ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage xx (2004) xxx–xxx

Expert knowledge-guided segmentation system for brain MRI

Alain Pitiot,[a,b,c,*] Hervé Delingette,[a] Paul Thompson,[b] and Nicholas Ayache[a]

[a]*EPIDAURE Laboratory, INRIA, Sophia Antipolis, France*
[b]*Laboratory of Neuro Imaging, UCLA School of Medicine, Los Angeles, USA*
[c]*Mirada Solutions, Ltd., Oxford, UK*

Received 1 July 2004; accepted 1 July 2004

We describe an automated 3-D segmentation system for in vivo brain magnetic resonance images (MRI). Our segmentation method combines a variety of filtering, segmentation, and registration techniques and makes maximum use of the available a priori biomedical expertise, both in an implicit and an explicit form.

We approach the issue of boundary finding as a process of fitting a group of deformable templates (simplex mesh surfaces) to the contours of the target structures. These templates evolve in parallel, supervised by a series of rules derived from analyzing the template's dynamics and from medical experience. The templates are also constrained by knowledge on the expected textural and shape properties of the target structures.

We apply our system to segment four brain structures (corpus callosum, ventricles, hippocampus, and caudate nuclei) and discuss its robustness to imaging characteristics and acquisition noise.

Keywords: Magnetic resonance images; Simplex mesh surfaces; 3-D segmentation system

Introduction

The rapid development of imaging technologies (Ayache, 2003) now routinely allows living organs and organisms to be explored noninvasively. One of the least accessible and most complex organs, the human brain, is a primary beneficiary of these new medical imaging techniques. Its complexity is expressed at a variety of scales. At the microscopic level, neurons, glial cells, and fibers form the support tissue for cerebral communication. At a more macroscopic level, the brain can be partitioned into several regions (e.g., brainstem, cerebellum, diencephalon, and cerebrum) which are associated with high-level mechanisms such as sensation, motor control, or affect and cognition. Within these regions, we distinguish substructures (e.g., the amygdala, hippocampus, basal ganglia, etc.) in view of whose functional importance the development of precise segmentation and labeling methods has become a major objective of neuroinformatics. The need, shared across many levels of description, for such correlation between brain structure and function is exemplified by the broad range of studies that have analyzed cortical structures: in particular, diseases such as schizophrenia (Narr et al., 2000), through development (Thompson et al., 2003), etc.

Although qualitative image analysis is often sufficient for diagnosis of disease, quantitative analysis, for which segmentation is a pivotal first step, is necessary for many applications: longitudinal monitoring of disease progression or remission (Rey et al., 2002), preoperative evaluation and surgical planning (Holly and Foley, 2003), radiotherapy treatment planning (O'Sullivan and Shah, 2003), or statistical analysis of anatomical variability or deficits (Collins et al., 2003; Thompson et al., 2000).Yet, effective segmentation is especially challenging, as a structure can present a wide variety of shapes and appearances.

Automated segmentation of brain structures

In spite of the high variability of brain structures, the delineation process calls for high precision as the quality of the analysis generally depends on how accurately the various structures can be identified. For instance, as argued in (Thompson et al., 1997), given the corpus callosum's key role as the primary cortical projection system, regional analysis of its structure is important in assessing several neurological disorders (Alzheimer disease, multi-infarct dementia, dysplasias, etc.). Nonetheless, subtle variations in shape, relative to a mean callosal delineation, are observed between and within patient and control groups, and this makes it difficult to detect and classify abnormal structural patterns. As a result, intense debate still rages on whether different callosal regions undergo selective changes in each of these disease processes and whether these are systematic differences in neuropsychiatric disorders such as autism or schizophrenia. These controversies may be alleviated by precise and reliable segmentations, applied to large image databases.

Segmentation has traditionally been tackled by human operators, but the many drawbacks of manual delineation (lack of reproducibility, strong a priori biases, unavailability of sufficient resources to handle ever-growing databases of images) advocate the use of automated methods. However, to reach the desired

* Corresponding author. Mirada Solutions, Ltd., Level 1, 23–28 Hythe Bridge Street, Oxford OX1 2ET, UK. Fax: +44 1865 265501.
E-mail address: apitiot@loni.ucla.edu (A. Pitiot).

Available online on ScienceDirect (www.sciencedirect.com.)

1053-8119/$ - see front matter © 2004 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2004.07.040

accuracy, many difficulties must be overcome: input images are noisy, poorly contrasted, and full of "decoys" (many structures are similar in shape and/or in intensity), and the target structures are variable in shape and intensity, etc.

A plethora of automated segmentation methods have been proposed in the literature to extract anatomical structures, using an array of feature descriptors and shape models. The choice of an adequate segmentation paradigm is central as it conditions the ability of the segmentation system to extract anatomically meaningful delineations. We favored deformable templates as a basis for our segmentation system, for the following reasons. First, they can adequately handle the various discontinuities and irregularities induced by sampling artifacts or noise along the boundaries of the target structures. Next, they can compactly describe a wide variety of shapes while minimizing the overall number of parameters or masking these behind a small and easily manageable set of physical principles. They also often provide a local, if not global, analytical representation of the segmented structure, which facilitates its subsequent analysis. Finally, a priori knowledge on the shape, location, or appearance of the target structure can guide the deformation process. Deformable templates are then the framework of choice for combining bottom-up constraints [computed from the input magnetic resonance imaging (MRI)] with a priori top-down medical knowledge.

Model-based segmentation using explicit knowledge

In many deformable template techniques, statistical analysis helps to introduce a priori knowledge on the shape or appearance of the target structures.

Most of these approaches fall in the *implicit knowledge* category: from a learning set of a priori segmented instances of an anatomical structure, they must automatically discover the relationships and functional dependencies of the various model parameters.

However, *explicit* information about the target structures is often available, based on anatomical expertise. For instance, the relative positions and topology of most of the key subcortical gray matter structures are fairly consistent across individuals, anatomical structures should not intersect, etc. From these observations, rules can be derived to better drive the segmentation process. Broadly speaking, explicit knowledge approaches may be regarded as a special case of implicit knowledge algorithms where additional biomedical expertise provides short cuts in searching for the target structure.

We submit that the use of this a priori medical expertise in general, and explicit knowledge in particular, is the key to a robust and accurate segmentation system.

Reviews of various implicit knowledge deformable template-oriented techniques can be found in (McInerney and Terzopoulos, 1996) and (Montagnat et al., 2001). Explicit knowledge approaches are more heterogenous as they usually combine shape and intensity descriptions in the same framework. Also, explicit information is often complemented or generalized by implicit information (for instance, a purely explicit position rule can be made more robust as a fuzzy condition, which introduces nonexplicit elements: the α parameter of the cut-off, the amount of diffusion, etc.).

These close interactions between implicit and explicit models are exemplified in the hierarchical active shape models (ASMs) of Bernard et al. (2001). Pioneered by Cootes et al. (1994), ASMs infer new shapes by linearly combining the eigenvectors of the covariance matrix which captures the variations from the mean shape. These eigenvectors encode the modes of variation of the shape. The shape parameter space then serves as a means to enforce limits and constraints on the set of admissible shapes. Although ASMs can handle disconnected shapes, it is easier to partition a complex shape (e.g., the vertebral column) into simpler and more manageable elements (the vertebrae). Noting this, Bernard et al. devised a two-level hierarchical scheme to model the shape and topology of the resulting composite representation. Each individual structure was controlled by its own ASM, subject to an overall global ASM encoding the relative positions and orientations of the set of components.

In another type of explicit approach, Amit and Kong (1996) used a graph of landmarks, automatically chosen from the input images, as a topological model to guide the registration process of X-ray images of the hand.

