

Adaptive Elastic Segmentation of Brain MRI via Shape-Model-Guided Evolutionary Programming

Alain Pitiot, Arthur W. Toga, and Paul M. Thompson*

Abstract—This paper presents a fully automated segmentation method for medical images. The goal is to localize and parameterize a variety of types of structure in these images for subsequent quantitative analysis. We propose a new hybrid strategy that combines a general elastic template matching approach and an evolutionary heuristic. The evolutionary algorithm uses prior statistical information about the shape of the target structure to control the behavior of a number of deformable templates. Each template, modeled in the form of a B -spline, is warped in a potential field which is itself dynamically adapted.

Such a hybrid scheme proves to be promising: by maintaining a population of templates, we cover a large domain of the solution space under the global guidance of the evolutionary heuristic, and thoroughly explore interesting areas.

We address key issues of automated image segmentation systems.

- The potential fields are initially designed based on the spatial features of the edges in the input image, and are subjected to spatially adaptive diffusion to guarantee the deformation of the template. This also improves its global consistency and convergence speed.
- The deformation algorithm can modify the internal structure of the templates to allow a better match.
- We investigate in detail the preprocessing phase that the images undergo before they can be used more effectively in the iterative elastic matching procedure: a texture classifier, trained via linear discriminant analysis of a learning set, is used to enhance the contrast of the target structure with respect to surrounding tissues.
- We show how these techniques interact within a statistically driven evolutionary scheme to achieve a better tradeoff between template flexibility and sensitivity to noise and outliers.

We focus on understanding the features of template matching that are most beneficial in terms of the achieved match. Examples from simulated and real image data are discussed, with considerations of algorithmic efficiency.

Index Terms—Adaptive algorithms, deformable templates, hybrid evolutionary algorithms, segmentation.

Manuscript received September 13, 2000; revised May 31, 2001. This work was supported in part by the NCRR under a P41 Resource Grant RR13642; RR05956, in part by the National Institutes of Health (NIH) through the National Institute of Mental Health (NIMH) under Grant MH65166, in part by the National Library of Medicine under Grant LM/MH05639, and in part by the Human Brain Project, which is funded jointly by NIMH and NIDA under Grant P20 MH/DA52176. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was R. Leahy. *Asterisk indicates corresponding author.*

A. Pitiot and A. W. Toga are with the Reed Neurological Research Center, Laboratory of Neuro Imaging, Department of Neurology, UCLA School of Medicine, Los Angeles, CA 90095-1769 USA.

*P. M. Thompson is with the Reed Neurological Research Center, Rm. 4238, Laboratory of Neuro Imaging, Department of Neurology, UCLA School of Medicine, 710 Westwood Plaza, Los Angeles, CA 90095-1769 USA (e-mail: thompson@loni.ucla.edu).

Digital Object Identifier 10.1109/TMI.2002.803124

I. INTRODUCTION

IMAGE segmentation, that is, the division of an image pixel set into a number of disjoint regions that are homogeneous with respect to a set of characteristics, is a major goal of computer vision and image processing. Effective identification and labeling of anatomical structures in possibly complex magnetic resonance imaging (MRI) proves to be especially challenging, given the wide variety of shapes and intensities each anatomical structure can present. Automated image segmentation can be used to assist medical diagnosis or analysis, and calls for high precision: indeed, the quality of the diagnosis often depends on how accurately the various structures in the image can be identified. Similarly, brain atlases (see [1] for a review), which can guide stereotaxic neurosurgical procedures and provide a precise quantitative framework for multimodality brain mapping, turn out to be rather tedious to build, as many components typically have to be interactively outlined. Thus, automated segmentation systems are powerful tools to help in drawing consistent diagnosis from a number of images [2], to classify pictorial data [3], or collect statistical information on anatomical variability [4].

This paper addresses the problem of retrieving the boundary of a brain structure from MRI data. We approach the issue of boundary finding as a process of fitting a series of deformable templates to the target contour. A template is modeled by using a parameterized curve whose coefficients are iteratively updated to minimize an objective function. This function measures the match between the deformed template and a modified edge image, and the elastic deformation energy required in the warping process. Each template evolves, within an adaptive gradient descent scheme, in a potential field which is itself progressively refined during the deformation process. We introduce an evolutionary heuristic to control the behavior of the various templates by 1) incorporating statistical constraints to bias the deformations toward a range of shapes derived from a statistical analysis of a learning set, and 2) selecting and favoring the most promising templates at each round.

In this paper, our goal is to determine and analyze the features of template matching algorithms that play a determining role in terms of their efficiency and accuracy. We try to address in a methodical fashion what we identify the following as the main issues of fully automated segmentation algorithms:

- initial positioning of the deformable template with respect to the target structure;
- capture range of the target structure;
- tradeoff between flexibility of the template and robustness to noise and decoys;

- use of *a priori* statistical information about the shape of the target structure;
- search for a global solution in a computationally effective fashion.

The remainder of the paper is organized as follows. We review in Section II template matching methods and hybrid evolutionary algorithms. In Section III, we introduce our elastic approach to shape matching and the associated template modification scheme in the two-dimensional (2-D) case. Section IV discusses the preprocessing stage and how it influences the efficiency and accuracy of the template matching scheme. Section V presents the shape model guided evolutionary heuristic that controls the deformable templates. Finally, several experiments are analyzed in Section VI, aiming to parameterize highly variable structures in a large medical image database. We discuss the strengths and limitations of our approach. We also consider an extension of our 2-D model to the three-dimensional (3-D) case.

II. BACKGROUND

A. Elastic Template Matching

Depending on their underlying structure, most elastic template approaches can be partitioned into two classes: *free-form* and *parametric*.

Free-form methods do not assume any specific global structure for the template and, therefore, can be applied to a variety of domains. The only possible constraints are local continuity and smoothness, which provide considerable flexibility to represent arbitrarily complex shapes. The potential field induced by a salient edge image is used to iteratively deform the template over an image. An early elastic deformable model in this category was proposed by Burr [5], who used a set of lines as a template. His matching approach consisted of designing an elastic model for the template, and letting it deform under local forces. Eventually, correspondence was defined by tracking an image point over the set of image data. In the active contour model proposed by Kass *et al.* [6], splines were used in favor of a set of lines. The “snake”—a deformable spline—was actively warped, seeking a local minimum in a potential field by adapting its shape and position. The potential field was made up of three factors: an internal spline force, the attraction force of the image, and an external constraint force. Among other deformable techniques, Sapiro [7] introduced color snakes evolving in a color or vector-valued image: object boundaries were then obtained as geodesic-weighted distance curves in a Riemannian space. Yet another approach, presented by Sethian and Osher [8], used level sets as a convenient means to handle complex topological changes: here, a raster image template was deformed and boundaries were inferred from the motion of its zero-level set. Lately, Leventon *et al.* [9] proposed a way to embed statistical information in the level set formulation.

Parametric methods, on the other hand, assume some prior knowledge of the geometrical shape, encoded in the form of a small number of parameters. This can be either as 1) the image of a prototype template under a parametric mapping or 2) a collection of parameterized curves or surfaces.

The first method consists of two elements: a prototype template that captures the basic shape, and a parametric mapping, which governs the template deformation during the warping process and determines the degree of freedom of the deformation. Its high versatility is due to the fact that different choices of the model template and of the deformation parameters will yield different warped outcomes. The prototype template may consist of a “mean” shape computed from a learning set. The parametric mapping may also be tuned so as to reward the most “natural” transformation—in the specific application domain—and penalize those which are less likely to be appropriate. When a learning set is unavailable, a uniform distribution of the possible deformations, a Gibbs prior, or a Tikhonov stabilizer is assumed to serve as prior knowledge. Grenander and Miller [10] developed a systematic pattern-theoretic framework to deal with variable shape structures under parametric mappings. In their paper on “active shape models,” Cootes *et al.* [11] used linear combinations of the eigenvectors of the covariance matrix of deviations from the mean shape to determine the parameters of the possible deformations. The mean shape was previously drawn from a training set of labeled objects. Their model implemented a learning process where the underlying structure of a shape class as well as the most relevant transformations that occur within a class are derived from the example set. A similar scheme was used by Jain [12]. The prototype template consisted of a bitmap that encoded the shape of the object. The whole bitmap was transformed according to probabilistic transformations. A Bayesian scheme controlled the matching process by maximizing a posterior probability based on both the severity of the template deformation and the agreement between the deformed template and the edge image. Another benefit with Bayesian models is that other statistical estimators (such as MVE and a measure of confidence in the solution) can be derived [13].

