A TIGHTLY COUPLED REGION-SHAPE FRAMEWORK FOR 3D MEDICAL IMAGE SEGMENTATION

Rui Huang, Vladimir Pavlovic, and Dimitris N. Metaxas

Department of Computer Science, Rutgers University, Piscataway, NJ 08854

ABSTRACT

Most hybrid 3D segmentation methods either heuristically couple the respective algorithm or combine a true 3D with a 2D algorithm due to computational considerations. In this paper we propose a new probabilistic framework for 3D image segmentation that combines tightly linked region- and shape-based constraints. Region-based label constraints are modeled by a 3D Markov random field, and are tightly coupled to shape-based constraints of a 3D Deformable Model. The full 3D nature of the combined model leads to a robust smooth surface segmentation that outperforms the single constraint, slice-based as well as the loosely coupled 3D methods.

1. INTRODUCTION

Image segmentation is one of the most important medical image processing tasks. Recently, increasing availability of high-resolution 3D volume data using modalities such as Magnetic Resonance (MR) and Computed Tomography (CT) has prompted the need for true 3D segmentation approaches. However, 3D segmentation remains an extremely difficult problem, due to the complex topology of 3D objects, the massive data, and demanding computational algorithms.

Most 3D methods solve the segmentation problem using either region- or shape-based approaches. Unfortunately, region-based methods often result in rough and inaccurate object boundaries as well as disconnected object regions. Shape-based methods, on the other hand, explicitly model object boundaries but are sensitive to noise and may lead to oversmoothing. One possible answer to improved 3D segmentation relies on hybrid methods that combine region- and shape-based approaches. Many hybrid methods, however, assume weak coupling of the two modalities, failing to fully benefit from the region-shape dependencies. Moreover, many 3D approaches are often 2D in nature (i.e., slice-based), further discarding potential benefits of the 3D data.

In this paper, we propose a new framework for 3D image segmentation based on the integration of 3D Markov Random Fields (MRFs) and Deformable Models (DMs). The proposed method is a true 3D method unlike the traditional 2D slice-based approaches. The algorithm is semi-automatic requiring only a shape prior or two manually picked seed points (one inside the object, the other outside but close to the object boundary). As a consequence, our approach is able to fully exploit the structure of the 3D data, resulting in improved object segmentation. The coupling also leads to computationally more efficient solutions.

2. 3D SEGMENTATION USING MRFS

MRFs have long been used for image segmentation because of their ability to model spatial contextual information of an image and deal with the noise. The simplest way to extend a 2D MRF-based segmentation method to 3D is to apply the 2D algorithm slice by slice to the 3D volume data [1]. The lack of interaction among individual slice solutions, however, leads to results that are inferior to true 3D-based solutions [2].

A true 3D MRF model is shown in Fig. 1. Each voxel of the 3D volume data is an observable node o_i, which is connected to a hidden node s_i representing the region label of that voxel (Fig. 1a). Each hidden node is also connected to 6 (or more if needed) neighboring hidden nodes (Fig. 1b).

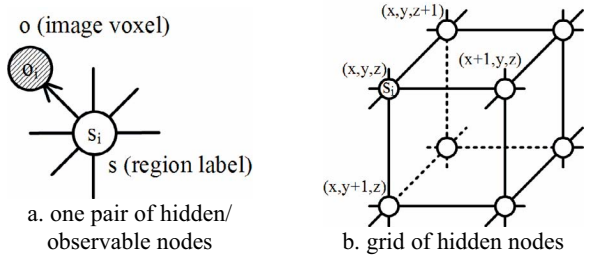

Fig. 1. 3D MRF models

The segmentation problem can be viewed as a problem of estimating the maximize-a-posterior (MAP) solution of the MRF model:

$$\mathbf{s}_{MAP} = \arg\max_{\mathbf{s}} P(\mathbf{s} \mid \mathbf{o}) \tag{1}$$

where

$$P(\mathbf{s} \mid \mathbf{o}) \propto P(\mathbf{o} \mid \mathbf{s})P(\mathbf{s}) \propto \prod_i \Phi_i(s_i) \prod_{(i,j)} \Psi_{ij}(s_i, s_j) \tag{2}$$

The compatibility function $\Phi_i(s_i)$ defines the relationship between a pair of hidden and observable nodes, and $\Psi_{ij}(s_i, s_j)$ defines the relationship between a pair of neighboring hidden nodes.