In view of its ability to represent and merge uncertain or imprecise statements, fuzzy theory also proved a popular choice. Among others, Chang et al. (2000) developed a fuzzy-controlled rule-based system to segment MR images of diseased human brains into physiologically and pathologically meaningful regions by incorporating expert knowledge on brain structures and lesions. Barra and Boire (2001) used information fusion to combine medical expertise with fuzzy maps of morphological, topological, and tissue composition data to segment anatomical structure in brain MRIs. Studholme et al. (1996) merged region labeling information with a classical iconic image registration algorithm via information fusion to align MR and PET images of the pelvis.

Anatomical atlases are also particularly well suited to model a priori knowledge. In Csernansky et al. (1998) for instance, fluid warping of a digital brain template helped study the relationship between schizophrenia and local changes in hippocampal morphology. The ANIMAL algorithm (Collins et al., 1995) deforms an MRI scan to match a previously labeled atlas MRI, and the nonlinear transformation is used to segment it by transferring the atlas labels on the individual scan.

When anatomic knowledge can be captured by a series of simple positional, geometric, or intensity rules, expert systems provide a convenient framework to assist in segmentation tasks. Ardizzone et al. (2001), for instance, developed a descriptive language to express the geometric features and spatial relationships among areas of images. Matesin et al. (2001) also used a rule-based system to organize and classify features (such as brightness, area, neighborhood, etc.) for regions that had been automatically extracted via region growing, and they segmented scalp, gray and white matter, CSF, and strokes. In Brown et al. (1998), lung boundaries were segmented in chest X-ray images by matching an anatomical model to the image edges using parametric features guided by a series of rules. Li et al. (1995) described a knowledge-based image interpretation system to segment and label a series of 2-D brain X-ray CT scans. Their model contained both analogical and propositional knowledge on the brain structures, which helped interpret the image primitive information produced by different low-level vision techniques. Finally, Poupon et al. (1998) used 3-D moment invariants to embed shape distributions in deformable templates. They devised a framework that could deal with several simultaneously deforming templates, with a fairly low updating cost, to segment deep gray nuclei in 3-D MRI.

Composite segmentation system for medical images

We propose an automated segmentation system for in vivo brain weighted MRI (see Fig. 1). We focused on devising a segmentation

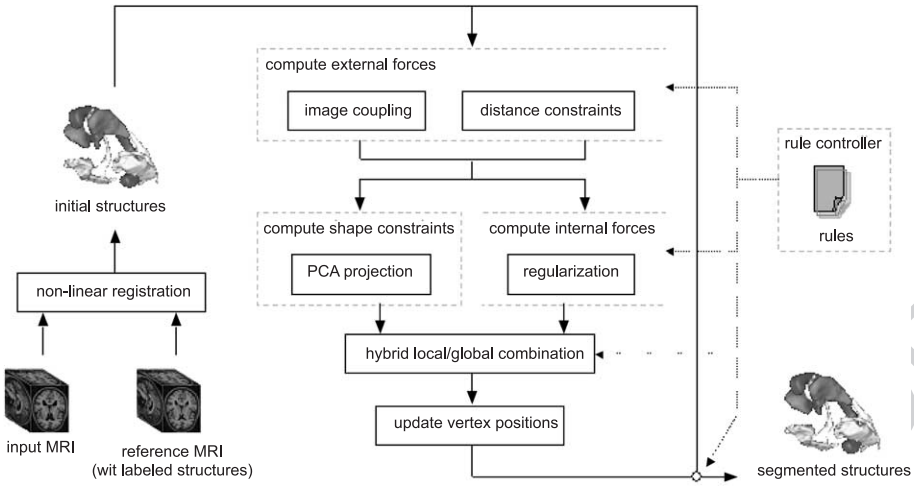

Fig. 1. Overview of our proposed segmentation system.

method that makes maximum use of available a priori anatomic expertise either in the form of implicit knowledge (the shapes of the structures, their appearances, ...) or of explicit information (the relative distance between structures, nonintersection rules, ...). These rules are implemented as constraints on the deformable templates, thereby incorporating several techniques from the above taxonomy. To illustrate the promise of our system, we selected a representative set of structures to segment: the lateral ventricles, the corpus callosum, the caudate nucleus, and the hippocampus.

We approach the issue of boundary finding as a process of fitting a series of deformable templates to the contours of the target structures. The templates are initialized by nonlinear registration of a hybrid MRI/structure atlas (built a priori) to the input MRI. Each initialized template is then iteratively modified to minimize a hybrid local/global energy which incorporates (1) an internal regularization energy, (2) an external term coupling the models to the underlying image features, and (3) a global shape-constrained term. The templates evolve in parallel within a rule-controlled framework whose purpose is to maximize the achieved match over each structure while respecting the distance, position, etc. constraints derived from neuroanatomical knowledge. For each structure, we also devise, from a learning set of already delineated instances in MRIs, a specific texture filter (here, we consider texture to be a function of the spatial variation, or distribution, of voxel intensities in a given window). This builds in a texture constraint to bias the evolution of the deformable templates towards the most texture probable boundaries.

Methods

We detail in this section the components of our segmentation system and how they interact with each other under the supervision of segmentation rules.

Deformation model

We chose simplex meshes (Delingette, 1999) to model the templates. They are discrete model representations (sets of vertices and edges) with prescribed vertex connectivity. Similar to triangular meshes (of which they are the duals), simplex meshes can represent surfaces of all topologies. To encode the surfaces of structures, we use closed 2-simplex meshes: each vertex is then connected to exactly three neighbors. This inherent topological simplicity makes it easier to impose constraints (internal and external) to guide the segmentation process. Finally, "zones" (subsets of vertices with their associated edges) can be defined on the simplex meshes to specify additional constraints (see Fig. 2).

Let $\Pi_j = \{P_j^i \in \mathbb{R}^3\}_{i=1}^{N_j}$ be such the mesh (a set of N_j points with constant connectivity matrix as we do not allow topological changes) associated with structure j (e.g., $j = 0$ for corpus callosum, $j = 1$ for caudate nucleus, etc.). We define the input MR image I by its intensity at each point. The algorithm's goal is then to find in I a pictorial object whose overall boundary fits that of Π_j. To guide the deformation and drive the template towards the required object shape, we introduce a compound energy functional E_I whose minimum we aim to determine. Classically, E_I is made up of three terms:

- an internal (or regularization) energy E_{internal} which characterizes the possible deformations of the template,
- an image coupling energy E_{image} which couples the template to the image, and
- a constraint energy $E_{\text{constraint}}$ which incorporates the various constraints (shape, texture, etc.).

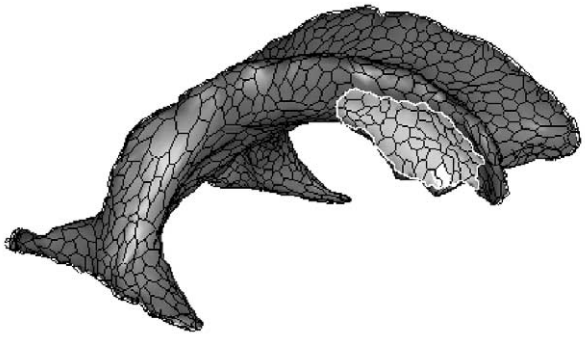

Fig. 2. Gouraud rending (gray) of the simplex mesh (black lines) associated with a model of the lateral ventricles, with a defined "zone" (white outline; see text).

E_I is then written:

$$E_I = \alpha.E_{\text{internal}} + \beta.E_{\text{image}} + \gamma.E_{\text{constraint}} \quad (1)$$

with $\alpha, \beta, \gamma \in \mathbb{R}$.