In the second parametric approach to elastic matching, the characteristics of the template are described by a set of curves/surfaces, controlled by a set of parameters. The template geometry entirely depends on the values of these parameters. The deformable template seeks local minima in a potential field based on salient image features. The warping process is conducted by updating the parameters. Staib and Duncan [14], for instance, used elliptic Fourier decomposition to encode the boundaries of a template. The parameters to be evolved were the Fourier coefficients. They introduced a likelihood functional, based on the correlation between the template and the edge image, to be maximized under a Bayesian model. Székely *et al.* [15] also chose Fourier contours to model an elastic template and segment 2-D or 3-D objects from MRI data. In [3], Del Bimbo and Pala favored fourth-order *B*-splines to model deformable templates, whose coefficients were modified in a gradient descent scheme.

Our elastic matching scheme falls into the second category of parametric methods.

B. Hybrid Evolutionary Algorithms

As opposed to sequential methods which progress step-by-step from an initial solution to the optimization problem (here, the initial estimate of the target structure) to a neighbor solu-

tion by performing some elementary modification (parameter updates), evolutionary techniques deal with a *population* of solutions [16]. The pivotal idea is to use the collective properties of a group of solutions to search the global solution space. Once an initial population of solutions has been generated, it is improved by a cyclic two-phase evolution process: an exploration phase (cooperation) and an exploitation phase (self-adaptation). In the exploration phase, solutions of the current population exchange information (recombining in a probabilistic fashion) with the aim of producing new solutions which inherit good attributes. In the exploitation phase, solutions change their internal structure without any interaction with the other members. Optimal solutions are thus evolved by iteratively producing new generations of solutions and selectively favoring the most successful ones.

A promising innovation in the field of general heuristics, hybrid evolutionary systems embed a sequential local search within the framework of a population-based strategy. The combination of a local search and a global control heuristic yields an algorithm with significantly better performance than that of both search methods running separately [17]. In a hybrid algorithm, the role of the sequential search is to explore thoroughly disjoint areas of the solution space, whereas the evolutionary procedure provides global guidance through the space.

Hybrid evolutionary techniques have been successfully applied to many combinatorial optimization problems: the traveling salesman problem [18], the quadratic assignment problem [19], the bin-packing problem [20] or the graph coloring problem [21] among others. To the best of our knowledge, there is, however, only a very limited body of applications related to template matching. In [22], Cootes *et al.* used genetic algorithms (a variety of evolutionary methods) to generate a number of starting positions to be evolved by their active shape model (ASM). They briefly mentioned the possibility of incorporating the ASM directly into the genetic search but did not elaborate on it.

Our approach proceeds along these lines by integrating a local elastic template matching method within a global evolutionary scheme.

III. AN ELASTIC DEFORMABLE MODEL

We describe here the deformable template scheme we implemented in the 2-D case. We will discuss later, in Section VI-D, how it could be extended to the 3-D case.

A. Deformable Template

The input data image is defined by its intensity value at each point $I: \Omega \rightarrow R; \Omega \subset R^2$. The initial template (prototype) consists of a 2-D curve $\tau: [0, 1] \rightarrow R^2, s \rightarrow \tau(s)$ parameterized with respect to a scalar, s .

The algorithm's goal is to find a pictorial object in I whose overall boundary shape fits that of τ . To decrease the discrepancy between the actual imaged object and the prototype template, we allow the latter to undergo deformation. Let $\theta: [0, 1] \rightarrow R^2, s \rightarrow \theta(s)$ be the parameterized deformation, then the deformed template is given by $\varphi: [0, 1] \rightarrow R^2$ s.t. $\varphi(s) = \tau(s) + \theta(s), \forall s$.

To guide the deformation process and drive the template toward the required object shape, we introduce a compound functional \mathcal{F} , whose minimum we aim to determine. \mathcal{F} is made up of two components: the match between the deformed template and a modified form of the edge image, and the elastic deformation energy spent in the warping process.

We use for the matching term an approach based on intensity features integrated along the template

$$M(\varphi) = \int_0^1 \Xi\{I_E\}(\varphi(s)) ds$$

where I_E is a binary edge image drawn from I and $\Xi: (R^2 \rightarrow R) \rightarrow (R^2 \rightarrow R)$ is a smoothing operator.

M is related to the likelihood of the contour of the warped template being the true boundary of the target pictorial element, once the parameters of φ are given.

Note that, as a measure of the match, M is particularly sensitive to local geometric variations in the template and the target pattern. Indeed, it does not encompass any information about the template's global shape, hence the need for a moderate smoothing filter Ξ which reduces the measure's local sensitivity without sacrificing its accuracy and capture range. As it is, M is a poor approximation of the human perception of shape similarity. Other measures based on shape decomposition have been proposed in the literature (see Veltkamp [23] for a survey of shape similarity measures). However, they often turn out to be computationally expensive and not as suitable as expected within an elastic matching framework.

Using the thin-plate under tension model [24], the elastic deformation energy of the template is

$$\begin{aligned} E(\varphi) &= \alpha.S(\varphi) + \beta.B(\varphi) \\ &= \alpha \cdot \int_0^1 \left[\left(\frac{d\theta_x}{ds} \right)^2 + \left(\frac{d\theta_y}{ds} \right)^2 \right] ds \\ &\quad + \beta \cdot \int_0^1 \left[\left(\frac{d^2\theta_x}{ds^2} \right)^2 + \left(\frac{d^2\theta_y}{ds^2} \right)^2 \right] ds \end{aligned}$$

where S and B respectively approximate the stretching energy and the bending energy associated with the warped template $\varphi = \tau + \theta$, and where α and $\beta \in R$ are weights.

The energy component can be regarded as a constraint placed on the extent to which the template can stretch and bend: by increasing α and β , one limits the range of possible evolution, that is the degree of freedom of the deformation process.

Finally, we get

$$\mathcal{F} = \alpha.S + \beta.B - \gamma.M$$

where γ is the relative weight of the attraction exerted by the image edges.

Note that since E makes use of the first and second derivatives of θ only, pure translations, achieved when θ is a constant, or rotations of the template (isometries) do not augment the deformation energy. Also, since E is independent of τ , the degree of complexity of the prototype template does not bias the warping process: this creates a general template matching scheme that can cope with a variety of shapes.

B. Numerical Solution

We use cubic B -splines to model both the prototype τ and the elastic deformation template θ . Our choice was motivated by the following observations.

1) Cubic B -splines are regular enough to cope with the first- and second-order derivatives involved in the compound functional \mathcal{F} (see Zhong [25] or Blake *et al.* [26] for a comprehensive study of the respective merits of B -splines, Fourier or wavelet decompositions).

2) They offer the potential for a large variability while still maintaining the required degree of smoothness, as opposed to Bézier curves for instance, for which additional constraints have to be added to preserve C^2 continuity. Besides, only the control points have to be warped, which results in a significant savings over the use of explicit parametric polynomial functions.

3) Finally, open curves are easily modeled by triplicating the end-points, and closed curves are readily accommodated by equating the last three control points with the first three.

We get the following expressions for τ and θ :

$$\tau = (\tau_x, \tau_y): [0, 1] \rightarrow R^2, \quad s \rightarrow \tau(s) = \sum_{i=1}^{N_\tau} B_i(s) \cdot P_\tau^i$$

and

$$\theta = (\theta_x, \theta_y): [0, 1] \rightarrow R^2, \quad s \rightarrow \theta(s) = \sum_{i=1}^{N_\theta} B_i(s) \cdot P_\theta^i$$

where $\{P_\tau^i = (P_{\tau_x}^i, P_{\tau_y}^i), 1 \leq i \leq N_\tau\}$ are the N_τ control points of the fourth-order B -splines τ ; $\{P_\theta^i = (P_{\theta_x}^i, P_{\theta_y}^i), 1 \leq i \leq N_\theta\}$ are the N_θ control points of the fourth-order B -splines θ ; and B_i are the fourth-order B -spline blending functions.

We now have to minimize the compound functional \mathcal{F} with respect to variables $\{P_\theta^1, \dots, P_\theta^{N_\theta}\}$. A variety of minimization techniques have been presented in the literature (see [27] for a review) using a combination of deterministic/stochastic discrete/continuous algorithms. In Section V, we will discuss how the use of an evolutionary heuristic enables an effective and broad exploration of the solution space by controlling the initial shape and pose of each template. Such a method manages to quickly discover the many attraction basins of the solution space. Here, we choose a deterministic adaptive gradient descent scheme to evolve the templates and reach the bottom of those basins.

Gradient Descent: At iteration $t_0 = 0$, $\Pi_\theta(t_0) = (P_\theta^1(t_0) = (0, 0), \dots, P_\theta^{N_\theta}(t_0) = (0, 0))^T$ is the parameter vector associated with the prototype template $\varphi = \tau$.

