Corpus Callosum Analysis using MDL-based Sequential Models of Shape and Appearance

Mikkel B. Stegmann[a,c], Rhodri H. Davies[b], Charlotte Ryberg[c]

[a]Informatics and Mathematical Modelling, Technical University of Denmark,
Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark

[b]Centre for Neuroscience, Howard Florey Institute,
University of Melbourne, VIC 3010, Australia

[c]Danish Research Centre for Magnetic Resonance,
H:S Hvidovre Hospital, Kettegård Allé 30, DK-2650 Hvidovre, Denmark

ABSTRACT

This paper describes a method for automatically analysing and segmenting the corpus callosum from magnetic resonance images of the brain based on the widely used Active Appearance Models (AAMs) by Cootes et al. Extensions of the original method, which are designed to improve this specific case are proposed, but all remain applicable to other domain problems. The well-known multi-resolution AAM optimisation is extended to include sequential relaxations on texture resolution, model coverage and model parameter constraints. Fully unsupervised analysis is obtained by exploiting model parameter convergence limits and a maximum likelihood estimate of shape and pose. Further, the important problem of modelling object neighbourhood is addressed. Finally, we describe how correspondence across images is achieved by selecting the minimum description length (MDL) landmarks from a set of training boundaries using the recently proposed method of Davies et al. This MDL-approach ensures a unique parameterisation of corpus callosum contour variation, which is crucial for neurological studies that compare reference areas such as rostrum, splenium, et cetera. We present quantitative and qualitative results that show that the method produces accurate, robust and rapid segmentations in a cross sectional study of 17 subjects, establishing its feasibility as a fully automated clinical tool for analysis and segmentation.

Keywords: registration, atlases, deformable models, active appearance models, MDL landmarking, corpus callosum, brain MRI

1. INTRODUCTION

Corpus callosum is the nervous tissue that connects the two cerebral hemispheres of the human brain. Many neurological studies indicate that the size and shape of the corpus callosum are related to gender, age, neurodegenerative diseases et cetera.[1–3] The gold standard for such morphometry studies is magnetic resonance imaging, which allows acquisition of accurate images of the anatomy and function of the human brain. However, obtaining manual tracings of the corpus callosum is both time-consuming, error-prone and operator dependent. Instead, medical image analysis should aim at replacing this task with fully automated and efficient methods eliminating subjectivity. In this paper we demonstrate that the generic Active Appearance Models (AAMs)[4,5] can be adapted to corpus callosum analysis. The efficacy of this approach will be assessed thorough a detailed quantitative and qualitative validation.

Contrary to manually defined contours, AAMs provide a reference coordinate system for every new example. This makes propagation of reference areas such as the rostrum, genu, truncus, isthmus and splenium a trivial matter. Further, regression analyses upon gender, age, motor abilities, lifestyle et cetera can be carried out directly since shape variation and size are both sensibly and compactly encoded.

This paper is accompanied by a technical report[6] containing background information, preliminary results and elaborations omitted here.

Corresponding author is M. B. Stegmann, E-mail: mbs@imm.dtu.dk, Web: http://www.imm.dtu.dk/~mbs/.

2. RELATED WORK

Many attempts on segmenting the corpus callosum have been carried out. However, the majority only gives qualitative results based on one or a few images. Below, recent contributions, which are either related to our approach, give quantitative results, or do both are chronologically summarised. All performance measures mentioned in this section, e.g. pt.pt./pt.crv., are described in Section 5.

Staib and Duncan[7] introduced a deformable model, which was based on a Fourier decomposition of object boundaries. Case studies were carried out on synthetic images and two MRI cases, one being a single mid-sagittal corpus callosum image. No quantitative results were given.

Székely et al.[8] developed shape models also based on a Fourier parameterisation and applied these for MRI corpus callosum segmentation in a case study carried out on 30 subjects. The models were initialised manually using an operator-specified AC-PC line as reference. No quantitative results were given.

Cootes et al.[9] proposed to implement deformable anatomical atlases using AAMs in a case study of 72 brain MRIs focusing on the ventricles, the caudate nucleus and the lentiform nucleus. A leave-one-out cross-validation using manual initialisation produced an average accuracy of 2.4 pixels pt.pt. and 1.2 pixels pt.crv.