Within a Newtonian framework, we get the following iterative point updating procedure:

$$\Pi_j^{t+1} = \Pi_j^t + (1-\delta)\left(\Pi_j^t - \Pi_j^{t-1}\right) + \alpha.f_{\text{internal}}\left(\Pi_j^t\right) + \beta.f_{\text{image}}\left(\Pi_j^t\right) + \gamma.f_{\text{constraint}}\left(\Pi_j^t, \{\Pi_k^t\}_{k\neq j}\right) \quad (2)$$

where t denotes the iteration step, $\{\Pi_k^t\}_{k \neq j}$ is the set of all structure meshes with the exception of Π_j, $\delta \in R$ is a damping coefficient and the f_i's are body forces applied to displace the mesh vertices at each iteration.

A number of image-based forces are available in the literature (Montagnat et al., 2001). They may be based on the gradient of the input image, on a smoothed version of its associated edge-image, on intensity profiles, etc. Here, we use a force that depends on the distance to the closest strong gradient in the underlying image, as this exhibits a good trade-off between precision and robustness (Delingette, 1999): f_{image} is then proportional to the distance to the strongest gradient along the direction of the associated normal to the simplex mesh, within a given exploration range, which depends on the expected distance between the point in the mesh and its final position in the target structure.

We implement an internal regularization by averaging the curvature of simplex vertices over a spherical neighborhood (which effectively modifies the position of these vertices).

Initialization

Once we have reduced the segmentation problem to an energy minimization task, we face a multimodal, nonlinear, and possibly discontinuous function of many variables. As the solution space is large and nonconvex, most minimization techniques would only lead to weak suboptimal solutions (where the deformation model adapts to noise or decoys or maybe only follows parts of the desired boundaries) if the search space were not drastically reduced by assuming that a good approximation to the solution was available. This may be either in the form of a set of pose parameters (position, orientation, scale) or shape descriptors.

Various approaches have been presented in the literature to overcome this robustness issue. In Blake and Zisserman (1987), for instance, a coarse to fine strategy, the Graduated Non-Convexity Algorithm, is implemented, where a scalar parameter controls the amount of "local" convexity in the model. Alternatively, the templates may be initialized at several locations and evolved in sequence: the deformed template with the best final match is then selected. In Pitiot et al. (2002b), a hybrid evolutionary algorithm controls a family of deformable templates that are evolved simultaneously and explore the search space robustly. Here, we use nonlinear registration to initialize the templates relatively close to their expected positions.

An MRI brain data set was selected for its "standard" appearance (the reference MRI), and in it, we carefully segmented the target structures (see Fig. 3a) following anatomical delineation protocols (Pitiot (2003)—Appendix B). Given an input MRI to be processed, we register the reference MRI to it first with a robust affine block-matching registration method (the "baladin" algorithm (Ourselin et al., 2001)) and second with a nonlinear registration algorithm with an elastic prior (the PASHA algorithm (Cachier et al., 2003)). The obtained transform is then applied to the meshes segmented in the reference MRI. Those transformed meshes serve as initial guesses for the segmentation of the target structures (Fig. 3b). Note that the PASHA regularization parameters were set so as to yield a particularly smooth transformation and prevent local sign changes of the Jacobian as these could cause the transformed meshes to self-intersect.

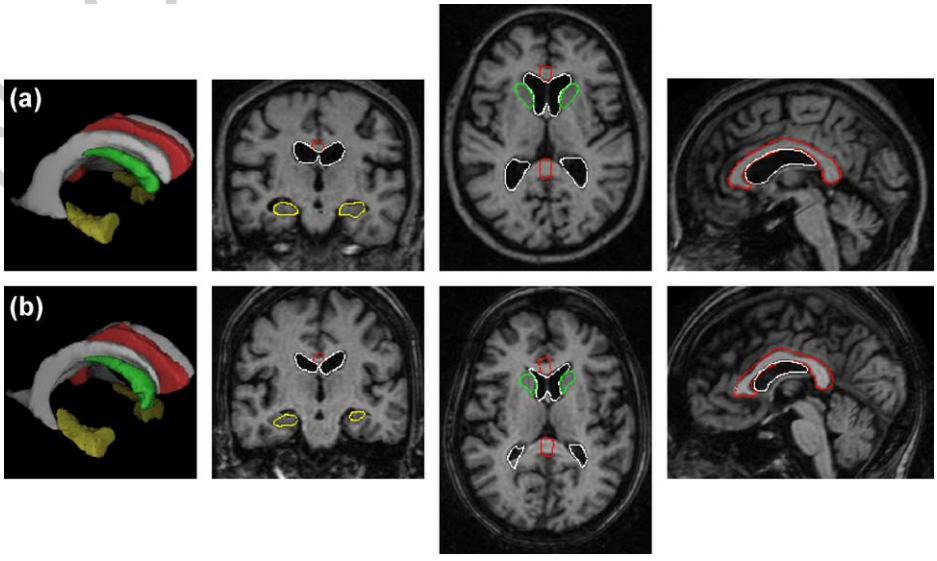

Fig. 3. (a) reference MRI with manually delineated structures superimposed (corpus callosum in red, ventricles in white, caudate in green and hippocampi in yellow); (b) reference in MRI registered to an input MRI registered to an input MRI and initialized structures. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

We found that the nonlinear registration of an average intensity atlas (Collins et al., 2003) (the average of several MRIs linearly or nonlinearly registered to themselves) used as reference MRI yielded inferior results since some of the strong features used by the PASHA algorithm (edges, ridges, ...) were not as clearly defined in the average atlas than in the "standard-looking" MRI.

Also, even though the affine registration of the reference MRI gave good initializations, better results were achieved with a nonlinear algorithm, especially when the MRI to be segmented was substantially different from the reference MRI: in this case, a global affine transformation was less effective in aligning the internal structures.

Knowledge-based constraints

The evolution of our deformable templates is guided by several constraints (energy terms and rules) towards shapes that are more probable with respect to the a priori anatomical knowledge we gathered on the target structures.

Statistical shape constraints

Even though a given structure can present a wide variety of forms, the notion of biological shape seems reasonably well explained by a statistical description over a large population of instances. Consequently, statistical approaches have attracted considerable attention (Cootes et al., 1994; Turk and Pentland, 1991; Staib and Duncan, 1992). A deformable template is then constrained not only by the number of degrees of freedom imposed by its geometric representation, but also in that it must be a valid instance of the shape model. Most of these approaches however require that correspondences between shapes be available a priori. We consequently reparameterize the meshes to form the shape learning set following Fleuté's methodology (Fleuté et al., 1999) which minimizes the distance between one of the input shapes and a second one registered with it (this assumes smooth transition paths in between them). Namely, the simplex mesh associated to the most average looking instance of each target structure is deformed (following the core deformation process described above, without external constraints) onto the other ones and the final deformed meshes then serve as reparameterizations.

Given a set $S = \{S_1, \ldots, S_N\}$ of N reparameterized instances of a target structure (the a priori learning set), we first align the structure's instances into a common coordinate frame with an iterative closest point algorithm. The eigenvectors of the covariance matrix of the positions of the structure's vertices then describe the modes of variation, and the vectors corresponding to the largest eigenvalues describe the most significant ones.

A statistical shape model is then available for each target structure. The deformable templates must then be constrained accordingly. In Cootes et al. (1998), the pose and shape parameters of the templates are adjusted by projecting the local deformation induced by the external energy onto the shape space. Let $d\Pi^t_{\text{image}} = f_{\text{image}}(\Pi^t_j) - \Pi^t_j$ be the deformation induced by the image coupling forces. Let $\bar{S}$ be the mean shape computed for the target structure, and $Q = \{q_1, \ldots, q_m\}$ its m first eigenmodes. The shape-constrained deformation is written:

$$d\Pi^t_{\text{shape}} = \sum_{i=1}^{m} \langle \Pi^t_j + d\Pi^t_{\text{image}} - \bar{S}, q_i \rangle . q_i$$

This however limits the range of possible shapes to be the projections onto the shape space. Alternatively, hybrid deformation models can be crafted where the shape constraints bias the deformation process, but less restrictively. We chose to adapt Montagnat's hybrid local/global scheme (Montagnat and Delingette, 1998). Deformations are then regularized by combining global (shape-constrained) and local (external) forces.