The parameters of θ are then iteratively updated as follows:

$$\forall t > 0, \quad \Pi_\theta(t+1) = \Pi_\theta(t) - \varepsilon \cdot \nabla \mathcal{F}$$

$$\text{with } \nabla \mathcal{F} = \left(\frac{\partial \mathcal{F}}{\partial P_\theta^1}, \dots, \frac{\partial \mathcal{F}}{\partial P_\theta^{N_\theta}} \right)^t.$$

Strictly speaking, we do not modify the B -spline φ point by point but rather update its control points. This can induce oscillations in the matching process since only the control points are modified, whereas the objective function is evaluated over the whole curve. Nonetheless, considering a slightly filtered edge image in the evaluation of the objective function \mathcal{F} moderates

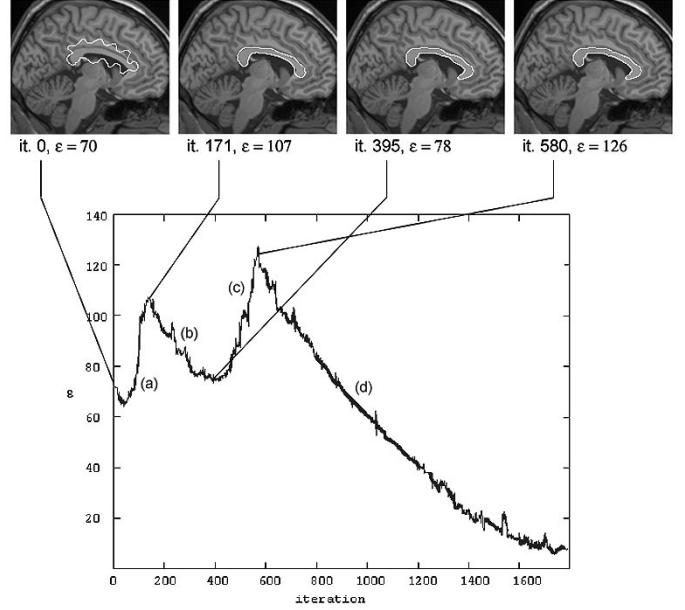


Fig. 1. Dynamics of the incrementation parameter ε . A deformable template (in white) evolves in a dynamic potential field (gray levels) within an adaptive gradient descent scheme controlled by ε .

this effect by introducing information about the neighborhood of the control points.

Adaptive Gradient Descent: In view of the complexity of the potential field induced by the edges of an input MRI and the form of the compound objective function \mathcal{F} , there are no general prescriptions for selecting an appropriate incrementation parameter ε for the gradient descent scheme, in order to avoid oscillation and converge to an acceptable local minimum quickly. Further, no fixed increment is usually suitable for the entire deformation process. Both problems can, however, be solved by adapting ε to the local structure of the potential surface.

We start with a given ε_0 and monitor the value of \mathcal{F} after each gradient descent step.

- If \mathcal{F} decreases, we assume that the minimization process evolves in the right direction and ε is increased by a multiplicative factor i_ε : the deformation process speeds up.

- Conversely, an increase in the value of \mathcal{F} (an “accident” during the minimization process) is taken as an indication that the step which occurred was too long: ε is decreased by a factor d_ε and the last change is canceled. This process of reduction is repeated until a step that decreases the energy value is found (this will inevitably happen since the search direction is that of the negative gradient). Note that with a standard gradient descent scheme, such a case cannot occur since the path followed always goes downhill. With $\varepsilon > 1$, on the other hand, there is no more guarantee as to the monotonicity of \mathcal{F} .

Fig. 1 displays the dynamics of the incrementation parameter ε in a typical example.

- ε is initialized with $\varepsilon_0 = 70$. At early stages of the deformation process [Fig. 1(a)], the template’s control points follow a regular path through the smooth potential field, since they are relatively far away from the edges: ε regularly increases.

- Later on, when the global shape of the template more accurately fits that of the pattern in the input image, the potential

field is refined a first time (see Section III-D) and the motion of the control points becomes more sensitive to local potential field variations: the incrementation parameter, ε , decreases accordingly to adapt the gradient descent scheme [Fig. 1(b)].

• Another increasing phase [Fig. 1(c)] follows, during which the template, which already achieves a good match over most of the target pattern, has to stretch to fit the U-shape in the bottom-right corner of the input image: ε increases until a new potential refinement occurs to allow a more precise adaptation: ε then progressively decreases [Fig. 1(d)] until the matching process is stopped. As it is, ε can be regarded as an indicator of good convergence.

Also, the value of i_ε and d_ε can be fixed once and for all. We chose i_ε slightly larger than unity ($i_\varepsilon = 1.05$) in order to avoid frequent occurrences of an increase of \mathcal{F} , since in such cases the evaluation is wasted. Similarly, we took d_ε to be significantly less than unity ($d_\varepsilon = 0.6$) so that the algorithm quickly reverts to finding a step which decreases \mathcal{F} , again to minimize the wasted computation.

Note that this approach bears resemblance to the “bold driver” technique developed by Vogl *et al.* [28] to train neural networks: in that case, an error function based on the discrepancy between the output of the neural network and the expected output has to be minimized by iteratively updating the various weights (parameters) of the network. It is also similar to the choice of kappa, the over-relaxation parameter, in successive over-relaxation.

C. Dynamic Number of Control Points

Intuitively, the larger the number of control points in the deformable template, the better the local fidelity of the match is likely to be. However, using a large number of points during early iterations gives an unnecessarily high computational cost and induces too many local minima thus complicating the search for a solution. Initial steps usually let the deformable template modify its global shape to that of the pattern in the image, and only later on are more points required to achieve a precise adaptation. This creates the need for a structural modification scheme in the form of a dynamic number of control points.

First, we need to supervise the dynamic behavior of the number of control points. The simplest method consists of regularly increasing the number of points every n_C iterations (n_C constant). The obvious drawback lies in the relative independence of this policy from the accuracy of the deformation process. A more sophisticated approach based on the quality of the match would introduce new points only when the match reaches a given threshold; this threshold would then be increased by a constant or variable factor and the process would go on until the match reaches the new threshold. Similarly, a statistical method would wait for the match to stabilize around a sufficiently high value, assuming that for a given α and β the template cannot achieve a better match because the required stretch and bending energies are too high. Thus, introducing more control points would enable the template to warp further toward the edges. Unfortunately, the potential of the matching functional M to oscillate makes it difficult for an automated system to distinguish between a stabilization phase and a slowly increasing one. Yet another technique would monitor the stretch and bend energies B and S , and introduce new points when




Fig. 2. Dynamic number of control points and regional insertion.

these energies reach a certain threshold. However, appropriate values depend on the nature of the target shape: a square, for instance, will require a high value of B at the corners and very low values along the sides, while a circle is more uniform. Presently, none of those approaches is a panacea. Incidentally, the template could very well be warping over “decoy features” (i.e., edge features which do not belong to the target object boundary) in which case adding new control points would enable it to achieve an even better adaptation to the decoy shape. We chose to increase the number of control points every n_C iterations (n_C constant) as this achieved the best results in practice: although this regime is independent of the state of the deformation process, it is also less sensitive to noise.

Once we have decided to add new points, we determine the segment (whose shape is parameterized by four control points) of the B -spline that achieves the lowest match by computing the partial match of each segment over the diffused edge image (to avoid getting a large number of null matches). We then locally refine that segment, inserting new control points.

Adding new control points somewhat overrides the influence of the weighting parameters α and β , since the first and second derivatives are now computed along a smaller segment. In Fig. 2(a), smoothness constraints are too high to allow the template to warp further toward the boundaries of the *caudate nucleus*. We introduce new control points along the segment that presents the lowest match (arrows). The new template then achieves a greater deformation along that segment and is able to match the inferior portion [Fig. 2(b)].

D. Adaptive Edge Potential

An elastic deformation method such as that of Jain [12], Cohen and Cohen [29], Davatzikos [30], or ours requires that the gradient of the edge features of the input image be extracted. That edge potential has to be tuned carefully, as its accuracy conditions the efficiency of the warping scheme. Namely, if the potential field is too sharp, the deformable template will not be driven toward the image contour if it does not overlap an area with nonzero amplitude [Fig. 3(a)]. Similarly, if the potential is too smooth, the template is misled toward amplitude maxima that do not correspond to actual edges but to areas in between them. In Fig. 3(b), the edge potential is not sufficiently well-defined: the capture areas of the main edges (green lines) are all “melted” together: the deformable template will be unable to discriminate between these and will warp toward arbitrary