Lundervold et al.[1] presented a method for MRI corpus callosum segmentation demonstrated on 10 subjects. The method consisted of a combination of a shape model and a multi-spectral tissue classification. The latter was based on three MR pulse sequences; FLASH, DESS, and FISP. Further, thresholding and a morphological cleaning step were needed. The method required no manual interaction and it was concluded that it was promising for clinical use, although no quantitative results were given.

Brejl and Sonka[10] reported a fully automated method for corpus callosum segmentation in MRI. An initial coarse position was determined using a shape-variant Hough transform based on an annotated training set. Refinement was carried out using i) snakes[11] or, ii) a dynamic programming approach, based on a cost function automatically designed from the training set. Trained on 15 subjects and tested on 75 subjects yielded an average segmentation error of 0.7 pixels semi-pt.crv. and 1.2 pixels semi-pt.crv. using either refinement method, respectively.

Van Ginneken et al.[12] augmented Active Shape Models (ASMs)[13] with a non-linear NN-classifier based on sequential forward and backward selection of moments of histograms over multiple scales. The method was evaluated on the data set in[10] using hold-out, i.e. 45 images for training and 45 images for testing. A standard ASM and the non-linear ASM yielded average shape overlaps of 0.62 and 0.81, respectively. Fitting the shape model directly to the ground truth yielded an upper bound of 0.89.

3. DATA MATERIAL

The data material used in this study comprises 17 cross-sectional, mid-sagittal magnetic resonance images (MRI) of the human brain obtained from 17 subjects using a whole body MR scanner (Siemens). MR pulse sequence: MPRAGE; matrix size: 256×256 pixels; subject age: 66/83/77 years (min/max/median).

Corpora callosa were annotated manually by drawing one closed, landmark-free contour in each image. These contours are shown in Figure 1, compensated for differences in size and orientation. The corresponding images are shown elsewhere.[6]

Figure 1. Aligned corpus callosum annotations from 17 mid-sagittal MRI, shown row-wise with the Procrustes mean shape in the lower right corner.

4. METHODS

4.1. Active Appearance Models

Active Appearance Models (AAMs)[4,5] are generative models capable of synthesising images of a given object class. By estimating a compact and specific basis from a training set, model parameters can be adjusted to fit unseen images and hence perform image interpretation. The modelled object properties are shape and pixel intensities (called *texture*). Training objects are defined by marking up each example image with points of correspondence. Variability is modelled by means of principal component analyses (PCA). Prior to PCA modelling shapes are Procrustes aligned and textures are warped into a shape-free reference frame and sampled. Let there be given P training examples, and let each example be represented by a set of N landmark points and M texture samples. Let $\mathbf{s}$ and $\mathbf{t}$ denote a synthesised shape and texture and let $\bar{\mathbf{s}}$ and $\bar{\mathbf{t}}$ denote the corresponding means. New instances can be generated by adjusting the PC scores, $\mathbf{b}_s$ and $\mathbf{b}_t$ in

$$\mathbf{s} = \bar{\mathbf{s}} + \mathbf{\Phi}_s \mathbf{b}_s \tag{1}$$

and

$$\mathbf{t} = \bar{\mathbf{t}} + \mathbf{\Phi}_t \mathbf{b}_t, \tag{2}$$

where $\mathbf{\Phi}_s$ and $\mathbf{\Phi}_t$ are eigenvectors of the shape and texture dispersions estimated from the training set. To recover any correlation between shape and texture and to obtain a combined parameterisation, $\mathbf{c}$, the values of $\mathbf{b}_s$ and $\mathbf{b}_t$ over the training set are combined in a third PCA,

$$\left[\begin{array}{c} \mathbf{W}_s \mathbf{\Phi}_s^{\mathrm{T}}(\mathbf{s} - \bar{\mathbf{s}}) \\ \mathbf{\Phi}_t^{\mathrm{T}}(\mathbf{t} - \bar{\mathbf{t}}) \end{array} \right] = \left[\begin{array}{c} \mathbf{W}_s \mathbf{b}_s \\ \mathbf{b}_t \end{array} \right] = \left[\begin{array}{c} \mathbf{\Phi}_{c,s} \\ \mathbf{\Phi}_{c,t} \end{array} \right] \mathbf{c} = \mathbf{\Phi}_c \mathbf{c}. \tag{3}$$

where $\mathbf{W}_s$ denotes a diagonal matrix weighting pixel distances against intensities. Synthetic examples, parameterised by $\mathbf{c}$, are generated by

$$\mathbf{s} = \bar{\mathbf{s}} + \mathbf{\Phi}_s \mathbf{W}_s^{-1} \mathbf{\Phi}_{c,s} \mathbf{c} \tag{4}$$

and

$$\mathbf{t} = \bar{\mathbf{t}} + \mathbf{\Phi}_t \mathbf{\Phi}_{c,t} \mathbf{c} \tag{5}$$

and rendered into an image by warping the pixel intensities of $\mathbf{t}$ into the geometry of the shape $\mathbf{s}$.

