Automatic Statistical Shape Model Building from A Priori Information

Jonathan Bailleul, Su Ruan, Daniel Bloyet

GREYC CNRS UMR 6072, Ensicaen, 6 bd Marechal Juin, F-14050 Caen cedex, France
E-mail: (Jonathan.Bailleul—Su.Ruan—Daniel.Bloyet)@greyc.ismra.fr

Abstract

We propose a method based on a priori knowledge provided by anatomical atlases to build almost automatically a Point Distribution Model of internal brain structures. A training set of 3D shapes is first constructed by registering an anatomical atlas over a MRI database. Then, the set shape instances are automatically landmarked in a Point Distribution Model optimization framework: according to the recent method developed by Davies et al. [1], a Minimum Description Length principle-based objective function drives the optimization process to the "best" annotation for input training set. Preliminary results of models built for several anatomical structures are encouraging.

1. Introduction

Magnetic Resonance Imaging (MRI) is a field that recently developed in Medicine and opened new horizons through allowing non-invasive in-vivo observations of the human body. This technical advance is particularly crucial in the human brain application for which it brings the opportunity to detect and possibly cure pathologies before they reach a dramatic stage. Though, the understanding of the human brain organization and functionality is still an emerging discipline that requires reliable cartography as far as common ways to study pathologies involve comparison of healthy and affected populations.

Nevertheless, the practical aspects of Brain Imaging applications raise several pitfalls. First of all, the contrast and resolution of MRI images are intrinsically limited compared to other sources. This fact is enworsed by the nature of internal brain structures of interest, consisting in mixtures of 3 pure anatomical tissues: *Cerebro-Spinal Fluid* (CSF) - appearing black on MRIs, *Grey Matter* (GM) and *White Matter* (WM). Thus, the segmentation (i.e. boundary determination) of most anatomical structures is a difficult and sometimes subjective process.

The intervention of human anatomical experts for performing segmentation is particularly delicate in the Brain MRI context. First of all, the amount of data to process is high (about 124 slices per volume). Furthermore, the inter and intra experts variability of the manual segmentation is a factor that will affect the consistency of any statistical study based upon these results. Thus, The need for an automated segmentation method becomes obvious to get stable results over large data sets.

The Brain MRI context is particularly challenging for classical image processing techniques, since most contrast-based ones proved to be rather ineffective in most cases (snakes among others). Modeling and reproducing some part of the anatomical expert "skills" becomes necessary to overcome limited boundary information. In most cases, this a priori information is inferred from a training set of shape instances, whose variation defines an allowable-shape space that can be used for determining conformity of any shape instance. Though, model training remains a tedious and often manual task that, for a long time, limitated the extent of such methods in dimension, considering that training 3D models demands far much work than for their 2D reductions.

Nevertheless, the recent advances achieved in the automatic learning of shape variation finally open new opportunities for building statistical shape models in a rather automatic manner: we took advantage of this to propose an - almost - automatic method for building a statistical 3D shape model using the very popular *Point Distribution Model* (PDM) formalism.

A priori information about Brain MRI structures shape (geometry) and position is initially provided by an expert-segmented anatomical atlas. Thus the first step of our method consists in building a training set of shape instances by projecting the atlas over a large set of patient MRIs using warp registration. All anatomical structures of interest are processed alltogether during this step, but afterwards each one will be considered as an individual object. The second step consists in landmarking the previous training set instances, i.e. reducing each object instance to a set of landmark points representing variation of studied shape, using the *Minimum Description Length* (MDL) based method of Rhodri Davies [1]. The final step consists in building a PDM from the previous landmark points, which will model the shape variation of studied object along our training set.

Deliberate efforts have been undertaken to design a scheme that should primarily apply to any kind of shape, without focusing on a specific applicative domain. Besides bringing a welcome automation, the presented method applies to a rather wide class of shapes that at least covers our structures of interest. As a consequence, we now can build a 3D PDM for any brain structure provided we can add its 3D segmentation into the anatomical atlas.