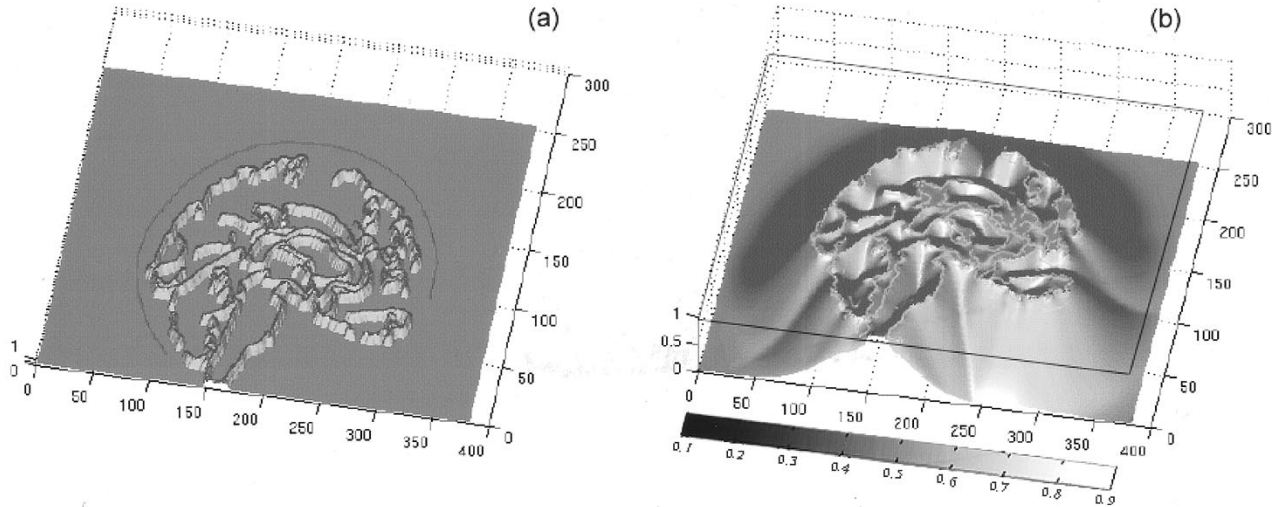


Fig. 3. Capture range of the potential field: (a) the potential is too sharp for the initial template to be captured and deformed; (b) the potential is too smooth for the adaptation process to be accurate.

maxima of amplitude. However, smooth edge potentials induce regular nets of iso-curves and their associated normals. These potential fields turn out to be particularly useful for adjusting the overall shape of the template to that of the pattern in the image at coarse levels.

The issue of the potential field's capture range has been addressed in the literature [29]–[31] and many methods have been developed to increase it. In most cases, the potential is computed before the deformation process starts, based on the initial edge image. Yet at any time during the deformation process, we wish to have an “optimal” potential field with a support just wide enough to let the template warp toward the edges but with sufficient local structure to achieve precise adaptation.

Diffusion of the Edge Potential: In Section IV-B, we present a texture filter capable of selectively enhancing the contrast of the target structure. That filter effectively decreases the fuzziness of the transition boundaries (essentially due to partial volume effects or line-type edges) between the structure to be segmented and the surrounding tissues: we can thereafter assume a step edge model and we choose the Canny–Deriche approach to compute the gradient magnitude of the input image. We then have to decide upon a scale for the Gaussian smoothing. In Section IV-C, we explain how to choose an initial scale based on the configuration of the image. We detail, here, how the scale dynamically evolves along with the deformation process.

Our approach is based on the following observation: points of the deformable template that belong to the B -spline segment with the lowest local match are those for which the mean distance to the closest edge is maximal (that is, they belong to the segment with the largest undirected Hausdorff distance). Every n_R iterations (n_R integer constant), the new scale is set to the 95% quantile of the Hausdorff distance. We chose a partial distance to prevent the edge potential from being too coarse in case only a small proportion of the template was far from the closest edges.

Fig. 4 shows six steps of the deformation process where the potential field is progressively refined to achieve a better match. By dynamically computing the Gaussian scale, we manage to diffuse while maintaining a good separability among features

of the potential field. Additionally, progressively refining the edge potential helps maintain the convexity of the search space and prevent the minimization algorithm from converging toward obvious local minima.

Monotone Remapping: In this scheme, the potential's amplitude decreases very rapidly with increasing distance from an edge. We thus used a monotone transfer function $f(x) = x^\alpha$ with $\alpha \in \mathbb{R}$, ($\alpha = 0.1 < 1$ for instance) to remap the interval $[0, 1]$ onto itself in order to derive an enhanced potential field. The remapped potential field still strictly increases toward the edges but presents better characteristics with globally higher gradient values and increased smoothness. Also, it re-equilibrates the speed at which each control point is driven toward the edges.

Note that for efficiency reasons, a number of potential fields are precomputed ahead of time.

IV. THE PREPROCESSING PHASE

We present here the various image treatments carried out before initiating the template deformation processes. In view of the complexity of the segmentation task, we believe that an adequate preprocessing phase determines to a large extent the quality of the final match, improves the convergence speed, and strongly influences the ability of the method to avoid and/or escape local minima.

A. Nonuniformity Correction

The removal of intensity nonuniformity proves to be an essential prerequisite for the quantitative analysis of MRI data. As it is, spatial variations of the image signal unrelated to anatomic information often significantly reduce the accuracy of computational procedures such as tissue classification or brain-surface extraction. A number of techniques have been discussed in the literature [32]–[34]. Following the results of the comparative study reported by Arnold *et al.* [35], we chose the $n3$ method [32] to treat the input MRI data.




Fig. 4. Adaptive potential field. 0: MRI and initial template; 1: initial potential; 2, 3, 4, 5, and 6: five deformation steps; and 7: MRI and final deformed template.




Fig. 5. (a) Two-input MRIs with the associated intensity profile (along the red line). (b) Texture-filtered version with sharpened profile. Shaded boxes correspond to the target structure.

B. Texture Filtering

Even though certain brain structures, such as the *corpus callosum*, are relatively smooth and well contrasted in brain MRI, most of them present fuzzy boundaries or line-type edges like valleys or ridges [Fig. 5(a)]. Furthermore, many anatomical structures are surrounded by decoy elements with equally sharp (or dull) contrast: in Fig. 5(a), the *fornix* is not part of the *corpus callosum*, yet it presents very similar intensity characteristics. Also, in spite of the nonuniformity correction step, all or part of the data may be corrupted by noise. Thus,

the use of only intensity information along the contour of the structure is not sufficient to drive an accurate segmentation. Mudigonda *et al.* [36] uses the distribution of the intensity along the normals to the template surface. We propose the use of texture to prepare for the segmentation of the input MRI by enhancing the contrast of the target structure.

Texture Descriptors: A number of texture descriptors have been discussed in the literature (see [37] for a review) making use of statistical measures, time-space-frequency decompositions, or fractal parameters for instance. As every descriptor has weaknesses and strengths, one usually combines them to improve on the classification performance. Following the review of [37] and comments of [38], we selected the following descriptors.

- *Cooccurrence Matrices* [39]: These embed in a single array the relative frequencies of gray level pairs of pixels for a series of translations. We used eight translations and five of the texture features introduced by Haralick (energy, contrast, entropy, homogeneity, and correlation).
- *Discrete Cosine Transform* [40]: We used the low frequency component of a 3×3 DCT without critical sampling (eight features in total).
- *Dyadic Gabor Filters* [41]: We chose three different frequency bands and four orientations, and derived the energy and entropy for each of the $3 \times 4 = 12$ filtered windows.
- *Local Hölder Exponent* [38]: This corresponds to an intuitive perception of regularity: it is computed by comparing the total variation of the gray levels in the subwindow with functions of the form: $\varepsilon \rightarrow \varepsilon^\alpha$.

Classification: Given a target brain structure, we want to combine the $N = 73$ texture descriptors ($8 \times 5 + 8 + 2 \times 12 + 1$) within a classification function in an optimal fashion. That is, we are looking for a function $F: R^N \rightarrow R$ which is maximal for the target pixels and minimal elsewhere. For each anatomical structure we compile a series of MR images together with their

associated manually segmented structures. For each image, we select at random a number of subwindows (8×8 pixels) belonging either to the target or not, to form a training database. Each subwindow is “normalized” by subtracting the intensity of the central pixel from the other ones. Then, we use linear discriminant analysis [42] to find the coefficients w_i of the linear discriminant function: $F = \sum w_i.TD_i$ where $\{TD_i\}_i$ are the outputs of the texture descriptors.

To process an input MRI, we consider a surrounding subwindow around each pixel and normalize the intensities before applying the various texture filters and computing the value of the discriminant function. Fig. 5(b) shows two texture-classified images: the target structures are the *corpus callosum* [Fig. 5 (top)] and the *hippocampus* [Fig. 5 (bottom)]. It appears that the combination of texture filters effectively removed the *fornix* and cleared the neighborhood of the *hippocampus* while homogenizing both structures.

Note that our approach is rather scanner and protocol dependent: the texture characteristics of the target structures may vary from one acquisition to the next if the scanner parameters change (as pointed out in [43]). In that case, a few test images have to be manually delineated to retrain the classifier. We are currently investigating the extent of that dependence in the hope of overcoming it by using nonlinear classifiers to better capture scanner/protocol effects.

C. Initial Potential Field

The texture filter successfully manages to sharpen the transition boundaries between the target structure and the surrounding tissues. Therefore, we can then assume a step edge model and, thus, choose the Canny–Deriche approach to compute the gradient magnitude of the input image, as it achieves a good approximation of the optimal product of signal-to-noise ratio and localization [44].