An AAM is matched to an unseen image using a least-squares criterion by an iterative updating scheme. This is based on a fixed Jacobian estimate,[14] or, originally, a principal component regression.[5] For this work we have chosen the former approach, which is treated in further detail elsewhere.[15] For further details on AAMs refer to.[5,14,16]

4.2. Landmark Placement

Establishing points of correspondence, so-called *landmarks*, on training examples by hand can not only be time-consuming, tedious and error-prone, but in some cases directly impossible when no "good" anatomical points are available. In this matter the term good refers to if landmarks can be identified with high degrees of reproducibility and repeatability.

We manually identify two landmarks on the corpus callosum; one at the rostrum and one at the splenium. Further, 78 semi-landmarks were interpolated along the contour using a uniform arc-length parameterisation.

To supplement this, a method for automatic landmarking proposed by Davies et al.[17,18] was applied. This is a minimum description length (MDL) framework that – in the spirit of Occam's razor – selects the landmarks, which lead to the minimum description length of the Point Distribution Model (PDM)[13] built on the training set. Procrustes alignment was used to obtain an initial alignment and this was optimised using the MDL objective function. Landmarks with a minimal description length were calculated for all shapes and the alignment was inverted so that the MDL-based landmarks were given in image coordinates.

4.3. Background Awareness

In medical images, organs and structures typically have a very homogenous appearance. Since AAMs only model the object appearance, one can find that the cost function easily can have a minimum inside the object, where both a plausible shape and good texture fit can be obtained. Unfortunately, corpora callosa are no exception to this. A partial remedy for this is to model the convex hull of the shape points. However, this only helps where the shape is concave. Further, this may lead to inclusion of unwanted texture variation, e.g. the appearance of the fornix brain structure in our case.

Consequently, we propose a simple scheme for modelling background variation by adding normals outwards from the shape, much inspired from the precursor to AAMs; the Active Shape Models.[13]   These normals are denoted *whiskers* and added implicitly during texture sampling with a scale relative to the current shape size. Texture samples obtained by sampling along whiskers are concatenated to the texture vector with a fixed and uniform weight compared to the conventional AAM texture samples. This provides a simple weighted method for adding background awareness to an AAM.

4.4. Sequential Relaxation of Model Constraints

To improve both convergence and robustness AAM search can be performed at different resolutions. A widely used scheme is to recursively decimate the input image thereby creating a pyramid representation of images at different scales. This can be carried out on the training set, thus producing one model for each level of the pyramid, see e.g. ref.[16]   The details of how this is implemented in the current framework is given and extended below.

Typically, model searches are initialised at the top level and propagated toward the full image resolution. Let $\mathbf{c}_i$ denote the converged model configuration at level i. The corresponding shape, $\mathbf{s}_i$ is propagated to level $i+1 = j$ by $\mathbf{s}_j = 2\mathbf{s}_i$ and the texture is sampled into $\mathbf{t}_j$. Assuming that the optimisation at level i produced an actual match, $\mathbf{c}_j$ can now be estimated by projecting $\mathbf{s}_j$ and $\mathbf{t}_j$ into the corresponding combined eigenspace

$$\hat{\mathbf{c}}_j = \mathbf{\Phi}_{c,j}^{\mathsf{T}} \left[\begin{array}{c} \mathbf{W}_{s,j} \mathbf{\Phi}_{s,j}^{\mathsf{T}}(\mathbf{s}_j - \bar{\mathbf{s}}_j) \\ \mathbf{\Phi}_{t,j}^{\mathsf{T}}(\mathbf{t}_j - \bar{\mathbf{t}}_j) \end{array} \right] \tag{6}$$

and model optimisation can continue at level j. If $\mathbf{\Phi}_{s,i} = \mathbf{\Phi}_{s,j} \ \forall \ i,j$ then $\mathbf{b}_{s,i}$ can be used to ease the calculation of $\hat{\mathbf{c}}_j$, but more importantly; $\mathbf{s}_j$ and $\mathbf{t}_j$ would lie in (or at least be very close to) the space spanned by $\mathbf{\Phi}_{c,j}$. This means that $\mathbf{t}_j$ can be used directly to calculate the first update of $\hat{\mathbf{c}}_j$.

