DEPARTMENT OF STATISTICS

University of Wisconsin

1300 University Ave.

Madison, WI 53706

TECHNICAL REPORT NO. 1105

February 15, 2006

Magnetic Resonance Image Segmentation with Thin Plate Spline Thresholding [1]

Xianhong Xie[a], Moo K. Chung[b,c], Grace Wahba[b, 2]

[a]Functional Neuroimaging Lab, Department of Psychiatry

Weill Medical College of Cornell University, New York, NY 10021, USA

[b]Department of Statistics, and Biostatistics and Medical Informatics

University of Wisconsin, Madison, WI 53706, USA

[c]Waisman Laboratory for Brain Imaging and Behavior

University of Wisconsin, Madison, WI 53706, USA

[1]Corresponding author address: Xianhong Xie, Functional Neuroimaging Lab, F-1307, Department of Psychiatry, Box 140, Weill Medical College of Cornell University, 1300 York Ave, New York, NY 10021. E-mail: xix2004@med.cornell.edu.

[2]Research of Xianhong Xie and Grace Wahba partially supported by NSF Grant DMS0072292 and NIH Grant EY09946.

Magnetic Resonance Image Segmentation with Thin Plate Spline Thresholding[‡]

Xianhong Xie[a], Moo K. Chung[b,c], Grace Wahba[b]

[a]Functional Neuroimaging Lab, Department of Psychiatry

Weill Medical College of Cornell University, New York, NY 10021, USA

[b]Department of Statistics, and Biostatistics and Medical Informatics

University of Wisconsin, Madison, WI 53706, USA

[c]Waisman Laboratory for Brain Imaging and Behavior

University of Wisconsin, Madison, WI 53706, USA

February 15, 2006

[‡]Corresponding author address: Xianhong Xie, Functional Neuroimaging Lab, F-1307, Department of Psychiatry, Box 140, Weill Medical College of Cornell University, 1300 York Ave, New York, NY 10021. E-mail: xix2004@med.cornell.edu.

Abstract

We propose a new method for the T1-weighted magnetic resonance image (MRI) segmentation. Thin plate splines are fitted to overlapping blocks of an image slice and thresholds are found. The knots and the smoothing parameters of the splines are chosen by a modified version of the generalized cross validation criterion. Each block is associated with a weighting function, which serves to blend the splines together as well as the thresholds in a smooth fashion. The blended image is then thresholded to get the boundaries between gray matter, white matter, cerebrospinal fluid, and others. We tested the method on MGH CMA 20 normal data. The results show that our method achieves good segmentation compared to human segmentation and SPM segmentation. Also our method generates subpixel results and handles the partial volume effect in the model. The new method has the advantage of being less dependent on image non-uniformity correction.

1 Introduction

Segmentation of magnetic resonance (MR) images is an important part of brain imaging research. The segmentation can facilitate the diagnosis of neurological diseases. It can also be used as a visualization aid for researchers, or as a preprocessing step for other studies. There are some inherent difficulties associated with image segmentation; among them are RF coil inhomogeneity, brain tissue susceptibility, and other systematic artifacts. Various preprocessing steps have been proposed to deal with some or all of these difficulties. After preprocessing, a segmentation method then can be used to classify the voxels in the whole brain volume into three different tissue types: grey matter (GM), white matter (WM), and cerebrospinal fluid (CSF). The segmentation methods we have seen so far can be roughly grouped into 2 categories: intensity based or surface based.

Regarding intensity-based classifiers, neural network classifier (Morrison and Attikiouzel, 1992; Ozkan et al., 1993; Kollokian, 1996; Wang et al., 1998), the k-nearest neighbor classifier (Bezdek et al., 1993) or a finite Gaussian mixture modeling (Bezdek et al., 1993; Kapur, 1995)

can be used for classifying each voxel into 3 different classes. In particular Gaussian mixture modeling assumes the image intensity values follow the mixture of two or more Gaussians and the unknown parameters of Gaussian distributions are estimated by maximizing the likelihood functions possibly via the expectation maximization (EM) algorithm or other optimization techniques. The widely used SPM'99 brain image analysis package (Wellcome Department of Cognitive Neurology, London, UK, URL http://www.fil.ion.ucl.ac.uk/spm) is based on a Bayesian Gaussian mixture modeling with a prior probability image generated by averaging the image intensity for large number of subjects (Ashburner et al., 1997; Ashburner and Friston, 2000). Based on a prior probability of each voxel being the specific tissue type, a Bayesian approach is used to get a better estimate of the posterior probability. This Bayesian update of the probability is iterated many times until the probability converges. The resulting probability is interpreted as the probability of each voxel belonging to one of three tissue types.

