Proceedings of the 29th Annual International
Conference of the IEEE EMBS
Cité Internationale, Lyon, France
August 23-26, 2007.

Statistical Shape Model-based Segmentation of brain MRI Images

Jonathan Bailleul, Su Ruan and Jean-Marc Constans

***Abstract*—We propose a segmentation method that automatically delineates structures contours from 3D brain MRI images using a statistical shape model. We automatically build this 3D Point Distribution Model (PDM) in applying a Minimum Description Length (MDL) annotation to a training set of shapes, obtained by registration of a 3D anatomical atlas over a set of patients brain MRIs. Delineation of any structure from a new MRI image is first initialized by such registration. Then, delineation is achieved in iterating two consecutive steps until the 3D contour reaches idempotence. The first step consists in applying an intensity model to the latest shape position so as to formulate a closer guess: our model requires far less priors than standard model in aiming at direct interpretation rather than compliance to learned contexts. The second step consists in enforcing shape constraints onto previous guess so as to remove all bias induced by artifacts or low contrast on current MRI. For this, we infer the closest shape instance from the PDM shape space using a new estimation method which accuracy is significantly improved by a huge increase in the model resolution and by a depth-search in the parameter space. The delineation results we obtained are very encouraging and show the interest of the proposed framework.**

I. Introduction

Correlation is shown between shape deformation of brain structures and some neuropathology. In order to carry out research studies about this correlation, neurologists need to perform a delineation of brain structures of interest (i.e. the accurate determination of contours) over large 3D brain MRI data sets. As known, a manual tracing in 3D MR is not only exceedingly time consuming, but also exhausting for experts, leading to human errors. Furthermore, structures of interest show weak contrast and high noise at boundaries since they are made of mixtures of 3 anatomical tissues: Cerebro-Spinal Fluid, Grey Matter, and White Matter. That will inherently introduce variability in delineation decisions: therefore an automatic segmentation procedure is mandatory to achieve consistent studies.

Manuscript received April 2, 2007. This work was supported in part by EEC via FEDER European Help Funds.

Jonathan Bailleul is with GREYC CNRS UMR 6072, ENSICAEN, 14050 Caen CEDEX, France (e-mail: bailleul@vectraproject.com).

Su Ruan is with CReSTIC, IUT de Troyes, 9 Rue de Québec, 10026 Troyes CEDEX France (phone : (+33) 3 25 42 71 01, e-mail: su.ruan@univ-reims.fr).

Jean-Marc Constans is with Unité de Résonance Magnétique, CHU de CAEN, Côte de Nacre, 14033, Caen CEDEX France (e-mail : constans-jm@chu-caen.fr)

To overcome this uncertainty, we selected the PDM (*Point Distribution Model)* statistical shape model from Cootes, Taylor et al. [2]: it infers a shape prototype and a basis of allowable deformation modes from a training set of valid shapes annotated by corresponding landmark points.

This *n-D* model widely proved its efficiency in medical imaging (e.g. [2][7]), but took a lot of time to address 3D studies since mandatory input delineations and landmarkings represent a tedious task for experts when available. We will summarize (chap. 2) our method to automatically build a 3D PDM via Atlas registration and MDL landmarking [1][9].

The associated delineation method, named *Active Shape Model* (ASM, or *smart snake*) [2], is led by a statistical *Intensity Model* of the gray-level environment around each landmark on training set images. The PDM *Shape Model* enforces rigorous shape constraints on the deformed object instance. Both models are iterated on the object model until eventual idempotence, which is assumed to designate an optimal delineation. Nevertheless, applying this framework on our 3D MRI case raised many unexpected problems.

All methods reviewed in the literature need priors like accurate delineations datasets [6] or other sources [10] to set up the intensity model. Since Atlas registration does not provide accurate delineations, we had to develop in [9] a new direct approach we summarize in chap. 3.1.

Also, the standard framework ties the model deformation capabilities to the presence of landmarks [2]: in other works, this limitation was only addressed by using shape models of increased complexity [5].

