

Non-Rigid Registration Based Active Appearance Models for 3D Medical Image Segmentation

Jan Klemencic*

[†]Faculty of Electrical Engineering, University of Ljubljana, Ljubljana, Slovenia

Josien P. W. Pluim and Max A. Viergever

Image Sciences Institute, University Medical Center Utrecht, Utrecht, The Netherlands

Hugo G. Schnack and Vojko Valencic[†]

Department of Psychiatry, University Medical Center Utrecht, Utrecht, The Netherlands

Active shape models and active appearance models are getting increasingly popular in medical image segmentation applications. However, they are not suitable for three-dimensional (3D) images in their original form. This is due to the underlying shape representation (a point distribution model, PDM), which becomes impractical in 3D. Recently, it was shown that nonlinear registration algorithms can assist in the automatic creation of a 3D PDM. Based on this idea, we built a 3D active appearance model of brain structures. The model extracts the mean texture and the image deformation variation information from the training set of images. A special benefit is the inclusion of an extended region of interest into the model, making it suitable for segmentation of structures with poorly defined edges. We evaluated the model by applying it to the task of automatic segmentation of the hippocampi from magnetic resonance brain images. We found high accuracy of the model, which is comparable to the accuracy of the underlying registration method. The main benefit of the model-based segmentation over the registration-based segmentation is time, which is reduced from many hours (for registering an atlas to the image) to only a few minutes (for fitting the model to the image).

Journal of Imaging Science and Technology 48: 166–171 (2004)

Introduction

The Active Shape Models (ASM)^{1,2} and Active Appearance Models (AAM)³ are well known and widely applicable algorithms for automatic image segmentation. At the core of the ASMs' functionality is their ability to process the pre-segmented training set of images and learn the mean shape of the Structure Of Interest (SOI), as well as the plausible variations of the mean shape. The shape is defined in terms of the Point Distribution Model (PDM),¹ which is a set of landmarks, placed on the SOI boundary. A key requirement for a functional ASM is the spatial correspondence of the landmarks over the training set (each SOI from the training set must be annotated by the same number of landmarks, and each landmark on a particular SOI must have a spatially corresponding landmark placed on each of the rest

of the training SOIs). The AAMs are similar to ASMs, but they also learn the mean SOI texture and its variability from the training image set. AAMs also use a different algorithm for fitting the model to the unknown image, in order to segment it.

Present day three-dimensional (3D) medical image acquisition techniques generate large amounts of 3D medical image data, which is impractical to objectively analyze without segmentation. One of the commonly used strategies for automatic 3D medical image segmentation is the *segmentation-by-registration* (SBR) paradigm. It involves a representative, accurately segmented image of the anatomy, called the *anatomical atlas*. The atlas segmentation usually adopts the form of binary or labeled images. The segmentation of the unknown patient image is carried out by registering the intensity atlas image to the patient image, and deforming the segmentation image along the way into the patient's shape, thus achieving segmentation.

Both ASMs and AAMs could be used as alternatives to SBR, avoiding some of the inherent drawbacks of the SBR methods (highly time consuming procedure, no *a priori* information used). However, most of today's ASM and AAM implementations are targeted at segmenting two-dimensional images only. While direct translation of the ASM and AAM concepts into 3D is straightforward,

Original manuscript received June 9, 2003

* jan.klemencic@fe.uni-lj.si

©2004, IS&T—The Society for Imaging Science and Technology

ward in theory, it becomes prohibitively impractical for real applications. The reason lies in the underlying shape representation, i.e., the PDM, and the related requirement for spatial landmark correspondence over the training set of images. Specifically, it is almost impossible for the operator (anatomical expert) to place a sufficient number of landmarks on each of the 3D SOIs from the training image set, and achieve satisfactory spatial landmark correspondence along the way. This is especially true for amorphous shapes, such as many biological structures (tissues, organs) with rounded boundaries. To solve this problem, an automatic method for placing spatially correspondent landmarks on a set of (manually) pre-segmented 3D SOIs is needed. This can be achieved by means of a nonlinear registration method: registering all the images of the training set to the atlas, i.e., warping them into a common shape, distributing a sufficient number of landmarks on each SOI in the common shape, and reverting all the warps (now including landmarks).⁴ This gives a set of SOIs with spatially corresponding landmarks. The degree of landmark correspondence is directly dependent on the accuracy of the registration method.