Instead of the above intensity-based segmentation techniques, surface-based segmentation techniques have begun to emerge. The advantage for surface-based segmentation methods is the possible reduction of the partial volume effect (Tohka et al., 2004), when triangular meshes are used. Triangular meshes are not constrained to lie on voxel boundaries. Instead the triangular meshes can cut through a voxel, which can be considered as correcting where the true boundary ought to be and reducing the partial volume effect. Deformable surface modeling (Terzopoulos et al., 1988; Davatzikos and Bryan, 1995; Dale and Fischl, 1999; MacDonald et al., 2000) can be used to segment tissue boundaries by either solving a partial differential equation or optimizing an objective function.

Recently isosurface modeling, also known as a level set method (Sethian, 1999; Osher and Paragios, 2003) seems to show promise in tissue boundary segmentation and has been used in segmenting the sagittal section of the corpus callosum (Hoffmann et al., 2004). A related approach to the image segmentation problem is the model proposed in Mumford and Shah (1985), where a piecewise smooth function is fitted to the image data, with the discontinuities happening only on the boundaries between different tissue types. The solution

can be obtained by optimizing an objective function iteratively.

In this paper, we propose and validate a new tuned thin plate spline thresholding method on image slices. In validating segmentation results, most studies compare the performance of their algorithms against expert manual brain segmentation (Boesen et al., 2004) or synthetic data sets. We will follow this rule as well.

2 Review of Thin Plate Splines

A thin plate spline (TPS) is the minimizer of the following optimization problem (Wahba, 1990, pp. 30–31)

$$\frac{1}{n}\sum_{i=1}^{n}(y_i - f(x_1(i),\cdots,x_d(i)))^2 + \lambda J_m^d(f), \tag{1}$$

where y_i is the i-th data point, $(x_1(i),\cdots,x_d(i))$ is a d-dimensional vector, n is the total number of observations, and J_m^d is a smoothness penalty functional involving m derivatives in d-dimensions.

The smoothing parameter λ will be chosen by some criterion to be described later. The penalty functional $J_m^d(f)$ for the special case $m = 2$, $d = 2$ is defined as

$$J_2^2(f) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}(f_{x_1x_1}^2 + 2f_{x_1x_2}^2 + f_{x_2x_2}^2)dx_1dx_2. \tag{2}$$

This is the one we will use in the paper. In the $m = 2$, $d = 2$ case for a unique minimizer to (1) exists it is necessary that a least square fit to these data at $\{t_i = (x_i(i), x_2(i)) : 1{\leq}i{\leq}n\}$ is unique, a condition obviously satisfied with the regular pixel data in two dimensions. Under this condition, the unique solution has the representation

$$f_\lambda(t) = d_0 + d_1x_1 + d_2x_2 + \sum_{i=1}^{n} c_iE_m(t - t_i), \tag{3}$$

where $t = (x_1, x_2)$ and $E_m(\tau) = \theta\|\tau\|^2 \ln\|\tau\|$. The constant θ will be absorbed into the

smoothing parameter λ. Details for general m and d can be found in Wahba (1990).

The coefficients $\{d_\nu\}$ and $\{c_i\}$ are linear in the y_i's, so we can write

$$\hat{y} \stackrel{\text{def}}{=} (f_\lambda(t_1), \cdots, f_\lambda(t_n))' = A(\lambda)y, \tag{4}$$

where $y = (y_1, \cdots, y_n)'$, and $A(\lambda)$ is called the smoothing matrix.

The smoothing parameter λ in the minimization problem (1) can be chosen by the generalized cross validation criterion with possibly a constant factor α which modifies the equivalent degrees of freedom of the spline (Luo and Wahba, 1997);

$$\alpha GCV(\lambda) = \frac{\|(I - A(\lambda))y\|^2/n}{[1 - \alpha \text{tr}(A(\lambda))/n]^2}. \tag{5}$$

The factor α should be a real number no less than 1. In the present work, where a function with jumps is being fitted, the conditions for optimality of the usual $\alpha = 1$ case are violated (Wahba, 1990). In this work α will be determined empirically once and for all for data from the present experimental setup.