The point updating rule becomes:

$$\Pi^{t+1}_j = \Pi^t_j + (1-\delta)\left(\Pi^t_j - \Pi^{t-1}_j\right) + \lambda\left\{\alpha . f_{\text{internal}}\left(\Pi^t_j\right) + \beta . f_{\text{image}}\left(\Pi^t_j\right) + \gamma . f_{\text{constraint}}\left(\Pi^t_j, \left\{\Pi^t_k\right\}_k\right)\right\} + (1-\lambda)\left\{d\Pi^t_{\text{shape}}\right\} \quad (3)$$

where λ is the "locality" parameter, which controls the contribution of the global shape-model constraint.

Distance constraints

The positions (and shapes) of nearby anatomical structures are not independent of each other. For instance, the caudate nuclei are juxtaposed to the lateral ventricles, so any change in the shape or position of one will affect those of the other. Information about the respective positions of structures can then help the segmentation process.

In Barra and Boire (2001), fuzzy logic was used to express distance and positional relationships between structures. In Tsai et al. (2003), a series of parametric models, built via principal component analysis of multiple signed distance functions, enabled the concurrent segmentation of anatomical structures, via minimization of a mutual information criterion. Interobject distance constraints were also used in Yang et al. (2003), where a maximum a posteriori estimator for anatomical shapes helped constrain the evolution of level set functions. We too chose distance maps here as they can model distance constraints with good precision and robustness (to guarantee nonintersection, for instance). Given a deformable template Π^t_0, we wish to impose on it a distance constraint with respect to template Π^t_1. We first compute the distance map D^t_1 associated with a discrete sampling of Π^t_1. We use a classical Chamfer map (Borgefors, 1984) algorithm to compute a signed distance map, positive outside the discrete sampling of Π^t_1 and negative inside. At each vertex $\Pi^t_{0,i}$ of Π^t_0, we then compute a "distance force" f_1 magnitude depends on the value of the distance map at the considered vertex.

Two types of constraints can be, and were, applied:

- We can cause the force to attract the vertex, along the direction of the gradient of the distance map, up to an exact distance d_{target} of the target mesh:

$$f_{\text{distance}}\left(P^t_{0,i}\right) = -\frac{\nabla D^t_1\left(P^t_{0,i}\right)}{||\nabla D^t_1\left(P^t_{0,i}\right)||} . \left(D^t_1\left(P^t_{0,i}\right) - d_{\text{target}}\right) \quad (4)$$

- Alternatively, we may want to only enforce that this vertex should remain at distance inferior or superior to d_{target} (to prevent intersections between structures for instance). We get:

$$\text{if } D^t_1\left(P^t_{0,i}\right) < d_{\text{target}} \text{ then } f_{\text{distance}}\left(P^t_{0,i}\right) = + \frac{\nabla D^t_1\left(P^t_{0,i}\right)}{||\nabla D^t_1\left(P^t_{0,i}\right)||} . \left(D^t_1\left(P^t_{0,i}\right) - d_{\text{target}}\right)$$

else

$$f_{\text{distance}}\left(P^t_{0,i}\right) = 0$$

to ensure a minimum distance ($d_{\text{target}} = 0$ for nonpenetration).

Note that this constraint does not guarantee non-interpenetration, it only favors it, which works adequately in our case. Should another application require absolute non-interpenetration, the norm of the gradient in the above formulation could be multiplied by $\|D^t_1(P^t_{0,i}) + \varepsilon\|$, which would yield an arbitrarily large force at contact.

Also these forces can also be applied to a subset of the mesh vertices (so-called zones; Fig. 2) to enforce more local constraints.

Texture constraints

Cerebral structures are not all equally well-defined in brain MRI. The corpus callosum for instance is remarkably smooth and contrasted in the midsagittal section of a typical T1-weighted MRI. Its mean intensity is clearly higher than that of most of the immediately surrounding tissues. Similarly, the intensity distribution of the lateral ventricles clearly distinguishes them from nearby structures (the caudate nucleus for instance). However, the borders of the hippocampus are significantly harder to outline in some areas. Furthermore, these structures are often surrounded by decoy elements with similar intensity distributions.

Finally, noise, partial volume effects and bias fields also impair the quality of the input images. Yet, the segmentation process relies on the minimization of an objective function to drive the templates towards the strongest edges of the input image, which should then correspond to those of the target structure. Unfortunately, these various difficulties jointly contribute to a poor edge map, which might impair the deformation. Interaction with neuroscientists prompted us to consider texture as a discriminating element for the target structures.

We therefore developed a series of texture filtering approaches to produce classification maps from the input MR data.

From a large pool of texture descriptors (Haralick descriptors computed from co-occurrence matrices, fractal measures, dyadic Gabor filters, etc.), a specialized feature selection algorithm first discards the least pertinent descriptors, for each target structure. This selection step is performed a priori, once and for all. The selected descriptors can then be classified. Three types of classifiers were investigated: linear (linear discriminant analysis), linear in a nonlinear projective space (support vector machines), and adaptive nonlinear (neural networks), with an increase both in performance and in the computing resources required (see Pitiot, 2003; Pitiot et al., 2002a for details). A priori information on the classification task is introduced in the form of a learning set of a priori segmented target structures.

For efficiency, a region of interest (ROI) is also identified a priori around each target structure in the "standard-looking" MRI that serves to initialize the deformable templates. Given a new MRI to segment, the texture filters are applied only inside the ROIs convected by the nonlinear transformation obtained with the PASHA algorithm by registering the standard reference MRI to the one to be segmented: this decreases the learning and processing times, and enhances the performance of the classifier as fewer decoy structures (similar-looking off-target tissues) must be discarded. As it performed best in practice, we selected the nonlinear classifier to extract the target voxels.

Fig. 4 displays a few classification results for 3 structures out of the 16 test corpora callosa, 20 caudate nuclei, and 20 hippocampi available. The use of a highly specialized neural network helped design a better classifier, owing to the ability of neural approaches to adapt the structure of the decision boundary in the search space to the classification problem as they search for the best fitting parameters (most especially due to a dynamic learning set). Note however that a set of target voxels adequate for our application can be obtained with straightforward linear discriminant analysis. The technicalities behind the nonlinear

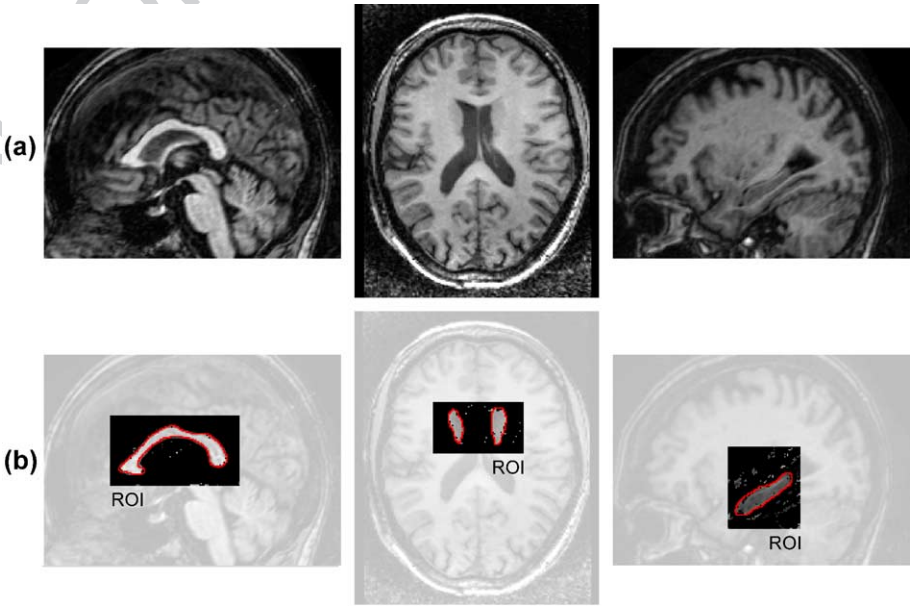

Fig. 4. Neutral classification of corpus callosum, caudate nucleus and hippocampus: (a) input T1-weighted 1 mm^3 MRI; (b) neutral extracted structures with true outlines superimposed.

approach only contribute to the final few percentage points in performance.