However, using the registration algorithm for achieving landmark correspondence can be brought one step further. Each of the training SOI registrations results in a transformation, defined as a vector field of voxel displacements. Having the displacement data for any chosen voxel, we can determine the spatially corresponding voxel in each of the remaining training images. Therefore, we can think of each voxel as a sort of a landmark. We can build the PDM, using voxels as landmarks. Such a PDM is quite different from the classic PDM: it does not represent the shape of the SOI anymore; rather, it represents the *shape of the image*. Instead of the mean SOI shape and its variations, the ASM or AAM then contains information about the mean deformation field (or mean warp) and its variations. By building the models in such fashion, more shape information is extracted from the registrations of the training images. The SOI shape, which is necessary for performing segmentation, can still be built into the model in the form of a binary image, which can always be warped along with the base intensity image.

A statistical shape model, based on a PDM as described above, was presented by Rueckert et al.⁵ Although demonstrative as a proof-of-principle, their implementation lacks the actual segmentation algorithm. An AAM, following similar principles, has been reported by Duchesne et al.,⁶ together with the full segmentation application, segmenting the medial temporal lobe structures (such as the hippocampus) from magnetic resonance images (MRI). The underlying registration method of Duchesne's model is ANIMAL (Automatic Nonlinear Image Matching and Anatomical Labeling).⁷ Our model (the building of which we already outlined in Klemencic et al.⁸) is in many ways similar to Duchesne's model. Our contributions include: a) building a more compact, fixed texture 3D AAM, b) building and evaluation of the model, based on a more accurate nonlinear registration method (confirming the consequent higher accuracy of our model), and c) some additional informative experiments, providing insight into the model's properties.

Data

To evaluate the performance of our model, we tested it on MRI brain images and used the right hippocampus

as the SOI. Our data consisted of 28 images: T1-weighted 3D fast field echo (FFE) scans with 160–180 1.2 mm contiguous coronal slices (TE = 4.6 ms; TR = 30 ms; flip angle = 30 degrees; FOV = 256 mm/80%; in-plane voxel sizes $1 \times 1 \text{ mm}^2$) of the whole heads of adult, healthy persons, obtained on a Philips 1.5T MR scanner. For each T1 image, a corresponding binary image was available, depicting hippocampi segmentations. All the images were provided by the Department of Psychiatry, University Medical Center Utrecht, The Netherlands. The segmentations were performed manually by a trained expert at the same institution.⁹

Methodology

Pre-Processing

Before building the model, all the training images were subject to four pre-processing steps: shading correction,¹⁰ intensity normalization,¹¹ affine registration¹² and extraction of Region Of Interest (ROI). The pre-processed images were used for building the model, and also as test images for evaluating the model performance via hippocampus segmentation.

The affine registration is needed in order to align all the images globally to the common atlas. To save on computer resources, it is sensible to model only a smaller ROI, not the entire head images. Modeling a relatively small region around the SOI also makes the models more specific. Our ROI is a block of voxels, centered on the SOI (hippocampus), $92 \times 92 \times 88$ voxels in size.

Creating a 3D AAM

As indicated in the introduction, a high-dimensional nonlinear registration method is needed as the foundation for our 3D PDM. We used the method of Rueckert et al.¹² It features a regular grid of control points (nodes), distributed within the image volume, and displaced during registration. Voxel displacements are computed via the B-spline interpolation of node displacements. Normalized mutual information is used as the voxel-based similarity measure. We will term this registration method a *free-form deformation registration* (FFDR).