2. Choosing a 3D Point Distribution Model

2.1. Atlas-based segmentation using Fuzzy Fields shape model

As our target is to segment anatomical structures from brain MRI volumes, we first concentrated on improving the method [2] that J.H. Xue developed in our lab. It primarily consists in segmenting structures of interest (for now: ventricles, putamens, caudate nuclei, thalami, hippocampi) from skull-removed brain MRIs using a priori information to overcome low contrast at edges.

This information is provided by an anatomical atlas $Atlas_a$ and its associated reference MRI volume Ref_a (fig.1), where $Atlas_a$ results from the segmentation of structures performed by an expert onto Ref_a. Given that such a task is really time-consuming, we only have a restricted number of such atlases (Cyceron, Harvard SPL, Talairach).

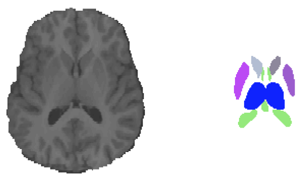

Figure 1. Corresponding volume slices of Ref_a and associated $Atlas_a$.

In that method, selected $Atlas_a$ is first registered* into a given patient MRI Mri_λ to provide approximative structures positions (seeds). Then, for each structure seed, a spatial 3D Fuzzy Field is generated in basically computing an Euclidean distance map whose values decrease significantly where we get closer to another structure. Thus these values can be interpreted as probabilities for a voxel to belong to related structure. To sum it up, spatial fields represent a statistical shape model expressing a priori information extracted from $Atlas_a$ concerning geometry, position and size of individual anatomical structures, and also spatial relationship between these.

Concurrently, Fuzzy Markov Random Field oversegmentation is performed onto Mri_λ, resulting in delimitation of homogeneous regions labeled with fuzzy values evaluating the mixture proportions of the 3 pure brain tissues (CSF, GM, WM).

*registration infers a transformation from Ref_a to Mri_λ which is then applied to $Atlas_a$.

Finally, a segmentation result is inferred slice-by-slice through conjunction of fuzzy spatial and tissue information sources. The method in [2] uses a rather stochastic approach employing a Genetic Algorithm (GA): this implies the use of a simplified fitness function to limit computation overhead, since GAs propose many unlikely configurations. In order to overcome this limitation, we rather adopted a sequence of morphologic filters, who progressively converge to the final solution in providing intermediate results, thus authorizing more flexibility and emphasizing deterministic behavior. The key steps of our method consist in determining safe cores for structures via spatial fields thresholding and then to dilate the cores in examining whether candidate voxels satisfy both tissue and spatial constraints.

2.2. Emerging need for an explicit shape model

Thanks to our collection of anatomical atlases, we managed to evaluate the relevance of our results, as far as segmentation of MRI $Ref_{b \neq a}$ using $Atlas_a$ can be compared with expert segmentation $Atlas_b$, here considered as *ground truth*. For both methods, results show good global positioning and coverage rate ($> 80\%$). Nevertheless, noticeable artifacts on low-contrasted structure edges still persist, except for ventricles whose segmentation appears satisfactory due to their high contrast.

One reason for explaining this limitation is the difficulty to establish proper tissue values thresholds in former step, though we project to improve them through establishing statistics on large MRI training sets. The key reason comes from the spatial fuzzy fields formulation which only estimates the probability of individual voxels to belong to structure: a segmented contour of unlikely distorted form cannot be discriminated provided that its voxels lie within rather high probability zones. Fortunately for our method, this flaw can only occur at local scale since we progress gradually from globally-registered seeds of correct shape, and can also count on tissue information.

Thus, in order to improve our results, we should extract more explicit shape information from $Atlas_a$ and enforce them along the segmentation process. Though other approaches have been examined, the Point Distribution Model from Cootes & Taylor [3] seems to best suit our purpose. Let's examine it in the most popular 2D case.