This above scheme represents a sequential relaxation of the constraints imposed by the texture model. At the top of the pyramid, i.e. the smallest image, only a weighted sum of the pixels in the original image is modelled, e.g. the mean of 16 pixels if a box filter and a three level pyramid is used. This is thus a more constrained model compared to the largest model, potentially capable of modelling 16 different pixel values.

We propose to generalise this coarse-to-fine resolution strategy by using the propagation scheme above to traverse sequences of AAMs designed to increase model flexibility. Means for doing so include i) gradual inclusion of more deformation modes in the shape, texture and combined models, ii) independent shape and texture models, iii) de-coupling, or exclusion of shape areas, e.g. omitting the convex hull, modelling only the actual shape area.

4.5. Initialisation using Pose Priors

Exploiting the convergence radius of the AAM search carries out model initialisation in this study. Within this radius the AAM will converge to a plausible fit. Parameters with variation outside this radius over the training set are semi-exhaustively searched, with a grid spacing less than twice the radius, by starting the AAM search at each grid point. Each search result represents a potential initialisation candidate. To add robustness this is implemented in a candidate scheme, where several candidates compete in an evolutionary process to become the initial configuration. See[19] for the details of our initialisation algorithm.

As the deformation, size, and orientation variation in our data set are within the convergence radius, the initialisation only needs to search the translational parameters (x, y) using the maximum likelihood estimate of the corpus callosum, i.e. the mean shape and appearance, having mean size and orientation. Further, this search is restricted by the bounding rectangle of all training shapes.

5. EXPERIMENTAL RESULTS

To assess the two landmarking methods, the first three principal modes of shape variation are shown in Figure 2. Their shares of the total shape variance were 33%, 30% and 29% (manual) and 35%, 25% and 15% (MDL). As these are ordered by decreasing variance it might surprise that the first manual mode seems to deliver less shape change than the second. The cause is clearly illustrated in Figure 3, which shows that variation in the manual case almost solely stems from landmark movements *along* the contour. However, in the MDL case deformations are, by design, much closer to being orthogonal to the contour. Figure 2 clearly illustrates that shape modes are *not* canonical, unless a very strict and well-defined landmarking scheme is used. The MDL shape model used only 27% of variance required for the manual shape model to represent the same contour variation. Thus, only MDL-based shapes are considered in the following.

Having established a specific and compact shape deformation basis that ensures correspondence of contour subparts between shapes in a well-defined manner, we tested the ability to actively recover these correspondences in unknown images by running an AAM search using the described initialisation method. Each test was carried out using leave-one-out cross-validation. Performance measures were the average distance between i) model and ground truth landmarks (pt.pt.), and ii) model landmark and ground truth contour (pt.crv.). Further, the shape overlap used in[12] has been calculated as the ratio between the ground truth area and the union of the ground truth area and the segmentation result area. The semi.pt.crv. measure used in[10] is the distance between each model landmark and the nearest ground truth landmark. Thus having, pt.crv. and pt.pt. as lower and upper bound, respectively. Texture samples from whiskers influenced the texture model by two thirds. The whisker length was three times the distance between landmark 1 and 2 on the mean shape. Shape, texture and combined eigenspaces were truncated to span 95% of the respective variation.