In order to create a 3D PDM, describing the complete 3D warp, the nodes of the FFDR can be used as PDM landmarks. The nodes are evenly distributed within the ROI of the atlas image, forming a rectangular mesh. Therefore, the pertinent PDM models the whole ROI, not only the SOI. Based on such a PDM, the corresponding 3D ASM can be built. The pose correction (as in original ASMs) is not necessary here, since all the images are affinely aligned before training. Similarly, global intensity normalization parameters are unnecessary, due to the pre-processing intensity normalization step. We will briefly review the model-building mathematical treatment in the following paragraphs.

Each deformation field can be decomposed into three orthogonal components of node displacement data, x , y , and z . The components can be concatenated into a single $3n$ -dimensional column vector \mathbf{s} , where n is the number of nodes in the deformation field.

A set of N training ROIs gives a set of N deformation fields \mathbf{s}_i ($i = 1, 2, \dots, N$), linking each ROI to the atlas ROI. The set of N deformation fields can be thought of as a cloud of N points in the $3n$ -dimensional space. By moving around within this cloud of points, new deformation fields can be generated, which are broadly similar to those from the training set. Through statistical analysis (*principal component analysis*, PCA¹), the high dimensionality of this space can be reduced: the shape




Figure 1. The variance distribution over 27 variation modes of the statistical deformation field model: (left) The variances for each mode of variation; (right) Cumulative variance; the horizontal dashed line marks the 99% threshold of the total variance.

of the cloud can be modeled by relatively few model parameters. To achieve this, the first step is to compute the mean deformation field $\bar{\mathbf{s}}$ by

$$\bar{\mathbf{s}} = \frac{1}{N} \sum_{i=1}^N \mathbf{s}_i. \quad (1)$$

The mean deformation field is subtracted from each original deformation field and the $(3n \times 3n)$ covariance matrix \mathbf{S} of the data is created:

$$d\mathbf{s}_i = \mathbf{s}_i - \bar{\mathbf{s}}, \quad (2)$$

$$\mathbf{S} = \frac{1}{N} \sum_{i=1}^N d\mathbf{s}_i d\mathbf{s}_i^T. \quad (3)$$

Then, the unit eigenvectors \mathbf{p}_k ($k = 1, 2, \dots, 3n$) and the corresponding eigenvalues λ_k of \mathbf{S} are computed and sorted, such that $\lambda_k \geq \lambda_{k+1}$. The eigenvectors corresponding to the largest eigenvalues describe the principal axes of the data cloud. In effect, they describe the principal modes of variation of the deformation field.

Most of the variation can usually be explained by a small number of modes t , where $t \ll 3n$. The value of t is chosen such as to explain a given proportion of the total variance exhibited in the training set (in our experiments, we used 99% of the total variance).

Collecting the eigenvectors \mathbf{p}_k in a matrix $\mathbf{P}_s = (\mathbf{p}_1 | \mathbf{p}_2 | \dots | \mathbf{p}_t)$, any deformation field can be approximated by

$$\mathbf{s} = \bar{\mathbf{s}} + \mathbf{P}_s \mathbf{b}_s, \quad (4)$$

where $\mathbf{b}_s = (b_1, b_2, \dots, b_t)^T$ is a t -dimensional vector of model parameters, given by

$$\mathbf{b}_s = \mathbf{P}_s^T (\mathbf{s} - \bar{\mathbf{s}}). \quad (5)$$

By varying the elements of \mathbf{b}_s , the model can generate synthetic deformation fields, using Eq. (4). The variance of the i -th parameter b_i across the training set is given by λ_i . By limiting each parameter b_i to $\pm 3\sqrt{\lambda_i}$, i.e., three standard deviations around the mean, we ensure that the generated deformation field will be similar to those from the training set.