Thin plate splines can be further approximated by fewer basis than those appearing in (3). The approximation has the form

$$\tilde{f}_\lambda(t) = d_0 + d_1 x_1 + d_2 x_2 + \sum_{s_l \in \Omega} c_l E_m(t - s_l), \tag{6}$$

where Ω is the set of knots ($\Omega \subset \mathbb{R}^2$). Each Ω specifies a knot configuration.

3 Method

TPS Thresholding Algorithm

Our method fits thin plate splines to overlapping blocks of an image slice, and blends the splines together smoothly. A similar idea was used in Wood et al. (2002). The main differences are that we choose the smoothing parameters differently and we use explicit subdi-

visions. In addition, our method obtains thresholds on every block; and the thresholds are blended the same way as the splines are blended.

A brief overview of the algorithm goes as follows. First, we divide the slice into overlapping blocks. Secondly, we fit thin plate splines to the image intensities at each block with different number of knots, and select the knot configuration that gives us the smallest αGCV score on every block. Thirdly, we fit the thin plate splines with the knot configurations found in the last step, and predict on a very fine grid. We also find the local thresholds with the k-means algorithm. Finally, we blend the predicted block images and the thresholds with smooth weighting functions. The thresholds are used on the blended image to get the boundaries between GM, WM, CSF, and other tissue types.

Step 1: Partitioning of Slice Image

With a typical slice of size 256x256 pixels, we first clip some surrounding empty space in the slice image. Many software programs can do this; and manual clipping is quite easy to do too. Once the clipping is done, we can divide the slice into blocks of size about 50x50 pixels. The users can determine the sizes of the blocks. A rule of thumb is to have all the tissue types (i.e. gray matter, white matter, CSF, and others) in every block. The same idea was used in Kovacevic et al. (2002). Following a similar line of thinking, we allow some degree of overlapping between adjacent blocks (horizontally, vertically, or diagonally). The overlapping proportion between each pair of horizontally or vertically adjacent blocks is about one half of the pixels in either of the blocks (see Figure 1(a)). This results in each pair of diagonally adjacent blocks having about one fourth of the pixels in the individual blocks overlapped. The histograms for the image intensities of two adjacent blocks are shown in Figure 1(b). We can see the bumps in the histograms.

Step 2: Finding Optimal αGCV Scores

After the partitioning is done, we fit thin plate splines to each block. Note that even for one block (of the size 50x50), there are more than two thousand data points. There can