| $-2.5\sigma_i$ | mean | $2.5\sigma_i$ | | $-2.5\sigma_i$ | mean | $2.5\sigma_i$ |
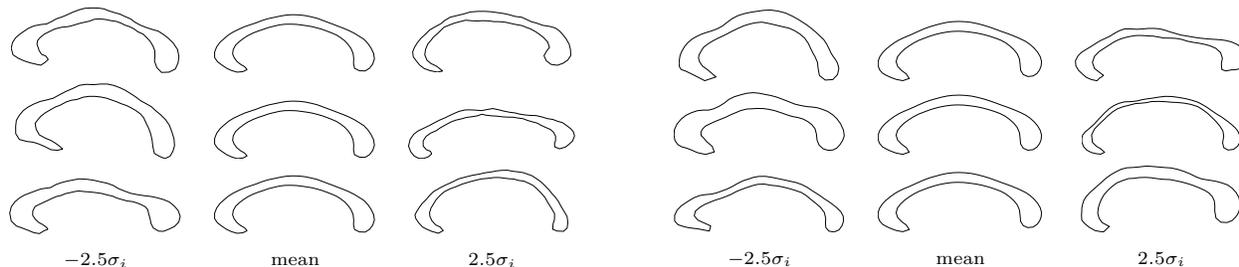
Figure 2. The three largest shape deformation modes of $\mathbf{\Phi}_s$ (top-down), manual (left), MDL (right).
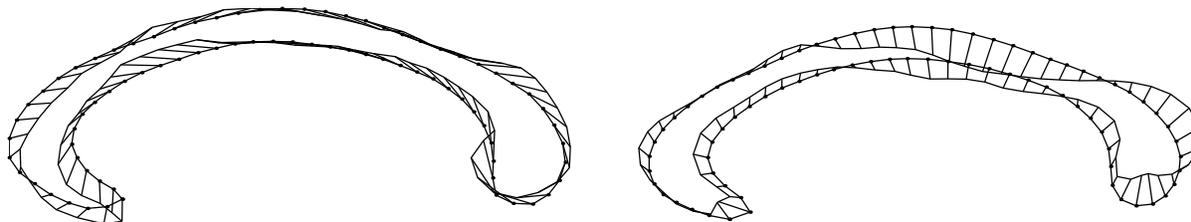
Figure 3. Displacement vectors of the largest shape mode in $\mathbf{\Phi}_s$, manual (left), MDL (right).

Table 1. Segmentation Results

Case	Whiskers added	Fine-tuning	AAM Seq.	Pt.pt. [pixels]		Pt.crv. [pixels]		Shape overlap	
				Mean	Median	Mean	Median	Mean	Median
A	No	No	I	2.08±0.32	1.95	0.87±0.07	0.80	0.69±0.02	0.68
B	Yes	No	I	1.24±0.11	1.09	0.73±0.06	0.71	0.76±0.01	0.75
C	Yes	Yes	I	1.14±0.08	1.14	0.68±0.04	0.68	0.76±0.01	0.75
D	Yes	Yes	II	1.13±0.08	1.14	0.67±0.04	0.68	0.76±0.01	0.76
E	Yes	Yes	III	1.07±0.06	1.08	0.60±0.03	0.62	0.78±0.01	0.78
F	Yes	Yes	IV	1.09±0.06	1.11	0.60±0.03	0.60	0.78±0.01	0.79

Sequence I: Size 1:1x1 (one model pixel corresponds to one image pixel);convex hull
Sequence II: Size 1:4x4;convex hull, Size 1:2x2;convex hull, Size 1:1x1;convex hull
Sequence III: (as seq. II), Size 1:1x1;no convex hull
Sequence IV: (as seq. III), Size 1:1x1;no convex hull;no combined model
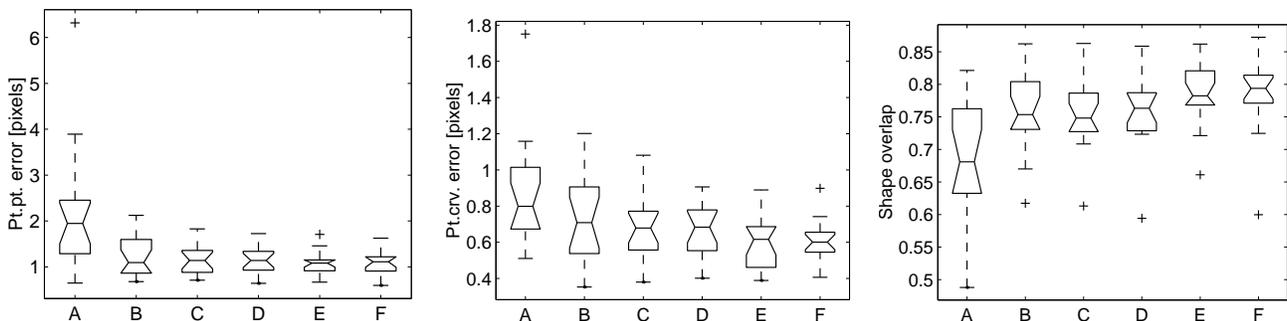
Figure 4. Box plots of pt.pt., pt.crv., and shape overlap for each case in Table 1. Whiskers are 1.5 IQR at maximum.