As mentioned in Section III-D, we still have to choose the scale of the Gaussian smoothing. We use a two-pass scanline algorithm (one pass/axis) to set the scale to the mean distance μ between any two nonadjacent local gradient maxima. We obtain a potential field with the notable property that two edges at distance μ from each other induce a potential value close to zero in their middle.

V. STATISTICAL SHAPE MODEL AND HYBRID EVOLUTIONARY HEURISTIC

Even though a given brain structure can present a wide variety of forms, it seems that the notion of biological shape is reasonably well explained by a statistical description over a large population. Consequently, model-based approaches have attracted considerable attention [10], [14], [45], [46] as they can achieve robust segmentation. A deformable template is then not only constrained 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.

Once we have reduced the segmentation problem to a minimization task, however, we face a multimodal, nonlinear, and possibly discontinuous function of many variables. Given the size of the solution space, most minimization techniques would

only lead to weak suboptimal solutions if the search space were not drastically reduced by assuming that a good approximation of the solution is available. This could be either in the form of a set of pose parameters (position, orientation, scale) or of shape descriptors (possibly those of the mean shape model). Nevertheless, given that medical images are often tessellated (that is, similar pictorial elements are present repeatedly in the image, at different positions scales or orientation), the energy functional landscape is highly nonconvex and the minimization algorithm highly likely to run into local minima.

In contrast with approaches which restrict the template deformation space by implementing a coarse to fine strategy, we chose an evolutionary heuristic owing to its ability to solve hard optimization problems in a timely fashion where little is known about the structure of the search space.

A. Building the Shape Model

Aligning the Training Set: Given a learning set of parameterized instances of the target structure, we wish to compute a mean curve and the associated variation modes. To obtain an average curve which actually represents the mean shape of the target structure, we need to apply a transformation t to each curve in the learning set to bring them in a comparable configuration with respect to a set of axes.

Specifically, if \mathcal{T} is the group of allowed transformations and $L = \{S_1, \dots, S_N\}$ the learning set of curves with $\{S_i\}_i$ vectors of R^k , then the mean shape \hat{S} is given by

$$\hat{S} = \inf_{S \in R^k} \inf_{t_1, \dots, t_N \in \mathcal{T}} \sum_{i=1}^N \|S - t_i S_i\|^2.$$

In [47], Pennec shows that \hat{S} corresponds to the arithmetical mean of the $\{S_i\}_i$ once they have been transformed to an optimal position. That is, once we have found $t_1, \dots, t_N \in \mathcal{T}$ such that

$$\sum_{1 \leq i < j \leq N} \|t_i S_i - t_j S_j\|_2^2 = \inf_{t_1, \dots, t_N \in \mathcal{T}} \sum_{1 \leq i < j \leq N} \|t_i S_i - t_j S_j\|_2^2$$

then $\hat{S} = \frac{1}{N} \sum_i t_i S_i$.

We restrict \mathcal{T} to the group of similarities (rotations, translations, changes of scale) and used the technique described in [48] to find the corresponding set of transformations $\{t_1, \dots, t_N\}$. This consists of an iterative procedure where the various instances at iterations t are registered to their mean at iteration $t - 1$. The process starts by selecting one of those instances as a target for the first round of registrations and at the end of each such round. We eventually obtain a mean structure with the desired property that there does not exist any transformation that could further decrease the distances between the learning instances and that mean. Note that such an approach is only an attempt to establish a good correspondence field across the learning set and does not guarantee that a point associated to a given index is homologous to the points with same indexes in the other instances. We are currently investigating other techniques, using partial differential equations, that could establish better correspondence fields.

Capturing the Statistics: We use the approach described in [48] and [49]. Given the set of N aligned instances $\{t_i.S_i\}_i$ and the mean shape \hat{S} , we compute the centered covariance matrix K . The eigenvectors of K , obtained by principal component analysis, describe the modes of variation and the vectors corresponding to the largest eigenvalues describe the most significant ones. Any shape $t_i.S_i$ in the transformed learning set can then be approximated

$$t_i.S_i = \hat{S} + Q.a^i$$

where $Q = (q_1|q_2|\dots|q_m)$ is the matrix of the first m eigenvectors (we take $m \approx N/10$) and $a^i = (a_1^i, \dots, a_m^i)^t \in R^m$ is the shape vector.

By varying a^i , we can generate new instances of the shape model.

B. Hybrid Evolutionary Algorithm

The shape model we have just described is used within an evolutionary heuristic to constrain the dynamics of a population of templates.

Our evolutionary approach is similar to the guided evolutionary simulated annealing (GESA) method detailed in [50]. We have, however, modified it to incorporate our local deformation scheme from Section III, thus turning it into a hybrid evolutionary heuristic which explores the solution space by interspersing periods of local optimization with phases of interaction and diversification.

Initialization: We consider a population of individuals $\{I^i\}_i$, each of which consists of a pose and shape parameter: $I^i = (s^i, \theta^i, T^i, a_1^i, \dots, a_m^i)$ with s^i scale, θ^i rotation and T^i translation.

The population is made up of f families. Each family F^i consists of a parent P^i and c_i children $C^{i,j}$. The pose parameter of each of the parents is selected at random to cover the entire input image. Their shape is that of the mean surface. Note that when an anatomical atlas (a collection of delineated structures) is available, we can alternatively register the atlas to the input image and use the position of the deformed atlas structures to initialize the parents closer to the target structure. The children in each family are generated by randomly modifying the pose parameters of the parents and choosing valid instances of the shape model

$$F^i: P^i = (s^i, \theta^i, T^i, 0, \dots, 0) \\ C^{i,j} = (\delta_s^{i,j}.s^i, \delta_\theta^{i,j}.\theta^i, \delta_T^{i,j}.T^i, \varepsilon_1^{i,j}.\sqrt{\lambda_1}, \dots, \varepsilon_m^{i,j}.\sqrt{\lambda_m})$$

with $\delta_s^{i,j}, \delta_\theta^{i,j}, \delta_T^{i,j} \in R$, and $\varepsilon_1^{i,j}, \dots, \varepsilon_m^{i,j} \in R$ drawn at random from a uniform distribution whose range decreases with each generation; and $(\lambda_1, \dots, \lambda_m)$ are the ordered eigenvalues of the covariance matrix K .

The hybrid evolutionary algorithm is given in Fig. 6.

1) Projection onto the Shape Space: Let $d\varphi$ be the deformation induced by the local template matching scheme. We want to adjust the pose and shape parameters of the template to the required deformation while insuring that the deformed template be a valid instance of the shape model.

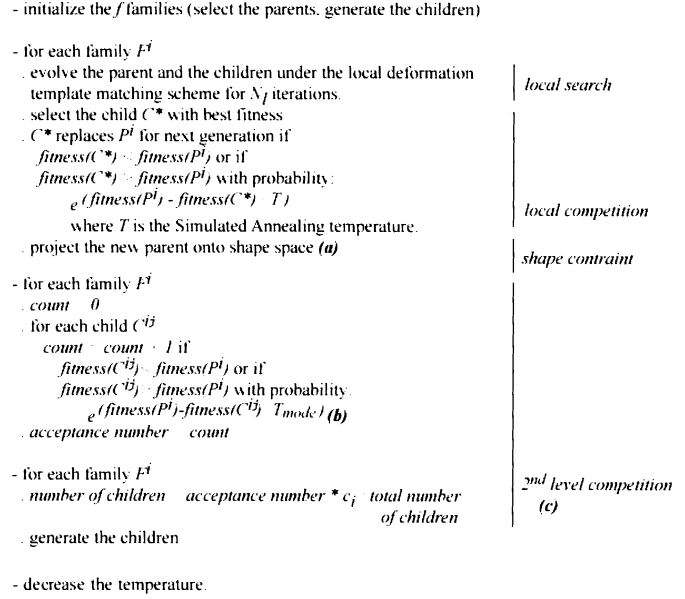


Fig. 6. Hybrid evolutionary algorithm.

Following the analysis of [22] and [48], we first register the template to the shape model. We then restrict the projection of $d\varphi$ onto the shape space so that the distance between the deformed template and the corresponding instance of the shape model is lower than three times the empirical standard deviation for each mode. We get

$$d\varphi_{\text{projected}} = \sum_{i=1}^m \text{sign} \left(\left\langle \varphi + d\varphi - \hat{S}, q_i \right\rangle \right) \cdot \min \left(\left| \left\langle \varphi + d\varphi - \hat{S}, q_i \right\rangle \right|, 3\sqrt{\lambda_i} \right) \cdot q_i.$$

2) To favor the children whose shapes are closer to the first eigenmodes, we vary the temperature of the Boltzman distribution that controls the acceptance number of each family as follows:

$$T_{\text{mode}} = \rho[\text{mode}] * T$$

where T is the SA temperature, ρ is an array of decreasing weights, and $\text{mode} = \arg \max_i |\langle d\varphi, q_i \rangle|$

3) As in [50], the approach implements two levels of competition. At a local level, the children of the same family compete with each other to generate the parent for the next generation. At a second level, there is a competition between the families themselves in that the number of children allocated to each family depends on the combined fitness of all the children, and is biased toward the first eigenmodes to favor the most likely shapes.