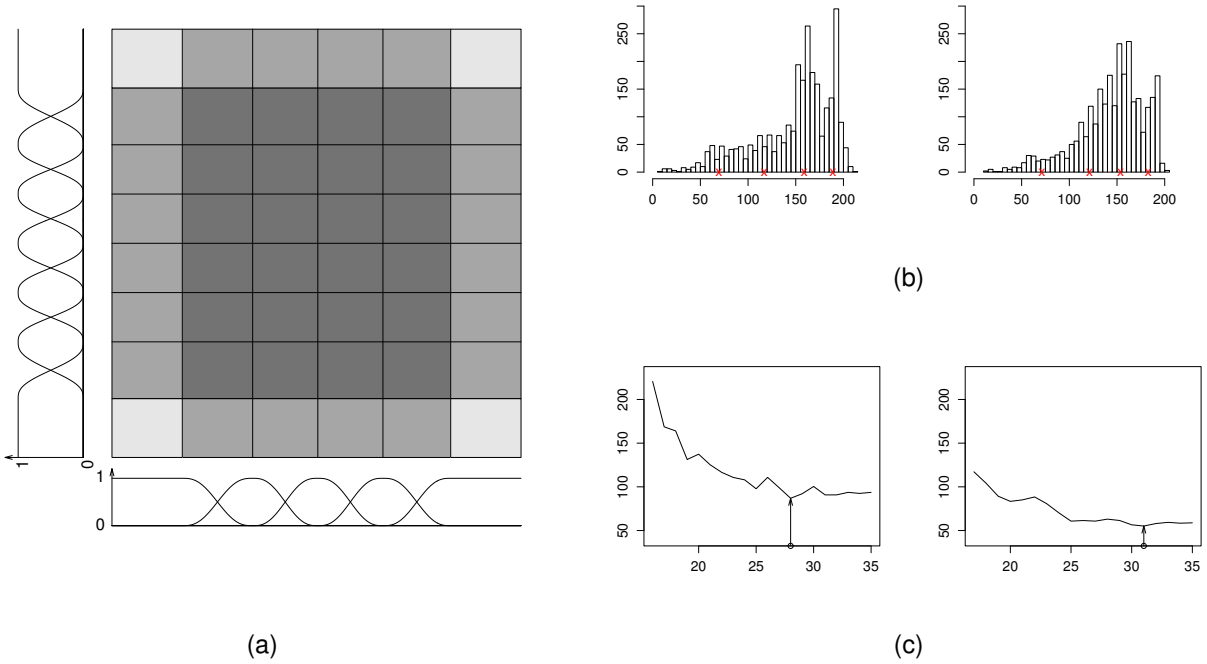

Figure 1: Overlapping scheme, αGCV curves and histograms: (a) overlapping scheme of one slice with 5 by 7 blocks (horizontally and vertically respectively). Each 4 adjacent shaded rectangles form one block, with different shades representing different numbers of overlapping (light grey=1, medium gray=2, dark gray=4). The weighting functions at each direction are given below and to the left of the plot. (b) histograms of intensities for two adjacent blocks (the centers found by k-means on predicted block images shown as red x's). (c) αGCV curves for the pair of adjacent blocks (minimum of αGCV shown with arrows).

be sharp boundaries between different tissue types within one block. However, the image should be considered relatively smooth within each tissue type in a block. To fit a thin plate spline with 2500 data points as knots is not only computationally ineffective, but also unnecessary. A remedy for this is to use a subset of the 2500 knots as an approximation to the original spline that uses every data point as knot (Luo and Wahba, 1997). Since the slice image is measured on a regular grid, a further approximation is to allow the knots to not fall on the pixel grid, but only require them to be equally spaced, where the knot grid size is approximately proportional to the pixel grid size. In our αGCV search, we bound the ratio of the number of the knots to the total number of pixels in every block within the interval $[0.02, 0.45]$, and find the lowest αGCV score for each given number of knots with respect to the smoothing parameter λ, we then minimize the scores over the number of knots. In this work an empirical value of $\alpha = 2$ was chosen once and for all, for all of the images. $\alpha = 2$ will result in a smoother signal and greater noise suppression than $\alpha = 1$. Although it may seem counterintuitive in the present context where one is trying to locate the boundary between two different regimes, each of which is roughly constant, this apparent oversmoothing reduces or eliminates any Gibbs effect and oversensitivity to noise and reduces the segmentation problem to a careful choice of threshold. The parameter α can be fixed at 2, or chosen once and for all for any particular experimental setup.

Step 3: Predicting TPS and Thresholding

Using the optimal knot configuration found in the last step, we fit a thin plate spline to each block with the given configuration. A fine grid is laid on the block, with every pixel divided into 8 by 8 subpixels and the thin plate spline predicted on the grid. The use of the fine grid is to get a smoother image, which will be beneficial for the thresholding later. Another advantage of the fine grid is that we can get subpixel level segmentation and smoother boundaries. To calculate the thresholds on every block, we use the k-means algorithm, which is simple, fast and efficient. The k-means algorithm is used with 4 centers corresponding to white matter (w), gray matter (g), cerebrospinal fluid (c), and empty space

(e), in the order of the intensity values of these regions appearing in a T1-weighted MR image from the highest to the lowest. Two examples of the centers found by k-means are given in Figure 1(b), which shows that the algorithm is doing a reasonably good job. We tried both 3 centers and 4 centers with k-means; it appears that the 4 center setting gives us better tissue boundaries. Once the centers have been found, we calculate the thresholds in the following way (Kovacevic et al., 2002)

$$t_{ec} = (m_e + m_c)/2 \tag{7}$$

$$t_{cg} = (m_c + m_g)/2 \tag{8}$$

$$t_{gw} = (m_g + m_w)/2, \tag{9}$$

where m_e, m_c, m_g, m_w are the centers found by the k-means algorithm, and t_{ec}, t_{cg}, t_{gw} are the thresholds to be used in the later step, where t_{ec} is the threshold between e and c, and so forth.