A distance map scheme similar to that established for distance constraints then serves to build texture constraints from the texture classification maps produced by these classifiers. Namely, we compute the distance map D_i^t of the thresholded classification map associated with each target structure and derive a "texture force" f_{texture} as follows:

$$\forall \Pi_j^t,\ \forall P_{j,i}^t,\ f_{\text{texture}}\left(P_{j,i}^t\right) = -\frac{\nabla D_j^t\left(P_{j,i}^t\right)}{||\nabla D_j^t\left(P_{j,i}^t\right)||}.D_j^t\left(P_{0,i}^t\right) \quad (5)$$

Since the texture maps are computed only inside regions of interest, the texture forces are also only available there.

Rule-controlled framework

In view of the complexity of the segmentation task, choosing a value for the various scalar parameters that control the contributions of the constraints and regularization energies is not trivial. Instead of setting a priori suboptimal values, these parameters could evolve dynamically along with the deformation process. Additionally, rather than segmenting the structures independently and running the risk of them intersecting one another, better segmentation results could be obtained by evolving the templates in parallel while controlling their interrelationships.

We therefore built a catalog of rules to control the dynamic properties of our deformable templates. For each target structure or pair of structures, a set of rules was developed that took into account recommendations from clinicians as well as low-level image observations.

Lateral ventricles. As the ventricles are fairly highly contrasted relative to the immediately surrounding tissues in T1-weighted MRIs, the nonrigid transformation obtained via registration of the reference MRI to the input MRI usually gives an excellent estimate of the true boundaries. The texture filter also delivers excellent maps and we set $\gamma_{\text{texture}} = 0.6$. With that in mind, and in view of the large variability of the ventricles, no shape constraint was used for their segmentation (Table 1 confirms that adding a shape constraint actually decreases the segmentation performance). For the same reason, only a small internal regularization energy was used. $\delta = 0.1$, $\alpha = 0.1$, $\lambda = 1.0$, $\gamma_{\text{distance}} = 0.0$, $\gamma_{\text{texture}} = 0.6$, $\beta = 0.3$.

Caudate nucleus. With the exception of the caudate tail, which the delineation protocol discards (see Pitiot (2003)—Appendix B for details), the caudate nuclei from our training set did not exhibit much variability. We consequently used a moderately high shape weight: $\lambda = 0.3$. To prevent intersections with the lateral ventricles, a distance constraint was added. We define on each caudate simplex mesh (left and right) a zone corresponding to the contact area with the adjacent lateral ventricle. A distance constraint with $d_{\text{target}} = 1$ mm ensures a good juxtaposition and prevents interpenetrations. $\delta = 0.1$, $\alpha = 0.1$, $\lambda = 0.3$, $\gamma_{\text{distance}} = 0.3$, $\gamma_{\text{texture}} = 0.3$, $\beta = 0.3$.

Corpus callosum. A fairly variable structure (at least based on the analysis of the 20 callosal instances in our training set), we did not use any shape constraint for the corpus callosum (here also, Table 1 supports this choice). A distance constraint with $d_{\text{target}} = 2$ mm ensures the nonintersection with the lateral ventricles (a 0-mm distance constraint would not prevent intersection since, as mentioned above, our distance constraints act as biases for the deformation process rather than as actual absolute constraints). $\delta = 0.1$, $\alpha = 0.1$, $\lambda = 1.0$, $\gamma_{\text{distance}} = 0.2$, $\gamma_{\text{texture}} = 0.6$, $\beta = 0.1$.

Hippocampus. The hippocampus shows poor contrast relative to its neighboring structures, so the use of a shape constraint proved necessary ($\lambda = 0.3$) to interpolate the missing information. Since the performance of the texture classifier was not particularly high, we gave the texture constraint a moderate weight. $\delta = 0.1$, $\alpha = 0.1$, $\lambda = 0.3$, $\gamma_{\text{distance}} = 0.0$, $\gamma_{\text{texture}} = 0.6$, $\beta = 0.6$.

Parameter dynamics

- A pyramidal decomposition of the gradient image (series of increasingly downsampled gradient images) was used to compute the external forces. This guaranteed deformation at early stages and later ensured a precise delineation (dynamic coarse-to-fine approach): the standard deviation of the 3-D Gaussian used to compute the gradient of I was initialized at 3.0 mm and decreased by 0.2 every 10 iterations.
- The locality parameter λ was slowly increased by 0.02 every 10 iterations as the deforming templates approach the borders of

Table 1
Performance of our segmentation system over the target structures for a set of 20 T1-weighted 1 mm^3 resolution MRIs

System	Distance (mm)	Corpus callosum	Ventricles	Caudate nucleus	Hippocampus
Basic framework	Mean 95% sym	1.3	4.4	4.2	3.5
		2.2	5.4	4.7	8.2
With shape constraint	Mean 95% sym	1.4	4.8	3.9	3.2
		2.4	5.5	5.5	7.8
With distance constraint	Mean 95% sym	1.2	4.3	3.8	N/A
		2.2	5.1	4.2	N/A
With texture constraint	Mean 95% sym	0.2	2.2	2.1	2.5
		0.4	3.3	3.0	3.5
With shape and distance	Mean 95% sym	1.3	3.5	2.2	N/A
		2.2	4.7	4.9	N/A
With all constraints	Mean 95% sym	0.2	1.9	1.8	2.3
		0.4	2.8	2.3	3.5
With feedback rule	Mean 95% sym	0.2	1.8	1.6	2.1
		0.4	2.6	2.0	3.0

their target structures to allow them to better adapt to these borders (for structures with shape constraints).

Leak prevention. While classical rules control the behavior of the deformable templates, feedback rules control the applicability of the rules themselves to ensure that no mistake is being made during the deformation process. As such, they may be considered as metarules. To ensure that the deformable templates do not "leak" outside of the correct boundaries, we checked at each iteration that their distances (mean distances averaged over all vertices) to their associated shape-constrained projections stayed reasonable. We defined four structure-dependent thresholds: 3 mm for the corpus callosum, 4 mm for the caudate nucleus, 4 mm for the ventricles, and 5 mm for the hippocampus. These were based on the computed variability of each structure. Each time the threshold was reached, we increased the amount of regularization (α was increased by 0.2) and the shape constraint if used (λ was decreased by 0.1). As the deformation process went along, the structure-dependent thresholds were increased to allow for finer-scale deformations.

Results

Here, we present some qualitative and quantitative segmentation results for the four selected target structures.

Delineation protocol

For each structure, a delineation protocol (see Pitiot (2003)—Appendix B) was devised by expert neuroscientists and used to build the training set of 20 manual delineations which served as ground truth (those were traced in 3-D on twenty $256 \times 256 \times 124$ 1 mm^3 resolution SPGR T1-weighted MRIs of a group of normal elderly subjects). One should however keep in mind that protocols are always designed towards a specific a priori goal: for instance, comparing diseased and normal individuals, or the longitudinal study of a pathology, etc. They also have to ensure that the delineations can be carried out with reasonable accuracy by trained operators. This may at times require that the least visible parts of a structure be discarded (lest the manual delineation should introduce spurious edges and yield artificially high variability). Consequently, the manual delineations from the training sets, which we consider our gold-standard, may not always conform to standard anatomical expectations about the shape of the target structures (for instance, our gold-standard caudate nuclei have a very short tail, and the inferior horns of the ventricles are missing, see Fig. 6).