The number of children actually reflects the interest in a given area of the search space: the more good candidate solutions in a given area, the more attention we devote to it. The entire algorithm can be viewed as parallel SA with competition: each family is a multiple-trial-parallel SA machine with the children contributing the trials in parallel: it is, therefore, a very efficient global search technique as argued in [50].




Fig. 7. Shape space versus local deformation matching scheme. The black mesh represents valid instances of the shape model.

Note that by projecting the new parents into the shape space, we bias the search for a solution toward the set of valid instances of the shape model. Yet, incorporating an independent local search allows a better traversal of the search space. As illustrated in Fig. 7, the considered individual could not leave its current location to reach the target minimum if it were restricted to move only in the lower dimensional shape space (as it would have to go “up-hill”) whereas the local deformable matching scheme can find a path (in white) in the solution space.

VI. RESULTS

In this section, we present and discuss quantitative results showing the performance of our approach and the solutions it provides to the issues we identified in Section I.

A. Initial Position and Capture Range

Figs. 2 and 4 demonstrate the ability of the local matching scheme to evolve a template when initially placed in the neighborhood of the target structure. By texture-filtering the input MRI, we considerably enhance the contrast of the target structure with respect to the surrounding tissues and help decrease the influence of the nearby anatomical structures, thereby allowing the templates to be initialized further away.

Additionally, as our evolutionary heuristic can efficiently explore highly nonconvex spaces and discover the many attraction basins in the MRI potential, it further decreases the issue of initial placement. Fig. 8 compares the dynamics of a population of templates (four families; five children/family, initially) when controlled by our heuristic [Fig. 8(b)] with that of a single template subjected to the only local search method [Fig. 8(a)]. The shape model was enabled but we did not texture-filter the image. The target structure was the *caudate nucleus*. As expected, the single template quickly reaches a local minimum, whereas the population eventually manages to segment the *caudate nucleus* (both the left and right *caudate* could be segmented here as the shapes of 2-D coronal slices of left and right caudate nuclei are very similar: this would probably not occur in three dimensions). Note how the families that are too far from the target structure quickly disappear as their children cannot achieve a good match.

B. Tradeoff Flexibility/Robustness to Noise

Flexibility: Xu and Prince describe in [31] two key difficulties, related to the target structure geometry and its capture range, that arise with parametric active contour algorithms: first, how close to the target structure the initial template must be for the deformation process to be likely to evolve it toward the right boundaries?; second the difficulties the parametric model encounters while progressing into boundary concavities.

The combination of an adaptive potential field and a dynamic number of control points enables a very flexible adaptation to the target shape. Fig. 2 shows how the template can stretch to achieve a good match. Fig. 9(a)–(c) shows three sequences depicting the iterative progression of three prototype templates when facing concave boundaries. The temporally adaptive potential field and the dynamic number of control points enable our approach to achieve good results [Fig. 9(a) and (b)], where standard methods, or a distance potential field [29], or pressure forces [51] may fail. Note that in Fig. 9(b), pressure methods would perform well but for the two openings at the top and bottom, which cause the template to bulge outside of the “natural” contours. However, our model fails to achieve a good match when the concavity is too high [Fig. 9(c)]. We are currently studying the use of a progressive gradient vector flow [31], as well as covariant deformable template approaches which bias the template dynamics in concave regions using covariant diffusion operators and an auxiliary tensor field [52].

Sensitivity to Noise: A series of experiments was conducted to evaluate the influence of noise on the ability of the texture classifier to sufficiently contrast the target structure to allow a correct adaptation of the template. We selected a series of five 2-D images from the same brain MRI and manually segmented the *hippocampus* to train the texture filter. We chose yet another image to conduct the test: we corrupted it by a zero-mean Gaussian noise with variable variance σ^2 and intensity λ ; we used the only local search, with no shape model guidance; the template was manually initialized on top of the target. Fig. 10(a) displays the root-mean-square error (RMSE) (using closest point distance) between the retrieved boundary and the manually segmented *hippocampus*. Note how the final match seems relatively independent from the noise as long as its level stays below ($\sigma^2 = 5.0, \lambda = 1.2$). We repeated the same experiment with various noise distributions and structure instances, and observed similar behaviors.

We then re-evaluated the performance of our texture classifier when the training images were also corrupted by a Gaussian noise ($\sigma^2 = 5.0, \lambda = 1.2$). The same test image was selected. Fig. 10(b) demonstrates how the linear discriminant analysis was capable of selecting a better combination of the texture filters to deal with the noise-corrupted MRI. To a certain extent, the texture filter can “learn” the noise characteristics of the training images and, as a result, performs optimally on a test image with similar characteristics. Here again, further experiments would be necessary to assess how performance depends on the acquisition protocols and scanners.




Fig. 8. (a) Dynamics of a single template subjected to the local search. (b) Dynamics of a population of templates under the supervision of the evolutionary heuristic.




Fig. 9. Convergence of our deformable model for three target objects with highly concave boundaries.

C. Accuracy

Methodology: The accuracy with which our approach detects the contours of a target structure was determined following the method presented in [53]. For each experiment, we selected a number of 2-D MR images and manually segmented the target structure. We then let the templates evolve. The continuous contour given by the final B -spline template was eventually converted into a number N of equidistant points where N was at least twice the number of points in the corresponding manually segmented anatomical structure. As in [54], we used the four following error measures:

- unsigned mean (L_1 norm): $\sum_i |d_i|/N$;
- signed mean: $\sum_i d_i/N$;
- RMSE (L_2 norm): $\sqrt{(\sum_i d_i^2)/N}$;
- maximum error (L_∞ norm): $\max_i |d_i|$.

with d_i the signed Euclidean distance between the i th point on the detected border and the closest point on the manual contour. As pointed out in [54], the signed mean is an indicator of consistent bias. All errors are reported as mean \pm standard deviation.

Robustness to Decoys: Fig. 5(b) shows how the texture filter efficiently isolates the *fornix* from the *corpus callosum*. A new series of experiments was carried out to assess the influence of

the model guided heuristic on the final match in presence of a tessellated MRI. 94 manually segmented *hippocampus* sections from four different brains with their associated MR images were selected. Table I reports the various error measures. We used four families with five children each, spread all over the image. The texture filter was disabled to maximize the tessellation. To assess the sensitivity of the performance to the design of the shape learning set, we conducted the experiment with a variable number of training samples (the training set was disjoint from the test set). Comparison between rows one and two clearly demonstrates the improvement induced by the shape model. Note how the number of training samples influences the performance: when too few samples are used, the shape model does not reliably encompass the possible modes of variation; conversely, once a sufficient number of samples has been used, adding more samples only marginally increases the performance.

A Few Brain Structures: The following three brain structures were selected to test the performance of our segmentation approach.

- *Corpus Callosum*: The test set consisted of 80 sections from ten different brains (T_1 -weighted MRI, 1-mm resolution). An additional five images were used to train the texture filter, and 30 to train the shape model.
- *Caudate Nucleus*: 70 sections from five different brains were selected (T_1 -weighted MRI, 1-mm resolution). We used four images to train the texture filter and 25 for the shape model.
- *Hippocampus*: We used 94 sections from four different brains (T_1 -weighted MRI, 1-mm resolution). Five images were used to train the texture filter and 40 to train the shape model.

For each structure, we initialized four families of five children all over the image and ran the method for 500 iterations.

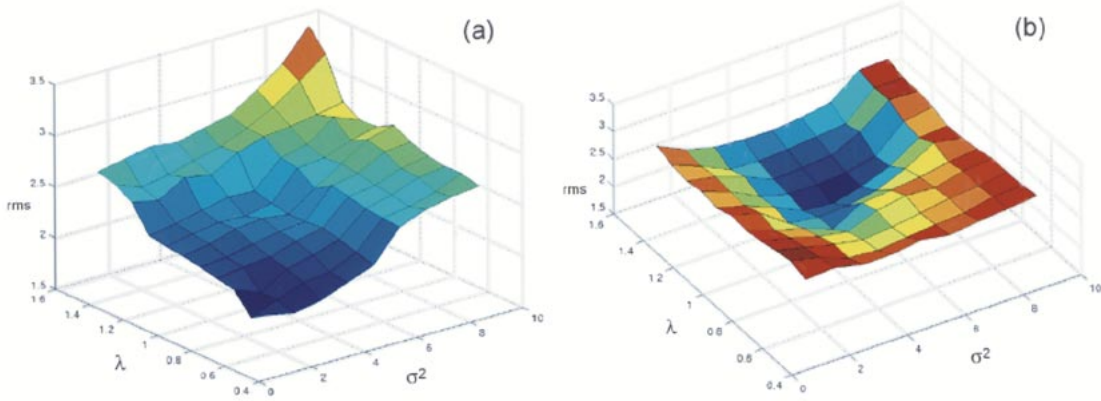


Fig. 10. RMSE of the segmentation of the *hippocampus* with (a) original training MRIs and (b) noisy training MRIs.