Step 4: Blending Block Images and Thresholds

Having done all the predicting and thresholding, we can now blend the block images together along with the thresholds using the weighting functions for each subblock, depicted schematically in Figure 1(a). In both the horizontal and the vertical direction, a pair of functions $1 - f(s)$ and $f(s)$ is used, where

$$f(s) = \begin{cases} 0 & , \text{ if } s \leq 0 \\ s^3(6s^2 - 15s + 10) & , \text{ if } 0 < s < 1 \\ 1 & , \text{ if } s \geq 1. \end{cases} \tag{10}$$

Note $f(s)$ goes smoothly (second order differentiable) from 0 to 1, and takes only nonnegative values (it is a 1-d quintic spline). In the case a subblock is covered by only 2 adjacent blocks, either horizontally or vertically, the weighting pair $1 - f(s)$ and $f(s)$ is used for left (lower) block and right (upper) block respectively. In the case a subblock is covered by 4

adjacent blocks, a tensor product weighting scheme is used. The weighting functions become $[1 - f(x)][1 - f(y)]$, $f(x)[1 - f(y)]$, $[1 - f(x)]f(y)$, and $f(x)f(y)$ for the lower left block, lower right block, upper left block, and upper right block respectively. For each overlapped subblock, the pixel coordinates are scaled so that the lower left endpoint of the subblock is mapped to $(0, 0)$, and the upper right endpoint of the subblock is mapped to $(1, 1)$.

For the blending of each of the three thresholds, we use the same scheme. The difference between the blending of the thresholds and the blending of the block images is that we have constant matrices in the place of block images.

When all the blending is done, we can display the segmentation result by drawing contours at level $= 0$ on the difference images between the blended image and the blended thresholds. Many software programs are available for this purpose.

Subjects and Image Acquisition

The 20 normal magnetic resonance brain data sets and their manual segmentations that we examine here were provided by the Center for Morphometric Analysis at the Massachusetts General Hospital and are available at the URL http://www.cma.mgh.harvard.edu/ibsr/. The coronal three-dimensional T1-weighted SPGR MRI scans were performed on two different imaging systems. Ten FLASH scans on four males and six females were performed on a 1.5 tesla Siemens Magnetom MR System (Iselin, NJ) with the following parameters: TR/TE 40/8 ms, flip angle 50, 30cm field of view, 3.1mm slice thickness, 256x256 matrix. Ten 3D-CAPRY scans on six males and four females were performed on a 1.5 tesla General Electric Signa MR System (Milwaukee, WI), with the following parameters: TR/TE 50/9 ms, flip angle 50, 24cm field of view, 3.0mm slice thickness, 256x256 matrix. Each image volume has about 60–65 slices. The data sets were used in Shan et al. (2002) and elsewhere.

4 Evaluation of TPS Thresholding Method

We used 5 subjects from the 20 normal data provided by MGH CMA. The selection scheme was as follows: we sorted the subjects based on their id's (1_24, 2_4, 4_8, $\cdots$), and selected the 2nd, 6th, 10th, 14th, 18th subjects. We applied our TPS segmentation method to one coronal slice near the middle of the brain for each subject. We then applied the SPM brain image analysis package, which works on the entire image of the subject, and extracted the corresponding slice for comparison. Note that we have the manual segmentation for the MGH data, so we have 3 segmentations in total. To compare each pair of segmentations, we used two measure of similarity. One is the correlation coefficient. The other one is the kappa index defined as

$$\kappa(S_1, S_2) = \frac{2|S_1 \cap S_2|}{|S_1| + |S_2|}, \tag{11}$$

where S_1, S_2 are the sets of pixels classified as one tissue type by two given segmentation methods, and $|\cdot|$ the number of elements in the set in question. This measure has been used in Shan et al. (2002), Kovacevic et al. (2002), and Zijdenbos et al. (1994). It has the nice property that two equally sized regions that overlap each other with half of their areas result in an index $\frac{1}{2}$. Also, the index is sensitive to both differences in sizes and locations of S_i's. A slight variation of the kappa index is the Jaccard index (Shan et al., 2002) which differs from the kappa index only in the constant and the denominator. These two criteria actually gave us the same conclusions, we will stick with the kappa index which seems to be the more popular criterion.