In the treatment above, PCA was used to build a compact statistical shape model (specifically, a statistical

deformation field model). Similarly, PCA can be applied to build a model of gray level intensities within the ROI, or statistical texture model. Instead of node displacement data (vector \mathbf{s}), voxel intensities are concatenated into the PCA input vector, denoted by \mathbf{g} . The dimensionality of \mathbf{g} now corresponds to the number of voxels within the ROI, while the rest of the mathematical treatment remains the same. The resulting model can generate synthetic, but plausible textures, using the texture equivalent of Eq. (4):

$$\mathbf{g} = \bar{\mathbf{g}} + \mathbf{P}_g \mathbf{b}_g, \quad (6)$$

where $\bar{\mathbf{g}}$ is the mean texture vector, \mathbf{P}_g is a matrix, describing the modes of texture variation (eigenvectors of the covariance matrix), and \mathbf{b}_g is a set of texture model parameters. The complete statistical appearance model is created by combining the deformation field model and the texture model into a unified model.

Following the mathematical framework outlined above, we built the statistical deformation field model, based on the MRI data provided (we will show in the next section, that the statistical texture model is not needed for successful segmentations). Using the set of 28 training images, we obtained 27 modes of deformation field variation. Figure 1 shows the variance for each mode of variation. Regular as well as cumulative charts are given. In the regular chart, we can observe that considerable amount of variance is still present in the highest modes of variation. This indicates that our training set is not large enough to adequately cover the anatomical variability of the population. Consequently, our models are overly constrained (too specific). In the cumulative chart, we can observe that the 99% variance threshold is reached in the highest mode.

The middle coronal cross-section of the model is shown in Fig. 2 (first mode of variation) and Fig. 3 (second mode of variation). The middle columns show the model in the mean state. The left and right columns show the variations of the model (offsetting the pertinent parameter by ± 3 standard deviations around the mean). In the bottom row, the difference images towards the mean are given. In the difference images, it is easier to visually perceive and assess the (occasionally rather slight) variations.

Appearance-Based Segmentation: Using Variable or Fixed Texture

Model-based segmentation application requires an algorithm for fitting the model to the unknown image.




Figure 2. The first mode of the deformation field variation around the mean. The middle column shows the mean image (upper row). The left and right columns show the deformed images, obtained by adjusting the first model parameter by ± 3 standard deviations around the mean. The lower row gives the difference images between the mean image and the corresponding deformed image. On all images, the middle coronal cross-section of the 3D ROI block is shown.

There is a fundamental difference in the original ASM and AAM fitting algorithm implementation. The ASM tries to align itself to some chosen image features, found at the SOI boundaries. On the other hand, the AAM is able to generate complete synthetic images, based on the set of model parameters. The goal of the fitting algorithm is to minimize the difference between the generated image and the image to be segmented.

Usually, the parameters of the AAM vary its shape as well as its texture. We note however that the statistical deformation field model (as proposed above) is also capable of generating synthetic images. The model parameters control the generated deformation field, which can subsequently be used to deform the (fixed texture) mean anatomy image. Such a model is a hybrid between the original ASM and AAM paradigms. While it does not model the SOI texture as classical AAMs do, its overall functionality seems closer to the AAM than to the ASM, due to the fact that it still contains the mean SOI texture, and its ability to generate synthetic images. Consequently, also the AAM-like fitting algorithm can be implemented for such a model.

In view of the above, we will refer to our model as to the 3D AAM model. Duchesne's model⁶ is similar, but includes the texture variation information into the model. The imminent practical advantage of our model over the variable texture model is size, since modeling the texture variations requires a lot of computer resources, even for a relatively small ROI. Also, the texture variations are rather slight and therefore may not contribute much to the accuracy and efficiency of the automatic segmentation method.

We built and evaluated the fixed texture 3D AAM, proving that the fixed texture segmentation scheme is practical and feasible. We implemented a similar model fitting algorithm to the one proposed for 2D images^{3,13} and adapted it for 3D images. A short outline is given in the following section.

Appearance-Based Segmentation: Model Fitting

The goal of the model-fitting algorithm is to find such parameters, that the magnitude of the difference vector




Figure 3. The second mode of the deformation field variation around the mean. The middle column shows the mean image (upper row). The left and right columns show the deformed images, obtained by adjusting the second model parameter by ± 3 standard deviations around the mean. The lower row gives the difference images between the mean image and the corresponding deformed image. On all images, the middle coronal cross-section of the 3D ROI block is shown.