TABLE I
SEGMENTATION ERRORS FOR THE *HIPPOCAMPUS* (WITHOUT TEXTURE FILTER)

number of training samples	mean	signed mean	rms	max
0	3.9 ± 1.5	0.2 ± 0.4	3.5 ± 1.0	6.5 ± 1.9
15	2.5 ± 0.9	0.3 ± 0.3	2.7 ± 0.7	4.5 ± 2.1
35	1.2 ± 0.8	0.2 ± 0.5	1.5 ± 0.7	3.2 ± 1.8
55	1.3 ± 0.6	0.1 ± 0.7	1.4 ± 0.8	3.0 ± 1.9

TABLE II
SEGMENTATION ERROR FOR VARIOUS ANATOMICAL STRUCTURES

structure	mean	signed mean	rms	max	rms operator	max operator
<i>corpus callosum</i>	0.7 ± 0.5	0.1 ± 0.3	1.0 ± 0.4	1.1 ± 0.3	1.1 ± 0.3	1.9 ± 0.7
<i>caudate nucleus</i>	0.8 ± 0.4	-0.2 ± 0.2	1.1 ± 0.3	0.9 ± 0.4	0.9 ± 0.4	1.5 ± 1.0
<i>hippocampus</i>	1.1 ± 0.7	0.1 ± 0.5	1.3 ± 0.6	2.7 ± 1.0	1.4 ± 0.3	2.8 ± 1.7

Table II reports the measured errors between the manually segmented structures and the outcomes of our algorithm. It also gives the RMSE and maximum error of the manual segmentation: as it is, a human operator may not always segment a structure in the same consistent manner from one MR image to another, and the amount of variability may change with respect to the target structure.

As expected, our approach had no difficulties segmenting the *corpus callosum* given its regular shape and clear contrast. The two other segmentations were also very good. Note that for the *caudate nucleus*, the method retrieved both *nuclei* (left and right) in every trial.

For all three structures, the automated segmentation compares favorably to the manual one in terms of rms error, even though the maximum error was occasionally substantially larger: the influence of the shape model was sometimes not strong enough to overcome the mistakes our local template matching algorithm made when a strong decoy edge prevented a correct deformation: the use of a better matching functional might alleviate that problem. Note that the delineation inconsistencies of human operators also affect the quality of the learning sets as both the shape descriptors and the texture classifier are dependent on the accuracy of the manually extracted structures.

We are currently investigating statistical measures that could better take into account these “intrinsic” errors.

Running Time: We give a few computational time measurements. On a SGI O2 graphics workstation (MIPS R4400 200 MHz), texture-filtering a 300×300 pixel image takes approximately 40 s. Each of the above segmentation runs (4×5 children initially) took about 6 min/image. Note that our implementations were not optimized for speed: the values we give here are simply *indications* of the running time.

D. Extension to Three Dimensions

We discuss here how our approach could be extended to three dimensions.

Elastic Deformation Model: The B -spline curves we use in the 2-D case can readily be replaced by 3-D tensor product B -spline surfaces. Accordingly, the deformation template becomes

$$\theta = (\theta_x, \theta_y, \theta_z): [0, 1]^2 \rightarrow R^3, (u, v) \rightarrow \sum_{i=1}^I \sum_{j=1}^J P_{\theta}^{j,i} \cdot B_i(u) \cdot B_j(v)$$

where $B_i(u)$, $B_j(v)$ are the B -spline blending functions and $P_{\theta}^{j,i}$ are the control points.

Following the analysis of [55], the objective function is written

$$\mathcal{F}(\varphi) = \int_0^1 \int_0^1 \Xi(I_E)(\varphi) + \int_0^1 \int_0^1 \|S(\varphi)\|_\alpha^2 + \|B(\varphi)\|_\beta^2$$

where S and B represent the first and second fundamental surface forms; and $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ are weighted matrix norms.

Under the thin-plate model [24], we simplify

$$\int_0^1 \int_0^1 \|S(\varphi)\|_\alpha^2 + \|B(\varphi)\|_\beta^2 = \int_0^1 \int_0^1 \sum_{i,j=1}^2 \alpha_{ij} \cdot \partial_i \varphi \cdot \partial_j \varphi + \beta_{ij} \cdot (\partial_i \partial_j \varphi)^2$$

where $\partial_i \varphi$ represents the partial derivative with respect to the i th parameter.

Strictly speaking, the approximation is accurate only near the actual minima, but it is nonetheless still well-behaved away from it. For a tensor product B -spline, we get a quadratic objective function that could still be locally minimized with our adaptive gradient descent algorithm.

Dynamic variation of the number of control points, however, is not easily extendable to the 3-D case, as the adaptive refinement of a 3-D tensor product B -spline surface is not a trivial issue.

Preprocessing Phase: Our texture filter could also be extended to three dimensions as most of the texture descriptors we chose have a 3-D formulation. Great care should nevertheless be taken in the implementation of the formulae or else processing an entire 3-D MRI may require too much time.

Working out a 3-D potential field should not cause any problem either as the Canny extractor also has a 3-D counterpart. However, given the often anisotropic nature of MRI data, one might favor the anisotropic detector of Brejl *et al.* [56] for it would then outperform the Canny detector.

Shape Model and Evolutionary Heuristic: Finally, the methodology we developed in Section V still holds in three dimensions. Note that the ability of our hybrid evolutionary algorithm to foster attraction basins in complex spaces would be especially interesting in a highly nonconvex 3-D extension.

VII. CONCLUSION

We have presented a general fully automated segmentation method, capable of dealing with a variety of anatomical structures. The use of a hybrid scheme combining a *global* evolutionary heuristic and a *local* deformable template matching algorithm proved very promising: while the evolutionary heuristic can effectively explore the solution space, the local matching scheme is capable of thoroughly exploiting each attraction basin the heuristic finds. Furthermore, the hybrid nature of our approach makes it easy to integrate a statistical shape model. As it is, we could easily replace the statistical term by another, or add other context-dependent constraints without having to modify the deformable template scheme.

We used such a framework to develop and study several optimization schemes and demonstrate the multivariate relations between the parameters of our deformation model, and how they affect the accuracy of the match. We believe that adaptive approaches which modify the control parameters of de-

formable templates in response to contextual information and/or using time-dependent criteria may present advantages in the progressive refinement and understanding of template matching methods.