A PDM model requires an input training set of n_s segmented instances of studied object, each instance being annotated by n landmarks. Some of these landmarks have been placed to loci best characterizing the object's shape (e.g points of high curvature) or of particular anatomical significance: they can be seen as shape invariants; the other ones are then regularly placed along the contour between the previous "key" landmarks. Landmarks are also labeled so as to designate for given label the very same locus along training set instances. Since the PDM will show whether some selected points tend to move together (or not) over training set instances, key landmarks must be placed with attention, even if correlation is mostly captured thanks to intermediate points. If we take the hand example (cf e.g[4]), anatomical landmarks will be placed at fingertips but correlation will be more obvious for intermediate finger points as the finger moves along the training set.

Since we focus our study on shape variation, we will consider our set of shapes as polygonal approximations defined by landmark points: these polygons generally provide accurate representations provided that enough landmarks were placed. Therefore our training set of shapes now becomes a set of shape vectors $\{s_i\}_{(i \in [1,n_s])}$ built in concatenating landmark points components ($s_i : \{x_1, y_1, ...x_n, y_n\}$), thus of dimension $n_p = 2n$ in our present case.

Nevertheless, to objectively compare shapes, we first need to align the shape vectors polygons to same scale, orientation and center to enable closest landmarks correspondence in a least-squares minimizing scheme. We adopted the classical Procrustes Analysis, though Cootes et al. recommend a slightly different approach [3]. *Principal Component Analysis* (PCA) is then applied to the covariance matrix derived from aligned $\{s_i\}$, resulting in n_p sorted normalized eigenvalues with corresponding eigenvectors of dimension n_p. Then, the PDM formulation allows to model our training set instances as:

$$s_i = \overline{s} + \sum_{m=1}^{n_p} p^m b_i^m \tag{1}$$

, where $\overline{s}$ is the training set mean shape, p^m the $m-$th eigenvector and b_i^m its associated coefficient. Hence, the PDM models linear variations from the mean shape whose modes are represented by orthogonal eigenvectors. Also, noticing we just get the mean shape when b_i coefficients are set to zero, we can consider these as identifiers of s_i in the current PDM model.

Admitting the eigenvalues λ_m can be considered as variances of corresponding modes p^m, we also notice that the amplitude of shape variation expressed by p^m depends on the value of λ_m. This is of particular interest since it appears that mode variances decrease dramatically fast: a subset of t first eigenvectors can then be selected to model approximate but proper shape variation.

So, any shape instance - belonging to the PDM training set or just unseen - can be characterized by its unique set of b_i coefficients. Since they are applied to corresponding eigenvectors, we can verify that each coefficient stays within the bounds of eigenvectors-related eigenvalues - we generally allow $+/-3\lambda_m$ variation - to determine whether the variation from the mean shape induced by current instance is likely or not regarding to current PDM. Thus, the PDM model defines an *Allowable Shape Domain* (ASD) under the form of a $n_p - D$ ellipsoid, henceforth providing a criterion for discriminating likely and unlikely shapes.

Although this model can basically represent shape variations of objects of any dimension, its applications - to brain MRI among others - are far more frequent in 2D than in 3D considering that building an annotated 3D training set raises new problems we will now examine.

3. Towards Automatic 3D PDM Building

3.1. Automatic building of a 3D training set

In Brain MRI domain, proper 2D PDMs were achieved using a limited number of training shapes ($\simeq 20$, cf [5]) implying reasonable expert effort. But modeling 3D shape variations with an extra degree of freedom should require more shape instances (e.g 82 in [1] for Hippocampus), which has dramatic consequences considering that building only one 3D shape instance already requires expert segmentation on about 50 volume slices in average. Therefore the need for an automatic training set building method becomes obvious.

In our previous method, we used AIR† package [6] to register $Atlas_a$ to new patient through looking for a warp transformation associating Ref_a to Mri_λ. This is first a global process involving the whole brain information, and mostly driven by strongest intensity differences. With the white cerebral cortex surrounding our internal structures and the black ventricles tightly linked to them, a good structure overall positioning might always be obtained. The local warping aspect also contributes to the good fitting of cortex and ventricles, since translations, rotations and scaling are not generally sufficient to have similar shapes overlap. Furthermore, it also enables less contrasted structures to match.

