bergner S., et al.: Deformable structural models. Proc. IEEE ICIP (2004) 1875–1878

[3] de Bruijne M., et al.: Shape particle filtering for image segmentation. In: Barillot C., et al. (Eds.) Proc. MICCAI'04, LNCS **3216**, Springer (2004) 168-175

[4] Cohen L.: On active contour models and balloons. CVGIP Imag Underst (1991) **53**(2) 211-218

[5] Cootes T., et al.: The use of Active shape models for locating objects in medical images. Imag Vis Comp **12**(16) (1994) 355–366

[6] Cootes T., et al.: Active Appearance Models. Proc. ECCV (1998) 484–498

[7] Cootes T. and Taylor C.: Combining point distribution models with shape models based on finite–element analysis. Imag Vis Comp **13**(5) (1995) 403-409

[8] Cremers D., et al.: Nonlinear Shape Statistics in Mumford-Shah Based Segmentation. Proc. ECCV (2002) 93–108

[9] Dornheim L. and Dornheim J.: Automatische Detektion von Lymphknoten in CT–Datensätzen des Halses. Proc. BVM (2008) 308–312

[10] Engel K., et al.: A two–level dynamic model for the representation and recognition of cortical folding patterns. Proc. IEEE ICIP (2005) 297-300

[11] Engel K. and Toennies K.: Segmentation of the midbrain in transcranial sonographies using a two-component deformable model. Proc. MIUA (2008) 3–7

[12] Felzenszwalb P.: Representation and detection of deformable shapes. IEEE Trans Patt Anal Mach Intell **27**(2) (2005) 208-220

[13] He L., et al.: A comparative study of deformable contour methods on medical image segmentation. Imag Vis Comp **26** (2008) 141–163

[14] Heimann T., et al.: A Shape-Guided Deformable Model with Evolutionary Algorithm Initialization for 3D Soft Tissue Segmentation. In: Karssemeijer N. and Lelieveldt B. (Eds.) Proc. IPMI'07, LNCS **4584**, Springer (2007) 1-12

[15] Hill A. and Taylor C.: Model-based image interpretation using genetic algorithms. Imag Vis Comp (1992) **10**(5) 295-300

[16] Kass M., et al: Snakes: Active contour models. Int J Comp Vis **1**(4) (1988) 321–331

[17] Kendall D.: A survey of the statistical theory of shape. Statistical Science **4** (1989) 87-120

[18] Maleike D., et al.: Lymph node segmentation on CT images by a shape model guided deformable surface method. Proc. SPIE **6914** Medical Imaging (2008)

[19] McInerney T., et al.: Deformable organisms for automatic medical image analysis. Med Imag Anal **6**(3) (2002) 251–266

[20] Poggio T., et al.: Computational vision and regularization theory. Nature **317** (1985) 314–319

[21] Rueckert D. and Burger P.: Contour fitting using stochastic and probabilistic relaxation for cine MR images. Proc. CAR (1995) 137–142

[22] Sclaroff S. and Pentland A.: Modal matching for correspondence and recognition. IEEE Trans Patt Anal Mach Intell (1995) **17**(6) 545-561

[23] Spall J.: Introduction to Stochastic Search and Optimization. John Wiley & Sons (2003)

[24] Stegmann M., et al.: Extending and applying active appearance models for automated, high precision segmentation in different image modalities. Proc. SCIA 2001