We then have to take these delineation protocols into account when computing the segmentation errors. For each target structure, we devised a semi-automated means (which often relied on semi-automated masking) to correct the computation of the misclassified voxels. Namely, for each structure, its associated manual delineation protocol was applied to the output of the automated algorithm to discard from the error computation voxels outside of the range defined by the protocol.

This correction step, although necessary, unfortunately introduces artificial imprecisions. One should therefore consider the segmentation results with caution. In particular, it seems wise to nuance the performances of an algorithm by taking the measured variability of the delineating human operators into account (Zou et al., 2002) described a means to compare the results of automated algorithms with those of experts when the latter exhibit substantial variability).

A few segmented target structures

Fig. 5 displays a few 3-D renderings of the target structures segmented with our system, along with the associated 2-D synchronized views, for a previously unseen T1-weighted MRI, with imaging characteristics similar to those in the learning set.

Fig. 5a illustrates the successful delineation of all four structures with the complete segmentation system (using all applicable constraints and rules). As explained earlier, the overall shapes of the targets might look different from standard anatomical expectations. However, those outlines conform with the established delineation protocols for our study and are thus considered satisfactory.

The relevance of the distance constraint is demonstrated in Fig. 5b: the lateral ventricles, caudate nuclei, and corpora callosa of the same MRI were segmented by our system *without* nonpenetration distance constraints (for the corpus callosum, ventricle, and caudate nucleus). As expected, these structures intersect.

In Fig. 5c, no shape constraints were used to segment the structures from the same MRI. Not surprisingly, comparison with Fig. 5a confirms the usefulness of shape models for controlling the deformation of templates when little or spurious intensity information is available: hippocampal segmentation greatly suffered from this lack of a priori shape knowledge (we even observed changes in topology). Conversely, in the absence of a shape constraint, we obtained, on the same MRI, a segmentation of the caudate nuclei which better agreed with anatomical expectation, in that they both presented longer tails (which were still within the guidelines of the delineation protocol).

Segmentation accuracy

The accuracy of our segmentation system was evaluated following the methodology presented in (Gerig et al., 2001). We used as error metrics the partial Hausdorff distance (defined below) and the mean absolute surface distance. We favored this error methodology over the computation of the false positive and false negative voxel ratios as it better illustrates the global behavior of the segmentation system (Gerig et al., 2001). In particular, it is less sensitive to small delineation errors As argued above, all segmentations were adjusted to take into account the delineation protocols (Fig. 6).

Given a deformed simplex mesh Π_j^t, its Hausdorff distance to a gold standard segmentation GS_j (represented by a set of 3-D voxels) is the largest distance between them both, computed in an asymmetric way, as the maximum (over all voxels v of a discretized version of Π_j^t) of the minimum Euclidean distance between v and its closest voxel w on GS_j:

$$H_{\text{asym}}\left(\Pi_j^t,\ GS_j\right) = \max_{v \in \Pi_j^t}\left(\min_{w \in GS_j} d_{\text{euclidean}}(v, w)\right) \tag{6}$$

This distance can be symmetrized by taking the maximum of both asymmetric measures. Finally, in view of its high sensitivity to outliers, we considered the 95% quantile of the symmetric Hausdorff distance. For efficiency reasons, we evaluated it by integrating the values of the Euclidean distance map of one surface along the contour of the other one, as described in (Gerig et al., 2001). A similar strategy allows a symmetric mean absolute distance to be computed between the deformed template and its target gold standard.

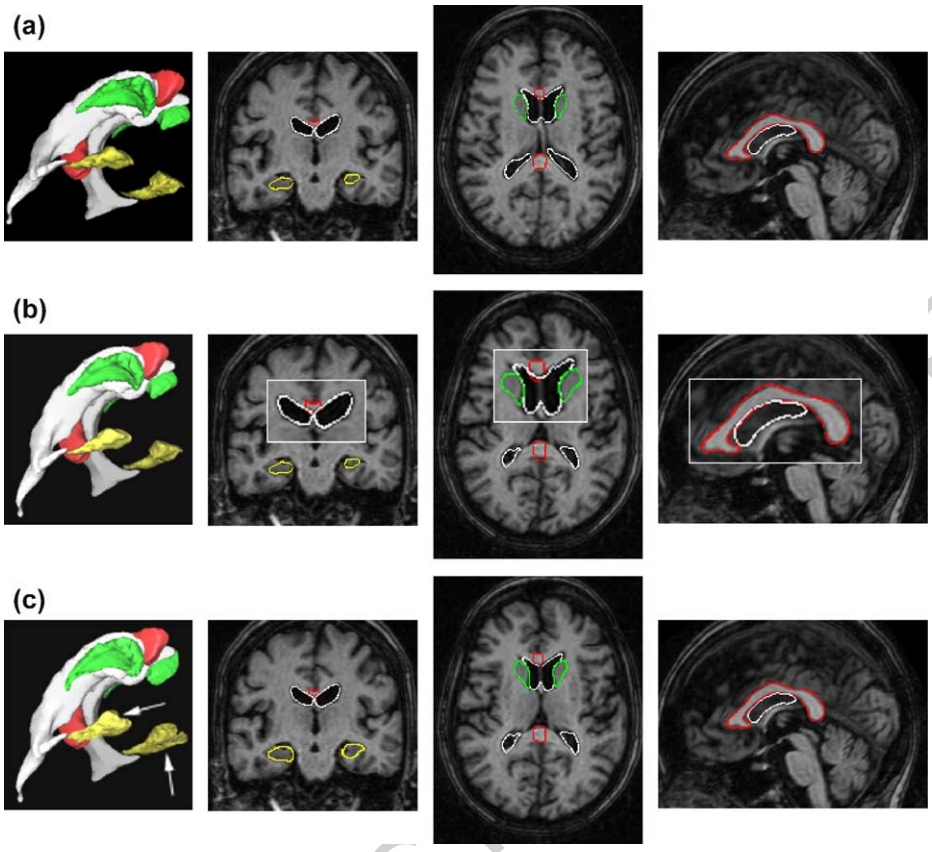

Fig. 5. Segmented target structures (in color) in a typical T1-weighted MRI: (a) with the complete segmentation system (all rules, all constraints); (b) without distance constraints (white rectangles shows magnified portions of the MRI where templates intersect); (c) without shape constraints (white arrows point to topological alterations of the hippocampi). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 1 reports both measures for all 4 structures, averaged over the 20 test instances (different from training instances). On average, segmentations were performed in approximately 6 min on a standard Pentium III, 1 GHz PC, for all four selected target structures. This does not include the training phase, which is done once and for all, in advance (and took about 20 h, mostly spent training the texture classifiers on all four structures). We present the accuracy of several versions of our segmentation system to demonstrate the influence of its various components.

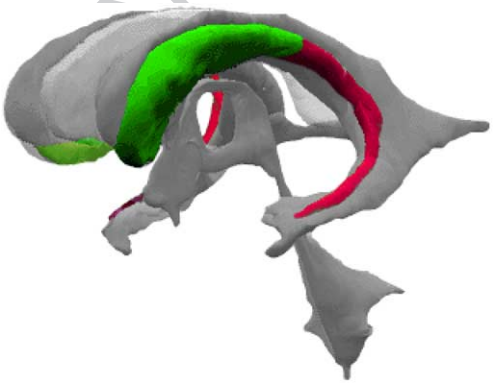

Fig. 6. Anatomically correct caudate nucleus (green + red) and manually segmented caudate nucleus (green) as obtained from the delineation variability. The nearby ventricles and corpus callosum are rendered in gray. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Noise robustness

The sensitivity of our segmentation methodology to imaging parameters was evaluated on a series of MRIs acquired on different scanners (from three different medical hospitals), with different sequences for various individuals. Twelve images were available with 3 MRIs per acquisition protocol.