To improve the AAM search result, the general-purpose optimiser BFGS (see e.g.[20]) was used in case C–F. The cost function was identical to the one used by the AAM search, namely the least squares fit of the model and image texture. AAM sequence designs and results are given in Table 1 and shown as box plots in Figure 4. Std. errors are symbolised using ±. The pose variations over the data set from the average were 86/130% in size and -13/19 degrees in orientation (min/max). A qualitative impression of the segmentation accuracy is given in Figure 5. Images are cropped for display, and do not show the full area searched by the initialisation algorithm.

Fitting the shape model directly to the ground truth shape points gave a lower bound on the pt.pt. accuracy in case F: 0.37/0.84/0.63 pixels (min/max/median). Another lower bound is estimated by a leave-all-in evaluation. On average in case F, this yielded 0.28 pixels pt.pt., 0.18 pixels pt.crv., and a shape overlap of 0.93. In an optimal optimisation process using shape and texture models spanning 100% variation, these would be zero, zero and one.

All results have been generated using an extended version of our public domain C++ AAM implementation.[15] The computation time incl. initialisation in case F was 1.6 seconds on average for each image using a 2.4 GHz Pentium 4 PC. This could easily be reduced to 0.5 seconds by optimisation of the extensions.

6. DISCUSSION

From Table 1 and Figure 4 it is observed that the addition of whiskers markedly improved the segmentation accuracy. The general-purpose optimisation and the AAM sequences brought a modest improvement in accuracy. However, sequence II–IV brought major improvements in initialisation time, due to the sixteen times smaller
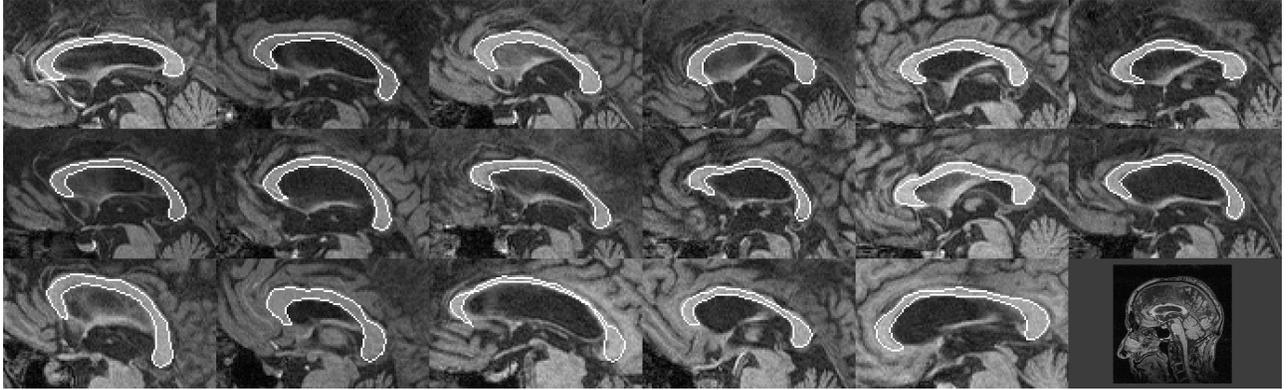
Figure 5. Segmentation results, case F (cropped). First image without cropping (lower-right).

texture model used for initialisation. Qualitatively, Figure 5 shows a very accurate match of each model to the corpus callosum. Quantitatively, these matches are on average only half a pixel from the lower bound. Observation 15 is slightly inaccurate, giving rise to outliers and whiskers in Figure 4. The limitations of the shape model are noticeable as subtle discrepancies in observation 3, 11 and 17, all showing small bumps not reachable by the shape model due to fairly limited number of training shapes. However, considering the large variations present in the data set, we consider the leave-one-out shape model performance impressive.

Alas, due to differences in data sets, comparisons to the related work are not possible. However, this study indicates that our approach easily compares with respect to segmentation accuracy, judging from training set sizes and performance measures.

This study is the only one known to the authors, which establishes compact, well-defined and unique corpus callosum contour correspondences. Shape models based on MDL optimality ensure this. Per se, this uniqueness is a prerequisite for, e.g. proper calculation of splenium areas, and meaningful morphometric regression analysis upon the recovered shape parameters, $\mathbf{b}_s$ in unseen images. Further, only one MR acquisition is required in the presented method, contrary to the work presented in.[1] Finally, this is also the first study that explores the (obvious) combination of MDL shapes models and Active Appearance Models.