REFERENCES

- [1] G. Subsol, "A scheme for automatically building 3D morphometric anatomical atlases based on feature lines: A list of references," in *DIKU International Summer School '98—Shape Variation*. Copenhagen, Denmark: Univ. Copenhagen, 1998.
- [2] S. Haney, P. M. Thompson, T. F. Cloughesy, J. R. Alger, and A. W. Toga, "Tracking tumor growth rates in patients with malignant gliomas: A test of two algorithms," *Amer. J. Neuroradiol.*, vol. 22, no. 1, 2001.
- [3] A. Del Bimbo and P. Pala, "Visual image retrieval by elastic matching of user sketches," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 19, pp. 121–132, Feb. 1997.
- [4] P. M. Thompson, R. P. Woods, M. S. Mega, and A. W. Toga, "Mathematical/computational challenges in creating deformable and probabilistic atlases of the human brain (invited paper)," *Human Brain Mapping*, vol. 9, no. 2, pp. 81–92, 2000.
- [5] D. J. Burr, "Elastic matching of line drawings," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-3, pp. 708–713, Nov. 1981.
- [6] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *IJCV*, vol. 1, no. 4, pp. 321–331, 1988.
- [7] G. Sapiro, "Color snakes," *Comput. Vis. Image Understanding*, vol. 68, no. 2, pp. 247–253, 1997.
- [8] S. Osher and J. A. Sethian, "Fronts propagating with curvature dependent speed: Algorithms based on Hamilton–Jacobi formulations," *J. Computational Phys.*, vol. 79, pp. 12–49, 1988.
- [9] M. Leventon, E. Grimson, and O. Faugeras, "Statistical shape influence in geodesic active contours," in *Proc. CVPR '00*, 2000, pp. 1:316–323.
- [10] U. Grenander and M. I. Miller, "Representations of knowledge in complex systems," *J. Roy. Statist. Soc. B*, vol. 56, no. 3, pp. 549–603, 1994.
- [11] T. F. Cootes, C. J. Taylor, and A. Lanitis, "Active shape models: Evaluation of a multi-resolution method for improving image search," in *Proc. Br. Machine Vision Conf.*, vol. 1, 1994, pp. 327–336.
- [12] A. Jain, Y. Zhong, and S. Lakshmanan, "Object matching using deformable templates," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 18, pp. 267–277, Mar. 1996.
- [13] J. Gee and L. Le Briquer, "An empirical model of brain shape," in *Maximum Entropy and Bayesian Methods*, G. Erickson and M. Knoll, Eds. Dordrecht, The Netherlands: Kluwer Academic, 1998, pp. 199–207.
- [14] L. H. Staib and J. S. Duncan, "Boundary finding with parametrically deformable models," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, pp. 1061–1075, Nov. 1992.
- [15] G. Székely, A. Kelemen, C. Brechbühler, and G. Gerig, "Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models," *Med. Image Anal.*, vol. 1, pp. 19–34, 1996.
- [16] Z. Michalewicz, *Genetic Algorithms + Data Structures = Evolution Programs*. Berlin, Germany: Springer-Verlag, 1999.
- [17] D. Costa, "An evolutionary tabu search algorithm and the NHL scheduling problem," *INFOR*, vol. 33, no. 3, pp. 161–178, 1995.
- [18] H. Muehlenbein, M. Gorges-Schleuter, and O. Kraemer, "Evolution algorithms in combinatorial optimization," *Parallel Computing*, vol. 7, pp. 65–88, 1998.
- [19] B. Freisleben and P. Merz, "New genetic local search operators for the traveling salesman problem," in *Lectures Notes in Computer Science 1141*. Berlin, Germany: Springer-Verlag, 1996, pp. 890–899.
- [20] E. Falkenauer, "A hybrid grouping genetic algorithm for bin packing," *J. Heuristics*, vol. 2, no. 1, pp. 5–30, 1996.
- [21] C. Fleurent and J. A. Ferland, "Object-oriented implementation of heuristic search method for graph coloring, maximum clique, and satisfiability," in *Proc. 2nd DIMACS Implementation Challenge*, vol. 26, 1996, pp. 619–652.
- [22] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam, "The use of active shape models for locating structures in medical images," *Image Vis. Computing*, vol. 12, no. 6, pp. 355–366, 1994.
- [23] R. C. Veltkamp and M. Hagedoorn, "State-of-the-art in shape matching," in *Principles of Visual Information Retrieval*, M. S. Lew, Ed. Berlin, Germany: Springer-Verlag, 2001, Ser.: Advances in Pattern Recognition.
- [24] D. Terzopoulos, "Regularization of inverse visual problems involving discontinuities," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, pp. 413–424, 1986.

- [25] Y. Zhong, "Object matching using deformable template," Ph.D. dissertation, Michigan State Univ., East Lansing, 1996.
- [26] A. Blake and M. Isard, *Active Contours*. Berlin, Germany: Springer-Verlag, 1998.
- [27] J. Montagnat, H. Delingette, N. Scapel, and N. Ayache, "Representation, shape, topology and evolution of deformable surfaces. Application to 3D medical image segmentation," INRIA, Res. Rep. 3954, 2000.
- [28] T. P. Vogl, J. K. Mangis, A. K. Rigler, W. T. Zink, and D. L. Alkon, "Accelerating the convergence of the back-propagation method," *Biological Cybern.*, vol. 59, pp. 257–263, 1988.
- [29] L. D. Cohen and I. Cohen, "Finite-element methods for active contour models and balloons for 2-D and 3-D images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 15, pp. 1131–1147, Nov. 1993.
- [30] C. Davatzikos and J. L. Prince, "Convexity analysis of active contour problems," *Image Visual Computing J.*, vol. 17, no. 1, pp. 27–36, 1999.
- [31] C. Xu and J. L. Prince, "Snakes, shapes and gradient vector flow," *IEEE Trans. Image Processing*, pp. 359–369, Mar. 1998.
- [32] J. G. Sled, A. P. Zijdenbos, and A. C. Evans, "A nonparametric method for automatic correction of intensity nonuniformity in MRI data," *IEEE Trans. Med. Imag.*, vol. 17, pp. 87–97, Feb. 1998.
- [33] B. H. Brinkmann, A. Manduca, and R. A. Robb, "Optimized homomorphic unsharp masking for MR grayscale inhomogeneity correction," *IEEE Trans. Med. Imag.*, vol. 17, pp. 161–171, Apr. 1998.
- [34] M. S. Cohen, R. M. Dubois, and M. M. Zeineh, "Rapid and effective correction of RF inhomogeneity for high field magnetic resonance imaging," *Human Brain Mapping*, vol. 10, pp. 204–211, 2000.
- [35] J. B. Arnold, J.-S. Liow, K. A. Schaper, J. J. Stern, J. G. Sled, D. W. Shattuck, A. J. Worth, M. S. Cohen, R. M. Leahy, J. C. Mazziotta, and D. A. Rottenberg, "Qualitative and quantitative evaluation of six algorithms for correcting intensity nonuniformity effect," *NeuroImage*, vol. 13, pp. 931–943, 2001.
- [36] N. R. Mudigonda, R. M. Rangayyan, and J. E. Leo Desautels, "Gradient and texture analysis for the classification of mammographic masses," *IEEE Trans. Med. Imag.*, vol. 19, pp. 1032–1043, Oct. 2000.
- [37] T. Randen and J. H. Husøy, "Filtering for texture classification: A comparative study," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 21, pp. 291–310, Apr. 1999.
- [38] J. Levi-Vehel and B. Guiheneuf, "Texture based video indexing," presented at the IASTED, 2000.
- [39] R. M. Haralick, "Statistical and structural approaches to texture," *Proc. IEEE*, vol. 67, pp. 786–804, May 1979.
- [40] I. Ng, T. Tan, and J. Kittler, "On local linear transform and Gabor filter representation of texture," in *Proc. Int Conf. Pattern Recognition*, 1992, pp. 627–631.
- [41] A. K. Jain and F. Farrokhnia, "Unsupervised texture segmentation using Gabor filters," *Pattern Recogn.*, vol. 24, no. 12, pp. 167–186, 1991.
- [42] A. M. Martinez and A. C. Kak, "PCA versus LDA," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 23, pp. 228–233, Feb. 2001.
- [43] EC COST B11 Action. Eur. Co-op Sci Technical Res. [Online]. Available: <http://www.uib.no/costb11>.
- [44] R. Deriche, "Using Canny's criteria to derive a recursively implemented optimal edge detector," *Int. J. Comput. Vis.*, vol. 1, no. 2, pp. 167–187, 1987.
- [45] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models—Their training and application," *Comput. Vis. Image Understanding*, vol. 61, no. 1, pp. 38–59, 1995.
- [46] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cogn. Neurosci.*, vol. 3, no. 1, pp. 71–86, 1991.
- [47] X. Pennec, "L'incertitude dans les problèmes de reconnaissance et de recalage—Applications en imagerie médicale et biologie moléculaire," Ph.D. dissertation, Ecole Polytechnique, Palaiseau, France, 1996.
- [48] N. Raynaud, "Approche Statistique pour la Segmentation d'Images Tridimensionnelles du foie," Ecole Normale Supérieure de Cachan (ENS), Cachan, France, Rep. DEA, 2000.
- [49] Y. Wang and L. H. Staib, "Boundary finding with correspondence using statistical shape models," in *Proc. CVPR '98*, pp. 338–345.
- [50] P. P. C. Yip and Y. H. Pao, "Combinatorial optimization with use of guided evolutionary simulated annealing," *IEEE Trans. Neural Networks*, vol. 6, pp. 290–295, Feb. 1995.
- [51] L. D. Cohen, "On active contour models and balloons," *CVGIP: Image Understanding*, vol. 53, no. 2, pp. 211–218, Mar. 1991.
- [52] P. M. Thompson and A. W. Toga, "Anatomically driven strategies for high-dimensional brain image warping and pathology detection," in *Brain Warping*, A. W. Toga, Ed. New York: Academic, 1998, pp. 311–336.
- [53] V. Chalana and Y. Kim, "A methodology for evaluation of boundary detection algorithms on medical images," *IEEE Trans. Med. Imag.*, vol. 16, pp. 642–652, Oct. 1997.
- [54] M. Brejl and M. Sonka, "Object localization and border detection criteria design in edge-based image segmentation: Automated learning from examples," *IEEE Trans. Med. Imag.*, vol. 19, pp. 973–985, Oct. 2000.
- [55] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, "Elastically deformable models," *Comput. Graphics*, vol. 21, no. 4, 1987.
- [56] M. Brejl and M. Sonka, "Directional 3-D edge detection in anisotropic data: Detector design and performance assessment," *Comput. Vis. Image Understanding (Special Issue on Analysis of Volumetric Images)*, vol. 77, pp. 84–110, 2000.