Furthermore, the generic shape-coercion procedure is very sensitive to outliers [2] that are very difficult to avoid even in 2D [7]. A few articles address this practical problem in amending the search procedure to reduce the influence of faulty data [7]. We decided to revamp the procedure in replacing the shape parameters increase by an exhaustive search, many optimizations being required to keep this scheme rather affordable (chap 3.2).

Finally, our approach revisits the whole ASM framework: a huge resolution increase, independent from landmarking, allows for an accurate and robust delineation.

II. Automatic 3D Shape Model Construction

A. Point Distribution Model

Every shape instance of the training set must be annotated

by l landmarks, each designating the same anatomical locus along the set. Thus, the set s becomes a collection of n shape vectors: $S=\{s_i\}_{i\in[1,n]}$ with $s_i=(x_0^i,y_0^i,z_0^i,\ldots x_{l-1}^i,y_{l-1}^i,z_{l-1}^i)$.

Their alignment removes differences due to pose, which are irrelevant in shape study. Applying a PCA to the covariance matrix of s provides eigenvectors, describing linear variation modes, and associated eigenvalues. Each eigenvector can be seen as a sequence of 3D euclideans vectors applicable to each landmark point. The PDM formulation [2] (eq. 1) defines an *Allowable Shape Domain* (ASD), centered on the mean shape $\bar{s}$. Within the ASD, any valid shape instance s_x can be obtained by linear combinations of eigenvectors, coefficients b_m being bounded depending on eigenvalues λ_m:

$$s_x=\bar{s}+\sum_{m=1}^{k} p_m b_m \qquad (1)$$

B. Automatic training set building and landmarking

Building delineated training sets is a tedious task for experts, who also often keep their results strictly private. To obtain some in an independent manner, we considered a public Atlas (Harvard SPL), i.e. the result of an expert delineation of structures onto a reference MRI, that we warp-registered (AIR, Woods, cf. [1]) over a wide set of MRI images. These are rather easy to obtain in general.

This scheme generates numerous correct instances representing proper and quantified shape variability (details in [1][9]): this fulfills the conditions for a proper PDM, but breaks the classical ASM framework since registration cannot ensure an exact delineation on MRI [2].

Landmarking a training set of a single structure consists in inferring and interpreting both the mean shape geometry and the ways it varies along the set. Davies et al. [3] developed a fully automatic method making use of a Simplex optimization framework to ensure that the resulting landmarking is the best that could be found among a wide class of hypothesis. This work represented a huge stride in the field, formalizing an inherently intuitive process often addressed with arbitrament or huge limitations.

Due to patent restrictions, we rather used the method developed by Kildeby et al. [6]. It demanded advanced processing to transform the voxellic structures produced by atlas registration into smooth and regular meshes [9].

III. Segmentation

The ASM segmentation framework [2] consists in the cooperation of two models: an intensity model and a PDM shape model. First, an initial shape guess is posed into the target image. Then, the intensity model is invoked to move points individually along their normal direction in search for a better boundary. Nevertheless, noisy and low-contrasted borders are frequent and cause incorrect proposals. The purpose of the PDM shape model is to amend these proposals if they threaten shape integrity, namely if the shape of the whole model escapes from the ASD.

A. Intensity Model

Though 2D/3D ASM implementations might vary significantly, the Mahalanobis model is commonly invoked. During training, for each landmark l, gray-level information is collected on all n expert-delineated images. This provides n segments g_l of $2v+1$ voxels each, centered on landmark l and oriented in curve/surface normal direction. Their mean $\bar{g}_l$ and covariance matrix S_l are then computed [2].

During delineation, segment G_l of size $2V+1$ ($V>v$) is extracted at landmark l position on estimated shape: the position of $\bar{g}_l$ onto G_l that minimizes the Mahalanobis distance (eq. 2) gives the new estimated landmark position:

$$D_M=(G_l-\bar{g}_l)^T S_l^{-1}(G_l-\bar{g}_l) \qquad (2)$$

A major limitation for this approach is that moves can only be computed at landmark positions. But considering our particular estimated training set isn't accurate enough anyway, we absolutely need to design a new intensity model requiring no intensity priors.