The EM algorithm is used to estimate both the MAP solution of the region labels and the model parameters (i.e., the parameters in the compatibility functions) [3]. In E-step, the MRF-MAP solution is estimated based on the current model parameters. Belief propagation (BP) is an efficient inference method for this problem [4]. In M-step, the MRF parameters are updated based on the MAP solution of the region labels.

The performance of the 3D MRF-based segmentation method, verified by our experiments, has advantages over slice-based methods. However, it still may suffer from the common problems of region-based methods.

3. COUPLING 3D MRFS AND DMS

Geometrically Deformed Models (GDMs) [5] and Finite-Element Method (FEM)-based Balloon models [2] are commonly used for the representation of 3D surfaces and segmentation of volume data. While a deformable surface is often sufficiently smooth, the oversmoothing may be excessive; and it often suffers from local energy minima in noisy images.

A hybrid 3D approach was proposed in [6], attempting to couple MRFs and DMs to alleviate the deficiencies of the individual methods. However, the MRF model is applied slice by slice due to the computational complexity, which may lead to inconsistencies between neighboring slices and fail to build the deformable surface. Moreover, the MRF and the DM are coupled in a loose manner (e.g., there is no direct guidance from the DM to the MRF). [4] proposed a method for tight coupling of the MRF and DM constraints in 2D segmentation, which tangible benefits. In the 3D case, however, the model expansion process far away from the true boundary can be time-consuming and needs frequent reparametrization, and may suffer from local energy minima in noisy images. We address these issues here and propose a true 3D hybrid segmentation framework.

3.1. The graphical model representation

Our new framework uses the graphical model theory to tightly couple 3D MRFs and DMs. This is achieved, as depicted in Fig. 2a, by constraining region labels **s** with the underlying 3D object surface **v**. To simplify the graph, only one pair of **s** and **o** nodes is drawn, which represents the whole grid in Fig. 1.

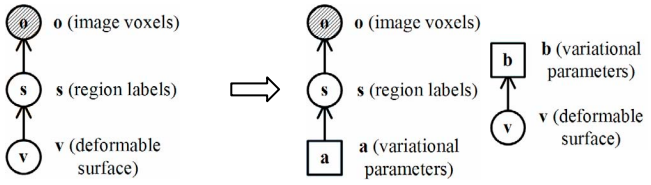

a. Integrated model b. Extended MRF c. Modified DM

Fig. 2. Integrated and decoupled models

DMs can often be viewed in a probabilistic framework, in which the internal energy $E_{\text{int}}(\mathbf{v})$ leads to a Gibbs prior distribution of the form:

$$P(\mathbf{v}) = \tfrac{1}{Z_i}\exp(-E_{\text{int}}(\mathbf{v})) \tag{3}$$

while the external energy $E_{\text{ext}}(\mathbf{v})$ from the data can be converted to a sensor model with conditional probability:

$$P(\mathbf{o} \mid \mathbf{v}) = \tfrac{1}{Z_e}\exp(-E_{\text{ext}}(\mathbf{v})) \tag{4}$$

where $E_{\text{ext}}(\mathbf{v})$ is some function of the image **o**.

Rather than directly conditioning the image **o** on the DM, we assume that the conditioning is indirect, through constraints on the image region labels:

$$P(s_i = inside \mid \mathbf{v}) = 1/(1+\exp(-dist(i,\mathbf{v}))) \tag{5}$$

$$P(s_i = outside \mid \mathbf{v}) = 1 - P(s_i = inside \mid \mathbf{v}) \tag{6}$$

induced by the signed distance of voxel *i* from surface **v**:

$$dist(i,\mathbf{v}) = sign(i,\mathbf{v}) \min_j \left\| loc(i) - \mathbf{v}_j \right\| \tag{7}$$

where $sign(i,\mathbf{v}) = 1$ if voxel *i* is inside surface **v**; $sign(i,\mathbf{v}) = -1$ when it is outside. Mathematically, this sign can be determined by checking the surface normals. $loc(i)$ denotes the spatial coordinates of voxel *i*, and $\mathbf{v}_j$ denotes the discretized spatial coordinates of a vertex of surface **v**.