The low contrast of our structures obviously introduces residual errors that make this registration not acceptable as a segmentation result. Nevertheless, the global nature of the process showed to preserve global shape geometry of registered atlas regarding to $Atlas_a$. Also, both global and local fitting obviously adapted $Atlas_a$ to Mri_λ, even in a coarse way. The idea to use registration of an anatomical atlas as an automatic mean for capturing variability from a set of patient MRIs might thus emerge.

Such a choice really needs to be defended. A first argument can be that the observations we performed during our first method showed surprisingly good correspondence except in special cases ($\simeq 20\%$) whose elimination still requires eye-examination. A second argument is that variations among generated shapes are obvious to the eye (fig 2) even after rigid realignment, and that a PDM should manage to capture and model it properly. A final argument is that the disputable quality of our training shapes should be compensated by their potentially high number, since a PDM models shape variability.

†version 5.05, http://bishopw.loni.ucla.edu/AIR5/index.html

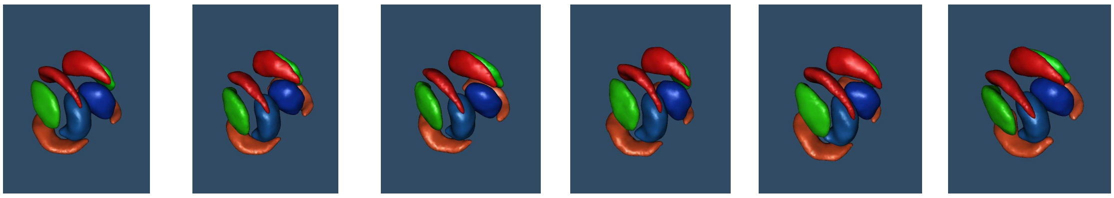

Figure 2. The original SPL atlas and 5 shape instances derived from warp registration.

Finally, we will assume we can use such shape instances as a training set for our PDM model in a Bootstraping approach, considering that the segmentation results we could get using the generated PDM might be used to refine the quality of the input set, until accordance to expert results converges to a satisfactory degree.

3.2. Automatic landmarking using the MDL principle

Shifting from 2D to 3D also complicates the landmarking process. 2D landmarking over limited training sets remains affordable for experts or trained users, helped by the fact that locating invariants over the same 2D shape is often rather intuitive. In 3D case, our first idea could be to iterate this process to each slice: this completely failed in our case. Hill et al. [7] managed to get probing results in subdividing a U-tubular shape into subshapes similar to cylinders, onto which manual landmarking was then performed at some evenly-spaced slices. Though effective, this method requires noticeable human intervention for subdivision and landmarking, and restricts the class of objects that can be modeled.

As variations of the 3D shape have an extra degree of freedom and visualizing instances is less straightforward, 3D annotation becomes completely unintuitive, thus raising the need for automatic methods. Some already exist to locate geometrically noticeable points on a single mean volume associated to its training set (e.g based on curvature, or resulting from mesh decimation strategies [8]). But in all cases, none of these methods ensures that landmarks are chosen in order to best represent shape variation over the training set, the key criterion to design a proper PDM.

Recent work of Rhodri Davies [1][9][10] formulated training set landmarking as a global optimization process. Each shape is mapped onto a corresponding sphere where a given number of landmarks is first evenly disposed. Then, their positions are blindly altered by continuous 3D reparameterization functions associated to each sphere. Their evolution is regulated by an objective function evaluating the “quality” of the PDM inferred by back-projection of displaced landmarks onto training set shapes. Thus, the whole process can converge to the individual shape landmarking that will raise the “best possible” PDM for given training set.