$\delta \mathbf{i}$ between the voxel intensities of the (pre-processed) unknown image and the voxel intensities of the generated image is minimized. Let \mathbf{b}_{itr} denote the vector of model parameters in a particular optimization iteration (a vector of zeros is used for the first iteration). Let \mathbf{b}_{opt} denote the vector of optimal model parameters (the goal of the optimization). Then, the error in model parameters for the pertinent iteration is $\delta \mathbf{b}_a = \mathbf{b}_{itr} - \mathbf{b}_{opt}$. A linear model, linking $\delta \mathbf{i}$ to $\delta \mathbf{b}_a$, can be constructed:

$$\delta \mathbf{b}_a = \mathbf{R} \delta \mathbf{i}, \quad (7)$$

where \mathbf{R} is a so-called prediction matrix. The computation of \mathbf{R} is time-consuming (around 9 minutes on a 1.6 GHz Pentium PC), but it can be approximately computed off-line (before the segmentation). In each optimization iteration, $\delta \mathbf{i}$ is first determined. Then, $\delta \mathbf{b}_a$ is computed using Eq. (7), and the model parameters are updated accordingly. As only an approximation of \mathbf{R} is available, a few such iterations (typically between 2 and 20) are usually necessary before the fit can not be improved anymore, and convergence is declared.

After convergence, the deformation field \mathbf{s}_{seg} , corresponding to optimal model parameters \mathbf{b}_{opt} is obtained. \mathbf{s}_{seg} is then used to deform the atlas segmentation of the SOI into the shape of the unknown image, thus segmenting it.

Results

We compared the 3D AAM segmentation method to the established SBR segmentation method (we used the FFDR for performing the SBR segmentation). The SOI was the right hippocampus. To assess the segmentation accuracy, manual segmentations were used as the ground truth with which to compare the automatic segmentation results. The same atlas of segmentations was used for all experiments, enabling an objective comparison of the approaches. To quantify the agreement between the segmentations, we applied a common binary-image similarity measure: $k = (2a)/(2a+b+c)$,

TABLE I. Results of the Two Segmentation Methods: SBR and AAM. Statistics of k Over all 28 Images is Given.

	SBR	AAM
mean k	0.84	0.80
std. dev. of k	0.03	0.05
maximum k	0.90	0.86
minimum k	0.76	0.68

where a is the number of intersecting, labeled voxels, b is the number of voxels only labeled automatically, and c is the number of voxels only labeled manually.

The results are listed in Table I. Mean and standard deviation of k as well as the maximum and minimum k over the 28 images is given. It is important to note that all 3D AAM experiments were performed on a leave-one-out basis, where the model was built from 27 images, excluding the segmented image (meaning that 28 distinct models were built, one for each of the 28 images). The processing time for 3D AAM segmentation was in the range of 1 to 3 minutes (excluding the affine registration), which is extremely low in comparison to the SBR algorithm (around 5 hours). All experiments were performed on a 1.6 GHz Pentium platform. Following these basic experiments, we performed some additional experiments to further assess the model.

First, we wanted to check how good the pre-computed approximation of the prediction matrix \mathbf{R} was. Therefore, we performed regular segmentations, re-computed the exact value of \mathbf{R} for the obtained model state, and used this value in an additional pass of the complete optimization procedure, thus refining the model fit. The results of these experiments are listed in the first column of Table II.

Secondly, we performed the segmentations using the model, built from all of the training images (including the segmented image). In contrary to all other, leave-one-out tests, we call these experiments *leave-all-in* experiments. The results are listed in the middle column of Table II.

In the third set of experiments, we used an alternative to the regular model fitting algorithm. Instead of the iterative parameter optimization, we computed the deformation parameters in a single step, using Eq. (5). This was possible because the deformation field \mathbf{s} was known in advance (from the model building stage) for each of the ROIs. Such approach is not feasible in a real segmentation application. The results of these experiments are listed in the last column of Table II.