This leads to a tightly coupled MRF-DM probabilistic model, $P(\mathbf{o},\mathbf{s},\mathbf{v}) = P(\mathbf{o} \mid \mathbf{s})P(\mathbf{s} \mid \mathbf{v})P(\mathbf{v})$. The segmentation can hence be viewed as a *joint* MAP estimation problem:

$$(\mathbf{v},\mathbf{s})_{MAP} = \arg\max_{(\mathbf{v},\mathbf{s})} P(\mathbf{v},\mathbf{s} \mid \mathbf{o}) \tag{8}$$

where

$$P(\mathbf{v},\mathbf{s} \mid \mathbf{o}) \propto P(\mathbf{o} \mid \mathbf{s})P(\mathbf{s} \mid \mathbf{v})P(\mathbf{v}) \tag{9}$$

3.2. Variational inference for hybrid MRF-DM

Despite the compact graphical representation of the integrated model, the exact inference in the model is computationally intractable. One reason is the large state space size and the complex dependency structure introduced by the Gibbs distribution of the prior $P(\mathbf{v})$. The second reason is the existence of loops in the graphical model precludes polynomial-time inference. To deal with these problems we use an approximate, yet tractable, inference solution.

Given an image **o**, structured variational inference finds a distribution $Q(\mathbf{v},\mathbf{s} \mid \mathbf{o},\theta)$, with an additional set of *variational parameters* θ, which is close to the original posterior:

$$\theta^* = \arg\min_{\theta} \sum_{\mathbf{v},\mathbf{s}} Q(\mathbf{v},\mathbf{s} \mid \mathbf{o},\theta) \log \frac{P(\mathbf{v},\mathbf{s} \mid \mathbf{o})}{Q(\mathbf{v},\mathbf{s} \mid \mathbf{o},\theta)} \tag{10}$$

In addition to being "close", the dependency structure of *Q must* allow a computationally efficient inference. In our case the *Q* factorizes the integrated model *P* into two independent distributions: an extended MRF model Q_M with variational parameters **a** and a modified DM Q_D with variational parameters **b** (Fig. 2b and 2c):

$$Q(\mathbf{v},\mathbf{s} \mid \mathbf{o},\mathbf{a},\mathbf{b}) = Q_M(\mathbf{s} \mid \mathbf{o},\mathbf{a})Q_D(\mathbf{v} \mid \mathbf{b}) \tag{11}$$

The *extended* MRF model means we have an additional part to the traditional MRF model to deal with the shape prior:

$$Q_M(\mathbf{s} \mid \mathbf{o},\mathbf{a}) \propto \prod_i [\Phi_i(s_i)P(s_i \mid a_i)] \prod_{(i,j)} \Psi_{ij}(s_i,s_j) \tag{12}$$

The modified DM uses the variational parameter **b** in place of the region labels:

$$Q_D(\mathbf{v} \mid \mathbf{b}) \propto \prod_i P(b_i \mid \mathbf{v}) Q_D(\mathbf{v}) \tag{13}$$

The optimal values of the variational parameters are obtained by minimizing the KL-divergence [7]:

$$\log P(s_i \mid a_i^*) = \sum_{\mathbf{v}} Q_D(\mathbf{v} \mid \mathbf{b}^*) \log P(s_i \mid \mathbf{v}) \tag{14}$$

$$\log P(b_i^* \mid \mathbf{v}) = \sum_{s_i \in L} Q_M(s_i \mid \mathbf{o},\mathbf{a}^*) \log P(s_i \mid \mathbf{v}) \tag{15}$$

Eq. (12) (13) allow us to apply efficient inference methods to the decoupled models (e.g., BP for MRF-MAP and FEM for DM fitting). Eq. (14) (15), on the other hand, ensure the new distribution is as close as possible to the original distribution. These four equations form a set of *fixed-point equations*. Solution of this fixed-point set yields a tractable approximation to the intractable original posterior.