The aforementioned objective function relies on the Minimum Description Length principle, which can be summed up as follows. Considering PCA analysis projects shape points - actually deviations from mean shape - to the space defined by computed orthogonal eigenvectors, we can assume the training set gets encoded by a centered multivariate Gaussian model. The idea, taken from the communication field, consists in attempting to pack together both model parameters and model-encoded values as a single message in the most compact form. Balancing model complexity and accuracy regarding to original data, DL is supposed to be minimal when the current training set interpretation - i.e. landmarking - is most both generalizable and compact.

Due to the choice of the spherical mapping associated to each shape instance, the class of 3D objects (2D objects are treated in [10]) is restricted to those topologically equivalent to a sphere: that limitation covers our anatomical structures and is also imposed on other considered methods.

Results (e.g [1]) show that automatically inferred PDM are more efficient than the ones established through tedious manual landmarking, thus demonstrating the high potential of that method. One reason for this apparent paradox is that landmarks are blindly chosen to optimize a PDM while manual landmarking tends to recover points of anatomical significance introducing a part of human subjectivity. The other reason is the ability to adapt landmarking to the given training set, as other methods treat all shapes in the same manner.

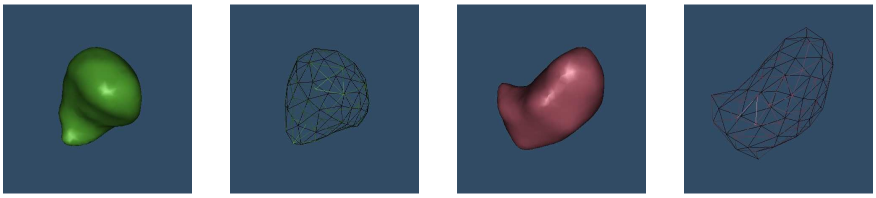

Figure 3. Two shape instances of Left Thalamus and respective landmark grids.

4. Results and extensions

This method was applied to 35 Brain MRI volumes provided by Cyceron. 27 aligned volumes (including those in fig2), featuring all of our structures of interest, were kept after the registration process. We plan to extend the number of volumes in our study, though it is already of the same order of some other 3D training sets [11].

The shape instances of each structure have then been landmarked by the MDL-based method of Rhodri Davies [1] previously described, and implemented by a program written by Allan Reinhold Kildeby [12]. The resulting shape vectors have then been rigidly aligned using the Procrustes method before a PCA was finally performed, raising variation modes partially illustrated in fig.4

Mode	Left Putamen	Left Thalamus
1	31.3109	33.2352
2	16.1430	20.0235
3	11.7009	7.2088
4	10.2925	6.9884
5	4.7667	3.2301
6	4.5503	3.1204
7	3.2502	2.5605

Figure 4. Percentages of variance of variation modes of Putamens and Thalami

Several distinct variation modes appear, which seem to indicate capture of proper shape variability. Thus, with the number t of significant modes set to 4 or 5 we can account for about 70% of total variation through a very compact model.

We also successfully managed to generate valid shape instances in varying values of b parameters within allowed bounds (cf fig 4 for a sample). This seems to show our PDM model is valid, though more exhaustive checkings presented for instance in[1] are considered.

5. Conclusion

With this new 3D PDM ready to use, we now have to set up a segmentation method: many ways have to be explored, though. The basic idea will probably be to project the mean shape onto MRI image and to deform it progressively within statistical bounds to have it fit the image information optimally. This approach is not original in itself as far as it describes the *Active Shape Model* (ASM) search procedure set up by Cootes and Taylor [3] and further specialized to Brain MRI in [5].