Discussion

From Table I the SBR method appears slightly more accurate ($k = 0.84$) than the 3D AAM methods ($k = 0.80$). We feel that the slightly lower AAM accuracy is mainly due to the lack of a more extensive training set. Twenty seven images (in the case of leave-one-out experiments) are clearly not enough to cover the anatomical variability of our chosen ROI. This explanation is also supported by the variance distribution chart (Fig. 1) and by our further experiments (see below). Still, even by using a relatively small training set, we showed that the much faster model based method gives results almost as good as the underlying registration method.

By comparing the first column of Table II to the regular experiments in Table I, we see that the on-line re-

TABLE II. Results of the Additional AAM Experiments (for Details see Main Text). Statistics of k Over All 28 Images is Given.

	updated \mathbf{R}	leave-all-in	pre-computed parameters
mean k	0.80	0.87	0.79
std. dev. of k	0.04	0.03	0.06
maximum k	0.86	0.92	0.89
minimum k	0.69	0.79	0.61

computation of the prediction matrix \mathbf{R} only gives very slight improvements. We conclude that the pre-computed approximation of \mathbf{R} is good enough. The on-line recomputation is not economical, as it extends the model fitting time by around 9 minutes on a 1.6 GHz Pentium PC for the given ROI.

The leave-all-in experiments give very good results ($k = 0.87$; see the middle column of Table II). This confirms that the model fitting algorithm works well. We can again conclude that in order to further improve the regular (leave-one-out) results, a model with greater anatomical variability information, i.e., a larger training set, is needed.

The same conclusion can once again be drawn from the results with pre-computed model parameters (the right column of Table II). A k of only 0.79 indicates that the model is too restricted in shape. It is interesting to note that the iterative model fitting algorithm finds a better match than computing the parameters from a pre-computed deformation field.

Since Duchesne et al.⁶ use the same binary-image similarity measure k and also segment the hippocampus from MR images, a direct comparison of the results is possible. Duchesne builds a variable texture AAM (we will term it a VT-AAM), and uses ANIMAL⁷ as the underlying registration method. Their training set comprises 70 MR brain images. Duchesne reports an accuracy of around $k = 0.66$ for the VT-AAM method, and an accuracy of around $k = 0.70$ for the SBR method using ANIMAL. We note the considerably lower accuracy of Duchesne's VT-AAM experiments in respect to our 3D AAM experiments ($k = 0.66$ vs. $k = 0.80$), in spite of his much larger training set of images (70 vs. 27). We also note the considerably lower accuracy of SBR using ANIMAL in respect to SBR using FFDR ($k = 0.70$ vs. $k = 0.84$). We observe that in ours as well as in Duchesne's experiments, the accuracy of the model-based methods is slightly lower than the accuracy of the underlying registration method. We can conclude that the accuracy of the model-based methods depends on the accuracy of the underlying registration method to some extent. Apart from the lower ANIMAL accuracy, another reason for lower accuracy in Duchesne's experiments might be that the manual segmentations of their training set are less consistent than ours. Also, differences in MR imaging might cause some of the difference in the outcome of the experiments.

Conclusions

Segmentation of 3D medical images is an ongoing research topic. We presented a method which is fully automatic and generally applicable to any 3D medical image modality; its basic requirement is a representative training set of manually segmented images. The method, which we term a 3D AAM, successfully joins the benefits of the two widely used segmentation approaches: SBR and AAM.

With respect to the regular AAM implementations, our 3D AAM adds automatic and straightforward 3D implementation. Additional benefit is the inclusion of the SOI surroundings into the model; consequently, the model-fitting algorithm does not rely on the SOI edges only, and is therefore suitable for segmentation of poorly defined SOIs (SOIs with fuzzy boundaries) and low quality images. With respect to the SBR, the 3D AAM features two major advantages: it embeds *a priori* knowledge about the appearance of the images, and most importantly, the processing time is brought down from several hours to only a few minutes.