Incorporating a DM in the MRF framework smoothes out the rough boundaries, and more importantly, introduces a shape prior into our model. The presence of the shape model allows us to perform the BP only in a band of model variables around the current surface estimates. The number of voxels in the band depends only

on the band thickness and the area of the object, instead of the whole image volume. In most of our experiments, this proportion is usually small (e.g., 20%). This significantly speeds up the whole algorithm. The hybrid MRF-DM algorithm is summarized:

1. if (shape prior **v** is not given) {
   1.1. Use the 3D MRF segmentation algorithm (Sec. 2) to generate an initial region segmentation;
   1.2. Use the Marching Cubes [8] algorithm to get an initial surface **v** from the region segmentation; }
2. while (error > maxError) {
   2.1. Calculate a band area **B** around **v**. Perform the remaining steps 2.2~2.5 inside **B**;
   2.2. Calculate $P(s_i \mid a_i)$ based on Eq. (14) using current surface model **v**;
   2.3. Estimate the MRF-MAP solution $Q_M(s_i \mid \mathbf{o}, \mathbf{a})$ based on Eq. (12) using $P(s_i \mid a_i)$, which is essentially the BP-based 3D MRF segmentation;
   2.4. Calculate $\log P(b_i \mid \mathbf{v})$ based on Eq. (15) using $Q_M(s_i \mid \mathbf{o}, \mathbf{a})$;
   2.5. Deform the deformable model **v** to fit the data based on Eq. (13) using the generalized force:
   $\partial \log P(\mathbf{b} \mid \mathbf{v}) / \partial \mathbf{v} = -\partial \log P(\mathbf{b} \mid \mathbf{v}) / \partial loc(i)$; }

4. EXPERIMENTS

Our algorithm is implemented in MATLAB with embedded C code, and tested on a 3GHz Pentium 4 PC. The running time varies from less than one minute to some dozen minutes depending on the data size and noise level.

4.1. Simulated images

The initial study of properties and utility of our method was conducted on a $64 \times 64 \times 64$ synthetic volume, which contains an "S"-like object of intensity 160 from frame 29 to 36, and the intensity of the background is 100. Gaussian noise with mean 0 and standard deviation 60 is added to the whole volume, and extra Gaussian noise with mean 160 is added to frames 32 and 33 (Fig. 3a). The 3D MRF model results in better segmentation than the 2D MRF, especially on frames 32 and 33, due to the interaction between neighboring frames. However, there is obvious error in the first and last frames because these two outermost frames suffer more interference from the background. This is improved by coupling the DM with the 3D MRF. Other frames are also smoother in Fig. 3d. The average error rates are shown in the figure.

We then tested our algorithms on simulated brain MRI data from BrainWeb [9]. There are two anatomical models: normal and multiple sclerosis. For both of these, full 3D data volumes have been simulated using three sequences (T1-, T2-, and proton-density- (PD-) weighted) and a variety of slice thicknesses, noise levels, and levels of intensity non-uniformity (RF inhomogeneity). Three different white matter segmentation results from normal brain data volumes using our hybrid method are shown. Fig. 4a is a slice from the ground truth data of the white matter, and 4b is the result from our hybrid method. Fig. 4c is a slice of the T1 image without noise and intensity non-uniformity. The segmented white matter (Fig. 4d) is slightly thicker than the results from the ground truth, because some of the grey matter is misclassified due to its similar grey value to the white matter. Same misclassification can be observed in Fig 4f, which is the segmentation result on the T1 image with 9% noise and 40% intensity non-uniformity (Fig. 4e). One possible solution to the misclassification problem is to initialize the 3D MRF with more accurate model parameters, which could be learned beforehand or specified by users.

We do not show the results of the slice-based method using 2D MRFs mainly because we cannot get satisfying results. One reason is that the white matter is a complicated object with high curvature, so there are more inconsistencies between the slice-based segmentation results. Another reason is that the traditional MRF model is sensitive to initialization, especially when there are multi-regions. The brain MRI data consists of many different regions (e.g., white matter, grey matter, skull, etc.), thus different slices may require different initializations. With our hybrid method, one only needs to consider the voxels close to the current boundary, i.e., we simplify the problem to only two regions (inside or outside). Even without a shape prior, the 3D MRF model is more robust than the 2D MRF to start with.