Table 2
Performance of our segmentation system on the MRIs with different image characteristics and different learning sets (from different scanners) "c.c." denotes corpus callosum

Distance (mm)	c.c.	Ventricles	Caudate	Hippocampus
Standard learning set				
Mean 95% sym	0.3	2.9	1.8	2.7
	0.6	2.9	2.5	4.2
Adapted learning set				
Mean 95% sym	0.3	2.0	1.6	2.3
	0.5	2.7	2.2	3.8
Mixed learning set				
Mean 95% sym	0.3	2.2	1.7	2.5
	0.9	2.6	2.6	4.3

We tested the algorithm with the standard learning set whose labeled samples belonged to the homogeneous batch of MRIs we have used so far, with a learning set adapted to each acquisition protocol (with labeled samples coming from similar acquisitions in terms of parameters and scanners), and with a mixed learning set with labeled samples coming from the entire set of heterogeneous MRIs. In all three cases, the performances were evaluated on a test subset different from that in the learning set.

Table 2 reports both distance measures for all four structures.

Discussion

The explosive growth in brain imaging technologies has been matched by a tremendous increase in the number of investigations focusing on the structural and functional organization of the human brain. A pivotal first step in elucidating the correlation between brain structure and function, the precise segmentation and labeling of cerebral structures is a challenging objective in neuroinformatics in view of the wide variety of shapes and appearances these structures exhibit.

Our approach to that challenge, the brain MRI automated segmentation system we have detailed in this paper relies on expert prior knowledge about the target structures, their interrelationships and the characteristics of the surrounding tissues to achieve increased performances.

Segmentation accuracy

As illustrated in Table 1, segmentation of caudate and callosum were good and further improved with the use of shape, distance and texture constraints. A few odd-looking caudates (far from the mean shape) worsened the performances when a shape constraint was added (the 95% symmetric Hausdorff distance was worse than the one computed with the basic framework while the mean error is better). This demonstrates the difficulty of designing a learning set representative enough for the shape model to cover all the encountered shapes adequately. Clearly, a compromise must be found between too exhaustive a learning set which would induce poor shape constraints overall, and too specific a learning set which might improve performance in a particular niche only, to the detriment of everywhere else. Incidentally, principal component analysis may not be optimal for building a shape model representative of the true anatomical variability. Here also, a priori information could be used to build a better shape model.

The less accurate segmentation of the ventricles is explainable because our deformable templates cannot reach as far as the anterior apex of the inferior horns as they would have to go through partial volume effect voxels. However, even though these voxels were included in the delineation protocol, manual delineations exhibit a large variability in this area, which should be understood when considering the relatively lower performances of our automated approach. Furthermore, the model-to-manual maximum Hausdorff distances $H_{\text{asym}}(IT_j^t, \text{GS}_j)$ were good (2.2 with all constraints), since our approach correctly segmented the "reachable" parts of the ventricles.

To improve the ventricle segmentation performances, we defined on the ventricle simplex mesh a zone which covered the apex area, and for the vertices in this zone we locally decreased the amount of regularization ($\alpha = 0.01$) and increased the influence of the image force ($\beta = 0.5$, $\gamma_{\text{texture}} = 0.49$). We obtained a better mean error: 1.5 mm with a 2.2-mm symmetric Hausdorff measure.

Overall, the feedback rule was particularly effective, especially in reducing the maximal errors. However, poor contrast and noise hampered the hippocampus deformable templates. The importance of the texture constraint was particularly evident for this structure.

These segmentation results should however be considered in the light of the intra/interoperator variabilities associated with the delineations. Several operator variability measurements can be found in the literature, though a consensus is still lacking as to which error measure to use, which makes for another difficulty in comparing algorithms and studies. Reported values are reasonably small for the corpus callosum (2.5% of the callosal area in a study of the choice of the midsagittal section around which the callosa slices are delineated (Rauch and Jinkins, 1996), 1 mm RMS error for interoperator variability in (Narr et al., 2000). Nonetheless, the average interoperator error can be as high as 13% volume difference for the hippocampus (Obenaus et al., 2001), a difficult structure to outline in an MRI. Those variabilities affect both the a priori delineated samples in the learning sets which are used to build the shape and texture constraints, and the gold standards that we use to evaluate the performances of the automated segmentation algorithm. In this regard, our segmentation system seems to be doing as well as a manual operator could do.

Noise robustness

Noise robustness was good for all structures with nonetheless a significant decrease for the hippocampus, mostly due to the decreased performances of the texture filter for that structure. As expected, the segmentation results were improved when an adapted learning set was used.

Comparable results were obtained with a mixed learning set. We however observed a slight decrease in segmentation performance relative to those obtained with the adapted learning set. Mixed sets also induced a greater variability in the segmentation quality, which is probably explained by the lack of representativity of the learning set. Clearly, the learning set is much harder to make representative when a large variety of MR characteristics must be represented. Additional experiments with more samples in the learning set confirmed our intuition (this is mostly due to the performances of the texture filter).

At a glance, the segmentation quality for the corpus callosum or the ventricles is somewhat independent of the imaging characteristics. However, these have to be taken into account much more cautiously when more difficult structures have to be segmented, such as the caudate nucleus or the hippocampus.

Assessing performances

Overall, the segmentation performance compared favorably with those reported in the literature (Gerig et al., 2001; Pizer et al., 1999; Styner et al., 2003). A detailed comparison of segmentation performances is however trickier. Clearly, in view of the complexity of the segmentation problem, there are no general prescriptions for selecting a "good" segmentation algorithm. This choice must not only be driven by the image characteristics (type of noise and signal–noise ratio, texture characteristics, contrast of the target object with respect to surrounding pictorial elements, bias fields, etc.) but also by the possible usage constraints (algorithmic complexity with respect to available memory/CPU resources, time

limits if real-time applications are envisioned, etc.). Downstream treatments that follow this segmentation step must be considered as well (diagnosis, morphometric analysis, shape recognition, etc.). Consequently, assessing the true performance of a given segmentation tool per se is a difficult, if not ill-posed, task, as ground truth is elusive. It seems more sound to compare segmentation algorithms by measuring the overall quality of the complete chain of processes of which they usually are part: the best segmentation technique then becomes that which maximizes the overall system performance.

Conclusion

We presented a general framework for the automated segmentation of anatomical structures in brain MRIs. A hybrid combination of external and internal energies, modeling a variety of aspects of prior neuroanatomical knowledge, drives a series of 3-D deformable templates towards the boundaries of these target structures. Explicit rules, also derived from medical expertise, further increase the overall accuracy and robustness of the method.

The validity of this approach was demonstrated on the four selected target structures. The developed framework could of course readily be extended to segment additional structures. A more in-depth study of the multivariate relations between the various parameters of the deformation scheme and how they affect the accuracy of the match should also be conducted.

A number of additional rules could also increase the overall performance. In particular, additional feedback loops could be devised to tackle the segmentation of difficult images where robustness is more pressing, when lesions are apparent for instance. We could also incorporate segmentation strategies (sets of metarules) to monitor the number of times the error-checking rules (leak prevention, for instance) have been triggered and either interact with the human operator (to alert them about a particularly difficult segmentation, or require assistance in an area of the image, etc.) or select an entirely different set of parameters and shape/texture constraints.

References

Amit, Y., Kong, A., 1996. Graphical templates for model registration. IEEE Trans. Pattern Anal. Mach. Intell. 18 (3), 225–236.

Ardizzone, E., Peri, D., Pirrone, R., Palma, A., Peri, G., 2001. Knowledge based approach to intelligent data analysis of medical images. Proceedings of Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP'02).

Ayache, N., 2003. Epidaure: A research project in medical image analysis, simulation and robotics at INRIA. IEEE Trans. Med. Imag. 22 (10), 1185–1201.

Barra, V., Boire, J., 2001. Automatic segmentation of subcortical brain structures in MR images using information fusion. IEEE Trans. Med. Imag. 20 (7), 549–558.