Our approach consists in finding the position of g_l onto G_l for which the v inner and v outer voxels are most dissimilar. For this, we examined our MRIs and selected some measures that proved their relevance in at least 33% of cases, expecting their mutual overlapping to be more reliable. Among these measures, *intensity difference* (fig 1: measure 1, green) expects an intensity change at the boundary while *means difference* (dark blue) expects it between inside and outside, *inner voxels regularity* (pink) expects the inner structure to be a regular continuum via standard deviation, while *outer/inner regularity difference* (light blue) assumes we can oppose some outer "chaos" to the supposed inner structure homogeneity (details in [9]).

On a single Putamen case study, combined measures found exact or very close boundary position in 75% of cases [9]. This result is very encouraging, since the PDM shape enforcement should manage to deal with remaining errors.

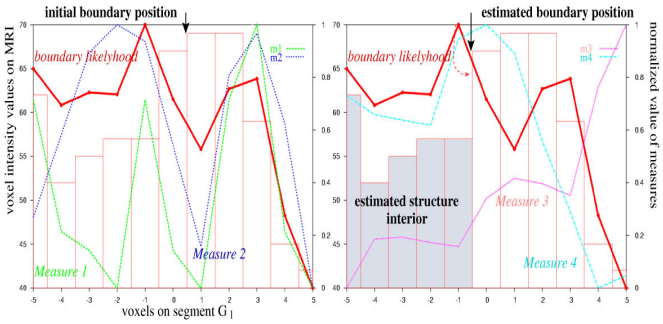

***fig 1**: the maximum value of the boundary likelihood (bold red), obtained by linear combination and filtering of measures 1-4 [9], designates the estimated boundary position for current model.*

B. Shape Enforcement

Within the ASM framework, we propose an updated scheme for model deformation and shape coercion that takes advantage of our flexible intensity model, of recent advances in computer graphic and mesh processing, and of the huge computational leap that occurred since the first ASM works.

1) Shape flexibility

The PDM is made of landmarks that, by definition, form a cloud of unorganized points: their only common feature is to lie at the object's surface. Generally, and in our particular case, these points are the vertices of a triangular mesh approximating the object's surface [10].

Though, landmarks posed by experts tend to focus on noticeable points regarding only anatomy or geometry, and the adjunction of *intermediate landmarks* [2] tends to trade sparsity with model consistency. This focus leads to excessive sparsity or uneven scattering of landmarks. Even in our ideal case where 258 landmarks can be automatically posed by a program, the scattering shows significant local variations. These fade off very slowly when more landmarks are posed, while the resulting PDM optimality decreases since compacity is a key criterion for MDL.

This induces bias in the object deformation freedom that can have dramatic consequences: e.g., how to bend a model toward a particular image without enough points likely to decently approximate such a nonlinear deformation? These problems seemed out of focus for MDL landmarking methods, which aimed at morphometric studies within the shape space rather than ASM segmentation [3][6].

We decided to use the *Qslim* [4] program to reconstruct an even and dense mesh from uneven landmark meshes, relying on progressive multiresolution 3D meshes. Most of the time, the mesh is refined so as to generalize the smallest triangles to the whole model with some interpolation in-between: the result typically shows 700+ vertices in our case. For each object, both the landmarking and the *Qslim* phase are tuned together so that the evenized object can show relevant local details regarding to the original training set.

This framework provides a detailed and flexible deformable model that avoids the density/scattering problems inherent to landmarked instances. It requires an intensity model like ours to infer boundaries from anywhere on the surface.

2) Shape coercion

Following the ASM framework, each vertex of the current estimate mesh *T"* is displaced by the intensity model towards a better boundary guess. Resulting mesh *T'* needs shape-coercion so as to preserve the learned shape, given that successive shape-consistent deformations will eventually lead to idempotence on the actual object boundary [2]. It is also assumed that this coercion consists in finding, within the ASD, the closest shape instance *T* that can affinely align to *T':* residual errors are simply declared shape-inconsistent for this iteration and cut-off [2].