4.2. Real medical images

The real world images usually have significant, often non-white noise and contain complex high-curvature objects, rendering the segmentation task a great deal more difficult. In this section we show experimental results of applying our methods to a real 3D volume [10], which is an MR image of a head with the skull partially removed to reveal the brain. We again segmented the white matter from the volume. Fig. 5a is one of the slices from the volume. The result of 3D MRF only and that of 3D MRF + DM are shown in Fig. 5b and 5c. To show the benefits of DMs, the upper-right parts of 5b and 5c are magnified in 5d and 5e. The arrows point to some incorrect patches that are eliminated by the deformable model fitting process (the colors of these patches are clearly different to their neighboring patches, which obviously shows the spatial discontinuities). To avoid oversmoothing, the parameters are usually chosen according to experts' opinions.

5. DISCUSSIONS

We have proposed a new tightly coupled hybrid framework to segment 3D medical images. The efficiency of the BP algorithm makes it affordable to apply a true 3D MRF model. DMs further improve the performance and decrease the computational complexity. A graphical model-based approach yielded an approximate, computationally efficient solution to otherwise intractable inference of region boundaries. We showed the advantages and utility of our hybrid method on a number of synthetic and real-world images, supported by the quantitative evaluations as well as qualitative expert opinions.

We are currently working on more precise quantitative evaluation of the segmentation algorithms, which requires manual segmentations by the experts and an appropriate metric. BrainWeb provides excellent ground truth for region segmentations; however, since our method is a region-boundary hybrid method, for those voxels that split by the boundary surface, the accurate evaluation of sub-voxel segmentation is still a problem using the traditional voxel based ground truth. Moreover, the quantitative metric should also be able to measure the smoothness of the surface, in addition to the accuracy of voxel classification.

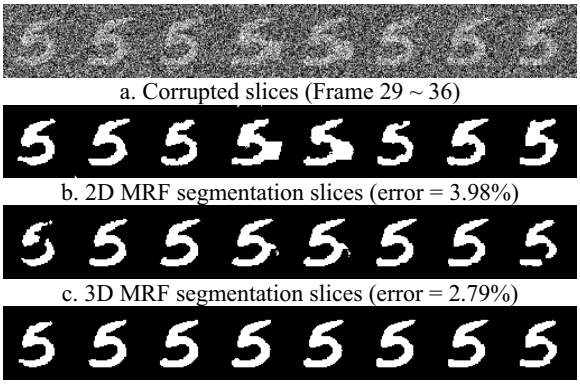

a. Corrupted slices (Frame 29 ~ 36)

b. 2D MRF segmentation slices (error = 3.98%)

c. 3D MRF segmentation slices (error = 2.79%)

d. 3D MRF + DM segmentation slices (error = 1.62%)