We evaluated the segmentation accuracy of the 3D AAM on the task of hippocampus segmentation from MR images. With respect to the accuracy of the underlying registration method, we found our model to be slightly less accurate, but we confirmed with additional tests that the main bottleneck for improving the accuracy of the model is simply a too restricted training set of images. This indicates that the 3D AAM method has the potential to replace the registration methods in many practical segmentation settings, especially due to its high speed.

Our future work includes testing the method on a larger training set of images. Also, an objective comparison of the 3D AAM model with fixed texture to the 3D AAM model with variable texture is needed. ▲

Acknowledgment. The data used in the study were kindly provided by the Department of Psychiatry, University Medical Center Utrecht, The Netherlands. We also wish to thank Marcel Quist of Philips Medical Systems, Best, The Netherlands, for providing us with the software for the intensity correction algorithm by Bostjan Likar.

References

1. T. Cootes, A. Hill, C. J. Taylor, and J. Haslam, The Use of Active Shape Models for Locating Structures in Medical Images, *Image and Vision Computing* **12**(6), 355 (1994); ISSN: 0262-8856.
2. T. Cootes, C. Taylor, D. Cooper, and J. Graham, Active shape models, their training and application, *Comput. Vision; Image Understanding* **61**, 38 (1995); ISSN: 1077-3142.
3. T. Cootes, G. Edwards, and C. Taylor, Active appearance models, *Proc. 5th European Conf. Comp. Vision; Lecture Notes Comp. Sci.* **1407**(2) 484 (1998).
4. A. F. Frangi, D. Rueckert, J. A. Schnabel, and W. J. Niessen, Automatic 3D ASM construction via atlas-based landmarking and volumetric elastic registration, *Proc. IPMI2001; Lecture Notes Comp. Sci.* **2082**, 78 (2001).
5. D. Rueckert, A. Frangi and J. Schnabel, Automatic construction of 3D statistical deformation models using non-rigid registration, *Proc. MICCAI2001; Lecture Notes Comp. Sci.* **2208**, 77 (2001).
6. S. Duchesne, J. C. Pruessner and D. L. Collins, Appearance-based Segmentation of Medial Temporal Lobe Structures, *NeuroImage* **17**, 515 (2002); ISSN: 1053-8119.
7. D. Collins and A. Evans, ANIMAL: Validation and application of nonlinear registration based segmentation, *Int. J. Pattern Recog. Artificial Intelligence* **11**, 1271 (1997); ISSN: 0218-0014.
8. J. Klemencic, J. P. W. Pluim, M. A. Viergever, H. G. Schnack, and V. Valencic, Three-dimensional appearance model for hippocampus segmentation from MRI, *Proc. SPIE* **5032**, 425 (2003); ISSN 1605-7422..
9. W. F. C. Baare, C. J. van Oel, H. E. Hulshoff Pol, H. G. Schnack, S. Durston, M. M. Sitskoorn, and R. S. Kahn, Volumes of brain structures in twins discordant for schizophrenia, *Arch. Gen. Psychiatry* **58**, 33 (2001); ISSN: 0003-990X.
10. B. Likar, M. A. Viergever and F. Pernus, Retrospective correction of MR intensity inhomogeneity by information minimization, *IEEE Trans. Med. Imaging* **20**, 1398 (2001); ISSN: 0278-0062.
11. L. G. Nyul and J. K. Udupa, On Standardizing the MR Image Intensity Scale, *Magnetic Resonance in Medicine* **42**, 1072 (1999); ISSN: 0740-3194.
12. D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, Nonrigid Registration Using Free-Form Deformations: Application to Breast MR Images, *IEEE Trans. Med. Imaging* **18**, 712 (1999); ISSN: 0278-0062.
13. M. B. Stegmann and R. Larsen, Multi-band Modeling of Appearance, *1st Int. Workshop on Generative-Model-Based Vision*, DIKU, Copenhagen, Denmark, 2002, p. 101.