The standard method [2][5] directly updates a set of shape parameters $\{b\}_i$ from $\{b\}_{i-1}$ and the residuals of alignment of T_{i-1} to T'_i. It was designed for speed (<1s) since it is invoked at each iteration. It is very efficient for estimating shape instances taken from the ASD, but becomes sub-optimal when residual errors are noticeable due to high sensitivity to outliers (cf. study in [7]). The tests we routinely performed with our 3D models being adapted to patient MRIs raised only such erratic configurations.

This motivated the development of a new shape estimation procedure that could be robust to our demanding conditions and take into account that, due to flexibility constraints, our models have an arbitrarily high number of vertices $>>l$.

Since we often managed to get better results in intuitively guessing parameters by hand, we eventually decided to formalize this procedure in performing an exhaustive search in the shape parameter space, aiming at finding the closest shape instance, no matter the time complexity. Fortunately, several factors might keep this brute-force approach affordable anyway. First of all, the shape space has the form of a gaussian closed hyperellipsoid, for which the first eigenvectors explain most of the variance. So, the shape space can be run in densely sampling the very first modes: sparsity on other modes has progressively less influence.

We know we should not find the best estimation this way, but all we need is a very close guess: we noticed that, depending on discretization conditions, very close shapes could be generated from very different coefficients values.

We start by setting $\{b\}_0$ to $\{0\}$, then we sample the first m-mode to find the values of b_m that will raise the *closest* shape instance S to the target mesh T' ($\{b\}_i$ generates S_i' with l vertices, and *Qslim* infers S_i). Running through 5 first modes is generally enough, but at least two passes are recommended since we start from $\{0\}$, a value arbitrarily far from the solution. Smaller steps and fewer samples can be computed at each pass: finally, about 100 samples should be enough to converge accurately to a close guess T of T'.

Nevertheless, if we know how to run sparsely through the shape space, we still need to quantify the quality of a shape sample *S.* The criterion will be a distance measure since we already invoked *closeness*. Since our meshes are both dense and even, we can consider a point-to-point distance between S and T', a standard in computer graphics.

Since accuracy matters, we selected the common symmetric Hausdorff distance: $H_{Sym}(S,T')=max(H(S,T'), H(T',S))$, considering $H(S,T')$ is the sum of Euclidean distances between each point s of S and its closest neighbor t in T'. Since shape is pose-independent, this distance must be measured once S and T' have been rigidly aligned: this is achieved by the standard Iterative Closest Point (ICP)

algorithm, which invokes H_{sym} distance over about 5 iterations, assuming both meshes are posed rather close.

So, computational overhead threatens, since one H_{sym} call requires up to $2n_S n_T$ iterations (n_S, n_T >700). Fortunately, such algorithms being of paramount importance in computer graphics, very optimized methods (e.g. *Mesh* [8]) were recently developed to limit the number of iterations: for instance, grouping vertices in clusters sets a short upper bound to neighbor search width.

TABLE I
Mean estimation error for both classical [2] and proposed shape estimation method

Test shapes	1	2	3	4	5	6	7
Std. method error	5.3	8.0	4.8	3.55	6.5	7.2	5.9
Time (s)	1	1	1	1	1	1	1
Proposed method error	2.5	1.5	2.1	1.2	0.8	1.3	2.2
Time (s)	512	603	540	523	607	561	533

Unit is % of shape instance bounding box diagonal.
Test shapes randomly selected among commonly used.

This estimation method raises very good results, since a close guess for all our instances can be found with robustness and sufficient accuracy (cf. fig. 2, Table 1), while the standard method [2] appears irrelevant in our configuration (cf. Table 2). Furthermore, our estimation method can work with arbitrary complex meshes, which takes advantage of our intensity model capabilities and makes ASM delineation more flexible and accurate.