Fig. 3. Synthetic volume

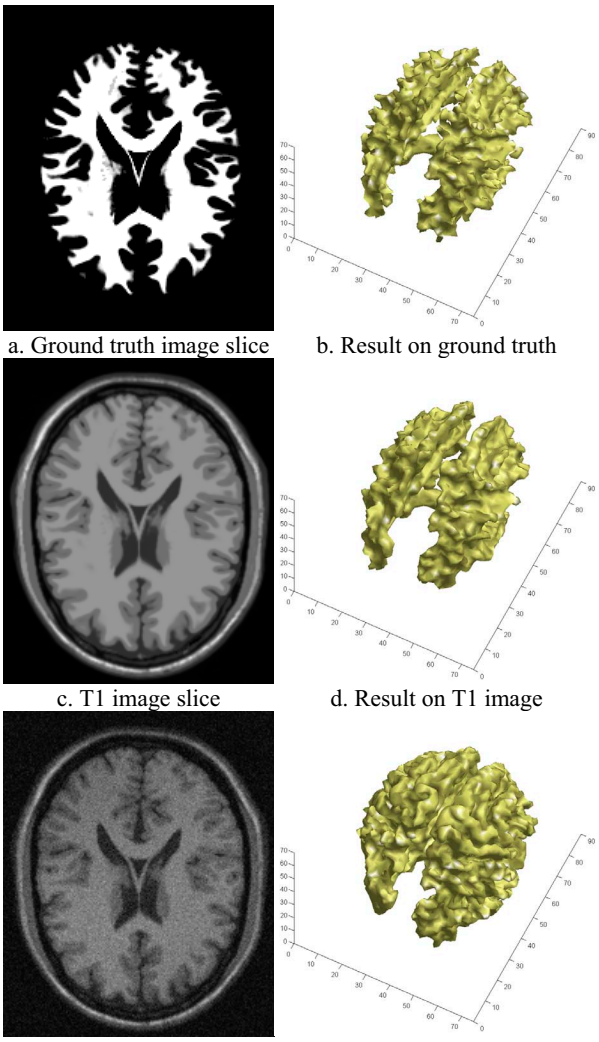

a. Ground truth image slice b. Result on ground truth

c. T1 image slice d. Result on T1 image

e. Noisy T1 image slice f. Result on noisy T1 image

Fig. 4. Simulated Brain Volume

Without the shape prior, our present method uses the 3D MRF and the Marching Cubes algorithm to build the initial deformable surface. An alternative initialization method is to start a GDM or a balloon from inside the desired region, making the initialization easier and less prone to initial outliers, even in the presence of multiple regions in the image [4]. However, this may increase the complexity of computation and reparametrization, which, in the 3D case, is still of main concern.

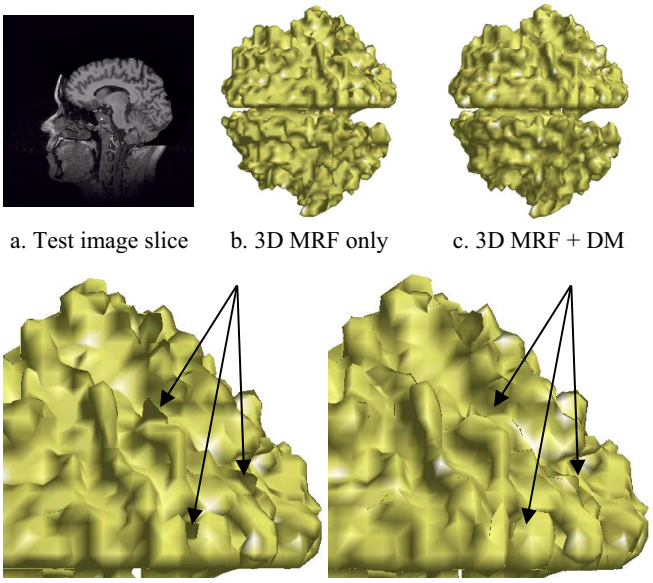

a. Test image slice b. 3D MRF only c. 3D MRF + DM

d. Upper-right part of b e. Upper-right part of c

Fig. 5. Real Brain Volume

7. REFERENCES

[1] S.M. Choi, J.E. Lee, J. Kim, and M.H. Kim, "Volumetric Object Reconstruction Using the 3D-MRF Model-Based Segmentation", *IEEE Transaction on Medical Imaging*, 16(6), 1997.

[2] L.D. Cohen and I. Cohen, "Finite-Element Methods for Active Contour Models and Balloons for 2-D and 3-D Images", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11), 1993.

[3] Y. Zhang, M. Brady, and S. Smith, "Segmentation of Brain MR Images Through a Hidden Markov Random Field Model and the Expectation-Maximization Algorithm", *IEEE Transaction on Medical Imaging*, 20(1), 2001.

[4] R. Huang, V. Pavlovic, and D.N. Metaxas, "A Graphical Model Framework for Coupling MRFs and Deformable Models", *Proceedings of CVPR*, 2004.

[5] J.V. Miller, D.E. Breen, W.E. Lorensen, R.M. O'Bara, and M.J. Wozny, "Geometrically Deformed Models: A Method for Extracting Closed Geometric Models from Volume Data", *Computer Graphics (Proceedings of SIGGRAPH '91)*, 25(4), 1991.

[6] T. Chen and D.N. Metaxas, "Gibbs Prior Models, Marching Cubes, and Deformable Models: A Hybrid Framework for 3D Medical Image Segmentation", *Proceedings of MICCAI*, 2003.

[7] Z. Ghahramani, "On Structured Variational Approximations", *Technical Report CRG-TR-97-1*, 1997.

[8] W.E. Lorensen and H.E. Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm", *Computer Graphics (Proceedings of SIGGRAPH '87)*, 21(4), 1987.

[9] BrainWeb: Simulated Brain Database. http://www.bic.mni.mcgill.ca/brainweb/

[10] The Stanford volume data archive. http://graphics.stanford.edu/data/voldata/