Bernard, R., Lika, B., Pernus, F., 2001. Segmenting articulated structures by hierarchical statistical modeling of shape, appearance, and topology. Proceedings of MICCAI (MICCAI'01). pp. 499–506.

Blake, A., Zisserman, A., 1987. Visual Reconstruction. MIT Press.

Borgefors, G., 1984. Distance transformations in arbitrary dimensions. Comput. Vis. Graph. Image Process. 27, 321–345.

Brown, M., Wilson, L., Doust, B., Gill, R., Sun, C., 1998. Knowledge-based method for segmentation and analysis of lung boundaries in chest X-ray images. Comput. Med. Imaging Graph. 22, 463–477.

Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N., 2003 (Feb.–March). Iconic feature based nonrigid registration: The PASHA algorithm. CVIU–Special Issue on Nonrigid Registration 89 (2–3), 272–298.

Collins, D., Holmes, C., Peters, T., Evans, A., 1995. Automatic 3D model-based neuroanatomical segmentation. Hum. Brain Mapp. 3 (3), 190–208.

Collins, D.L., Zijdenbos, A.P., Paus, T., Evans, A.C., 2003. Use of registration for cohort studies. In: Hajnal, J., Hawkes, D., Hill, D. (Eds.), Medical Image Registration. Kluwer.

Cootes, T.F., Hill, A., Taylor, C.J., Haslam, J., 1994. Use of active shape models for locating structures in medical images. Image Vis. Comput. 12 (6), 355–366.

Cootes, T., Edwards, G., Taylor, C., 1998. Active appearance models. Proc. of ECCV. pp. 484–498.

Csernansky, J.G., Joshi, S., Wang, L., Haller, J.W., Gado, M., Miller, J.P., Grenander, U., Miller, M.I., 1998. Hippocampal morphometry in schizophrenia by high dimensional brain mapping. Neurobiology 95 (19), 11406–11411.

Delingette, H., 1999. General object reconstruction based on simplex meshes. Int. J. Comput. Vis. 32 (2), 111–146.

Fleuté, M., Lavallée, S., Julliard, R., 1999. Incorporating a statistically based shape model into a system for computer-assisted anterior cruciate ligament surgery. Med. Image Anal. 3 (3), 209–222.

Gerig, G., Jomier, M., Chakos, M., 2001. Valmet: A new validation tool for assessing and improving 3D object segmentation. Proc of MICCAI. pp. 516–528.

Holly, L., Foley, K., 2003. Intraoperative spinal navigation. Spine 28 (15), 554–561.

Li, H., Deklerck, R., De Cuyper, B.D., Hermanus, A., Nyssen, E., Cornelis, J., 1995. Object recognition in brain CT-scans: Knowledge-based fusion of data from multiple feature extractors. IEEE Trans. Med. Imag. 14 (2), 212–229.

Matesin, M., Loncaric, S., Petravic, D., 2001. A rule-based approach to stroke lesion analysis from CT brain images. Proc. of Second International Symposium on Image and Signal Process. and Analysis. pp. 219–223.

McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. Med. Image Anal. 1 (2), 91–108.

Montagnat, J., Delingette, H., 1998. Globally constrained deformable models for 3D object reconstruction. Signal Process. 71 (2), 173–186.

Montagnat, J., Delingette, H., Ayache, N., 2001. A review of deformable surfaces: topology, geometry and deformation. Image Vis. Comput. 19, 1023–1040.

Narr, K., Thompson, P., Sharma, T., Moussai, J., Cannestra, A., Toga, A., 2000. Mapping morphology of the corpus callosum in schizophrenia. Cereb. Cortex 10 (1), 40–49.

Obenaus, A., Yong-Hing, C., Tong, K., Sarty, G., 2001. A reliable method for measurement and normalization of pediatric hippocampal volumes. Pediatr. Res. 50, 124–132.

O'Sullivan, B., Shah, J., 2003. New TNM staging criteria for head and neck tumors. Semin. Surf. Oncol. 21 (1), 30–42.

Ourselin, S., Roche, A., Subsol, G., Pennec, X., Ayache, N., 2001. Reconstructing a 3D structure from serial histological sections. Image Vis. Comput. 19 (1 – 2), 25–31.

Pitiot, A., 2003. Automated Segmentation of Cerebral Structures Incorporating Explicit Knowledge. PhD thesis, Ecole des Mines de Paris.

Pitiot, A., Toga, A., Ayache, N., Thompson, P., 2002a. Texture based MRI segmentation with a two-stage hybrid neural classifier. World Congress on Computational Intelligence / INNS-IEEE International Joint Conference on Neural Networks WCCI-IJCNN'02.

Pitiot, A., Toga, A., Thompson, P., 2002b. Adaptive elastic segmentation of brain MRI via shape-model-g-uided evolutionary programming. IEEE Trans. Med. Imag. 21 (8), 910–923.

Pizer, S., Fritsch, D., Yushkevich, P., Johnson, V., Chaney, E., 1999. Segmentation, registration, and measurement of shape variation via image object shape. IEEE Trans. Med. Imag. 10 (18), 851–865.

Poupon, F., Mangin, J.-F., Hasboun, D., Metrology, I., Frouin, V., 1998. Multi-object deformable templates dedicated to the segmentation of brain deep structures. Proceedings of MICCAI (MICCAI'98). pp. 1134–1143.

Rauch, R.A., Jinkins, J.R., 1996. Variability of corpus callosal area measurements from midsagittal MR images: effect of subject placement within the scanner. Am. J. Neuroradiol. 17 (1), 27–28.

Rey, D., Subsol, G., Delingette, H., Ayache, N., 2002. Automatic detection and segmentation of evolving processes in 3D medical images: application to multiple sclerosis. Med. Image Anal. 6 (2), 163–179.

Staib, L., Duncan, J., 1992. Boundary finding with parametrically deformable models. IEEE Trans. Pattern Anal. Mach. Intell. 14 (11), 1061–1075.

Studholme, C., Hill, D.L.G., Hawkes, D.J., 2003. Incorporating connected region labelling into automated image registration using mutual information. IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA'96). pp. 23–31.

Styner, M., Gerig, G., Lieberman, J., Jones, D., Weinberger, D., 2003. Statistical shape analysis of neuroanatomical structures based on medial models. Med. Image Anal. 7 (3), 207–220.

Thompson, P., Narr, K., Blanton, R., Toga, A., 1997. Mapping structural alterations of the corpus callosum during brain development and degeneration. In: Zaidel, E., Iacoboni, M. (Eds.), The Corpus Callosum. MIT Press.

Thompson, P.M., Woods, R.P., Mega, M.S., Toga, A.W., 2000. Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain. Hum. Brain Mapp. 9 (2), 81–92.

Thompson, P.M., Narr, K.L., Blanton, R.E., Toga, A.W., 2003. Mapping structural alterations of the corpus callosum during brain development and degeneration. Proceedings of the NATO ASI on the Corpus Callosum, .

Tsai, A., Wells, W., Tempany, C., Grimson, E., Willsky, A., 2003. Coupled multi-shape model and mutual information for medical image segmentation. Proceedings of Information Processing in Medical Imaging (IPMI'03). pp. 185–197.

Turk, M., Pentland, A., 1991. Eigenfaces for recognition. J. Cogn. Neurosci. 3 (1), 71–86.

Yang, J., Staib, L., Duncan, J., 2003. Neighbor-constrained segmentation with 3D deformable models. Proceedings of Information Processing in Medical Imaging (IPMI'03). pp. 198–209.

Zou, K., Wells, W., Kaus, M., Kikinis, R., Jolesz, F., Warfield, S., 2002. Statistical validation of automated probabilistic segmentation against composite latent expert ground truth in MR imaging of brain tumors. Proceedings of MICCAI (MICCAI'02). pp. 315–322.