As a consequence, our method can be also used to directly provide an excellent initial pose for the model in aiming at the very dense mesh (>3000 vertices) inferred by atlas registration. In the standard method [2][5], the guess would be by the mean shape, posed according to mean parameters learned from the training set. All these are arbitrarily far from the correct guess and thus do not ensure proper convergence for ASM delineation, especially in 3D.

IV. CONCLUSION

We presented a framework that significantly improves ASM delineation. Our intensity model can operate from any location at the object boundary and does not require learning from expert delineations. Thus, we can increase the resolution of the shape model for improved accuracy in boundary search. This resolution increase is also exploited by our new shape estimation method to enforce shape constraints in an accurate and robust manner for all realistic cases encountered while deforming a model towards a 3D MRI. This also solves the sensitive initial model pose.

Though, this estimation procedure has a significant time cost that remains rather affordable: each ASM iteration takes about 10 minutes on a low-end PC (Athlon 64 2.4ghz), 4-5 iterations being required in average.

Our ASM results are quite encouraging (cf. fig 3) and always better than atlas registration, though more precise quantification demands very accurate reference delineations

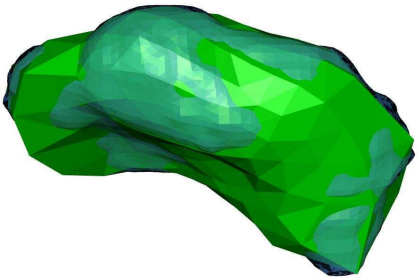

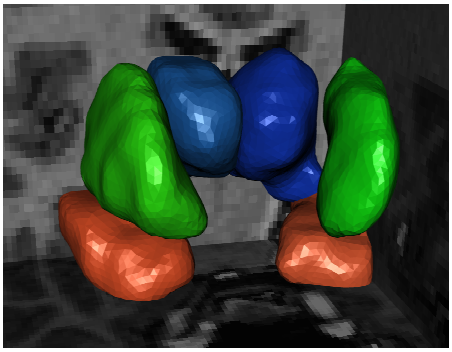

Fig 2 (up): fit of 258 points shape instance (full green) onto 3000 points initial model (transparent blue) obtained from atlas registration.
Fig 3 (down): MRI delineation example (green: putamens, blue: thalami, orange: hippocampi)

ACKNOWLEDGEMENTS

Cyceron GIN (brain MRI samples), Allan Reinhold Kildeby [6] (landmarking program), Sebastien Fourey (tesselation from voxels), Nicolas Aspert [8] (*Mesh*), Laurent Rineau (INRIA), and the whole VTK community.

REFERENCES

[1] Bailleul. J., Ruan S, and Bloyet D., "Automatic statistical shape model building from a priori information", In ICISP 2003 Proceedings, volume 2, 540--547, Agadir, Morocco, June 2003.

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

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

[4] Kircher S., Garland M., Progressive Multiresolution Meshes for Deforming Surfaces". ACM/Eurographics Symposium on Computer Animation, 191--200, 2005.

[5] Kelemen A., Szekely G., and Gerig G., "Elastic model-based segmentation of 3-d neuroradiological data sets", IEEE Transactions on Medical Imaging, 18(10):828-839, October 1999.

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

[7] Lekadir K., Merrifield R., Yang, G. Z., "Outlier Detection and Handling for Robust 3D Active Shape Models Search", IEEE Transactions on Medical Imaging, 26(2), February 2007.

[8] Aspert N., Santa-Cruz D. and Ebrahimi T., "MESH: Measuring Error between Surfaces using the Hausdorff distance", Proceedings of the IEEE ICME, I:705--708, 2002.

[9] Bailleul J., Ruan S., Bloyet D., "Segmentation of Anatomical Structures from 3D Brain MRI using automatically-built Statistical Shape Models", ICIP'04 Proceedings, 2741--44, Singapore, October 2004.

[10] Pitiot, A., "Automated Segmentation of Cerebral Structures Incorporating Explicit Knowledge", PhD thesis, Ecole des Mines de Paris, INRIA, November 2003