7. CONCLUSION

We have presented a both generic and holistic framework, which optimises well-defined measures based on well-described methods, thus avoiding a composition of various techniques, e.g. tissue classification, thresholding, morphologic operators, et cetera. The problem of model initialisation is addressed rendering the presented approach fully automated. Using our freely available AAM implementation, the segmentation time is vastly superior to the reviewed methods, allowing for interactive use, fast analysis of retrospective studies, et cetera.

We conclude that AAMs with the proposed extensions, and in conjunction with MDL-based landmarking, have a great potential as a clinical tool for rapid, accurate and fully automated analysis and segmentation of the corpus callosum in MRI.

Acknowledgements

REFERENCES

1. A. Lundervold, N. Duta, T. Taxt, and A. K. Jain, "Model-guided segmentation of corpus callosum in MR images," in *Computer Vision and Pattern Recognition*, IEEE Comput. Soc, 1999.

2. F. L. Bookstein, "Landmark methods for forms without landmarks: localizing group differences in outline shape," *Medical Image Analysis* **1**(3), pp. 225–244, 1997.

3. A. Dubb, B. Avants, R. Gur, and J. Gee, "Shape characterization of the corpus callosum in schizophrenia using template deformation," in *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, **2**, pp. 381–388, 2002.

4. G. J. Edwards, C. J. Taylor, and T. F. Cootes, "Interpreting face images using active appearance models," in *Proc. 3rd IEEE Int. Conf. on Automatic Face and Gesture Recognition*, pp. 300–5, IEEE Comput. Soc, 1998.

5. T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," in *Proc. European Conf. on Computer Vision*, **2**, pp. 484–498, Springer, 1998.

6. M. B. Stegmann and R. H. Davies, "Automated analysis of corpora callosa," Tech. Rep. IMM-REP-2003-02, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, mar 2003.

7. L. H. Staib and J. S. Duncan, "Boundary finding with parametrically deformable models," *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **14**(11), pp. 1061–1075, 1992.

8. G. Szekely, A. Kelemen, C. Brechbuhler, and G. Gerig, "Segmentation of 2-d and 3-d objects from mri volume data using constrained elastic deformations of flexible fourier contour and surface models," *Medical Image Analysis* **1**(1), pp. 19–34, 1996.

9. T. F. Cootes, C. Beeston, G. J. Edwards, and C. J. Taylor, "A unified framework for atlas matching using active appearance models," in *Information Processing in Medical Imaging. IPMI'99*, pp. 322–33, Springer-Verlag, 1999.

10. M. Brejl and M. Sonka, "Object localization and border detection criteria design in edge-based image segmentation: automated learning from examples," *Medical Imaging, IEEE Transactions on* **19**(10), pp. 973–985, 2000.

11. M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. Jour. of Computer Vision* **8**(2), pp. 321–331, 1988.

12. B. van Ginneken, A. F. Frangi, J. J. Staal, B. M. Ter Haar Romeny, and M. A. Viergever, "Active shape model segmentation with optimal features," *IEEE Transactions on Medical Imaging* **21**(8), pp. 924–933, 2002.

13. T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models – their training and application," *Computer Vision and Image Understanding* **61**(1), pp. 38–59, 1995.

14. T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," *IEEE Trans. on Pattern Recognition and Machine Intelligence* **23**(6), pp. 681–685, 2001.

15. M. B. Stegmann, B. K. Ersbøll, and R. Larsen, "FAME – a flexible appearance modelling environment," *IEEE Trans. on Medical Imaging* **22**(10), pp. 1319–1331, 2003.

16. T. F. Cootes and C. J. Taylor, *Statistical Models of Appearance for Computer Vision*, Tech. report, University of Manchester, 2001.

17. R. H. Davies, *Learning Shape: Optimal Models for Analysing Natural Variability.* PhD thesis, Department of Imaging Science and Biomedical Engineering, University of Manchester, Stopford building, University of Manchester, 2002.

18. R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor, "A minimum description length approach to statistical shape modeling," *Medical Imaging, IEEE Transactions on* **21**(5), pp. 525–537, 2002.

19. M. B. Stegmann, "Object tracking using active appearance models," in *Proc. 10th Danish Conference on Pattern Recognition and Image Analysis, Copenhagen, Denmark*, **1**, pp. 54–60, DIKU, 2001.

20. R. Fletcher, *Practical Methods of Optimization*, John Wiley & Sons, 1987.