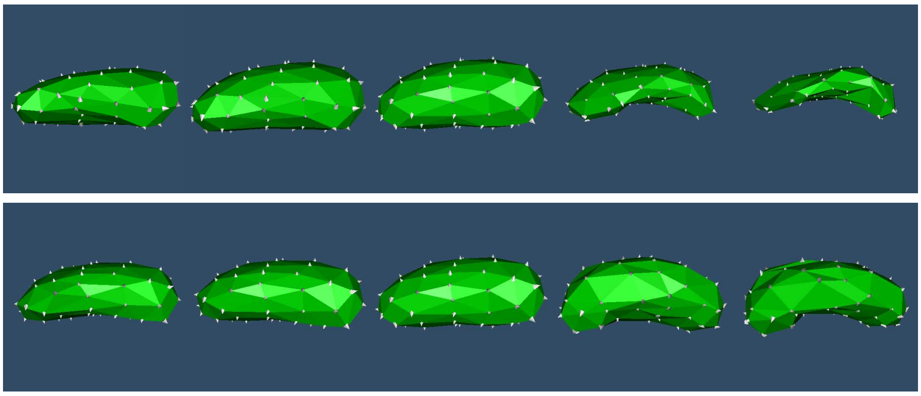

Figure 5. Generation of 4 unseen shape instances of Left Putamen in manipulating first (line 1) and second (line 2) variation modes. Middle shape is the meanshape. Camera positions might differ slightly.

Also, more recent works concern *Active Appearance Models* (AAM)[13] that combine both shape and image intensity constraints in a unified statistical model. The fact that our primary segmentation method managed to retrieve good start positions can be appreciated since this is a critical factor for the success of ASM and particularly AAM search. Furthermore, we should consider setting up a different search procedure to account for the complementary data sources we have.

6. Acknowledgments

We would like to thank Cyceron center for their support in providing patient Brain MRI data. We also thank Allan Reinhold Kildeby and Rasmus Larsen from IMM Copenhagen who provided helpful advice about their work and kindly granted us access to their 3D landmarking program [12].

7. References

[1] R. H. Davies, *Learning shape: optimal models for analysing natural variability.* PhD thesis, University of Manchester, Division of Imaging Science and Biomedical Engineering, 2002.

[2] J.-H. Xue, S. Ruan, B. Moretti, M. Revenu, and D. Bloyet, "Knowledge-based segmentation an labeling of brain structures from mri images," *Pattern Recognition Letters* **22**, pp. 395–405, 2001.

[3] D. Cooper, T. Cootes, C. Taylor, and J. Graham, "Active shape models - their training and application," *Computer Vision and Image Understanding* (61), pp. 38–59, 1995.

[4] M. B. Stegmann and D. D. Gomez, "A brief introduction to statistical shape analysis," tech. rep., Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, march 2002.

[5] N. Duta and M. Sonka, "Segmentation and interpretation of mr brain images using an improved knowledge based active shape model," *IEEE Transactions on Medical Imaging* **17**(6), pp. 1049–1062, 1998.

[6] W. RP, G. ST, H. CJ, C. SR, and M. JC, "Automated image registration: I. general methods and intrasubject, intramodality validation," *Journal of Computer Assisted Tomography* (22), pp. 139–152, 1998.

[7] A. Hill, A. Thornham, and C. Taylor, "Model-based interpretation of 3d medical images," in *British Machine Vision Conference*, Springer-Verlag, ed., pp. 339–348, 1992.

[8] A. Frangi, D. Rueckert, J. Schnabel, and W. Niessen, "Automatic 3d asm construction via atlas-based landmarking and volumetric elastic registration," in *Information Processing in Medical Imaging*, M. Insana and R. Leahy, eds., *Lecture Notes in Computer Science* **2082**, pp. 78–91, 2001.

[9] R.H.Davies, C.J.Twining, T.F.Cootes, J.C.Waterton, and C.J.Taylor, "3d statistical shape models using direct optimisation of description length," in *ECCV2002*, **3**, pp. 3–20, 2002.

[10] 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," *IEEE Transactions on Medical Imaging* **21**, May 2002.

[11] A. Kelemen, G. Szekely, and G. Gerig, "Elastic model-based segmentation of 3-d neuroradiological data sets," *IEEE Transactions on Medical Imaging* **18**, pp. 828–839, october 1999.

[12] A. R. Kildeby, "Building optimal 3d shape models," Master's thesis, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, 2002. Supervisor: Rasmus Larsen.

[13] T. Cootes and C. Taylor, "Statistical models of appearance for medical image analysis and computer vision," in *SPIE Medical Imaging*, 2001.