Since the TPS segmentation gives us subpixel level results, we need to convert them to the pixel level to be easily compared to the manual and SPM segmentations. The way we did it was to calculate the proportion of the number of subpixels in every pixel that belongs to each class (Figure 6(c)). Thus, we obtained a 4-tuple at every pixel, which sums to 1. The pixel level proportions were then used to calculate the correlation coefficients between the TPS and the other two methods. To get the kappa index, we thresholded both the

TPS proportion outputs and the SPM probability outputs by some given constant (0.5 for both GM and WM) to get the classification, then equation (11) was used. For example, the formula for calculating the kappa index for TPS vs SPM is

$$\kappa(P_T, P_S) = \frac{2|(P_T > 0.5) \cap (P_S > 0.5)|}{|P_T > 0.5| + |P_S > 0.5|}, \tag{12}$$

where P_T and P_S are the proportion output from TPS and the probability output from SPM respectively. We are only interested in the gray matter and the white matter proportions. So we computed the above mentioned indices for each pair of the segmentations on gray matter and white matter only.

5 Results

We display the segmentation results for the manual method, the TPS method, and the SPM method applied to one image slice in Figures 3, 4, and 5 respectively. The gray level image of the original slice is shown in Figures 2. Note that the manual segmentation gives the discrete classification (GM, WM, CSF, and other); the TPS method generates a predicted image with 3 thresholding fields; while the SPM method produces one probability image for each tissue class. These are reflected in the contour plots. With the subpixel property built in the algorithm, the TPS segmentation shows smoother boundaries than the other methods. Even at local levels (Figure 6(a) and 6(b)), the TPS method still traces the boundary between GM and WM well without being too wiggly. The similarity measurements (with mean and standard deviation summary) between the three methods on all 5 subjects are given in Table 1. We can see that the numbers are close, with the mean coefficients for the TPS against manual and those for the SPM against manual all within one standard deviation of each other. If we count the number of times TPS is doing better than SPM in terms of each index, and vice versa, we can find that there is no definite winner. Note that for one of the subjects (subject 4), both the TPS and the SPM failed on the segmentation. But TPS did a better job than did the SPM. Our method seems to be less sensitive to image non-uniformity.

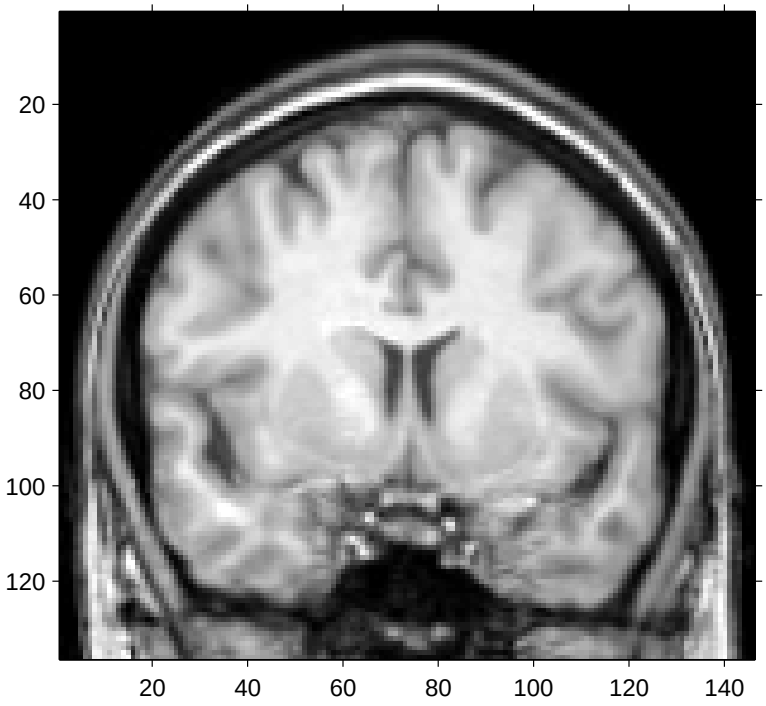

Figure 2: Original slice image in the gray scale

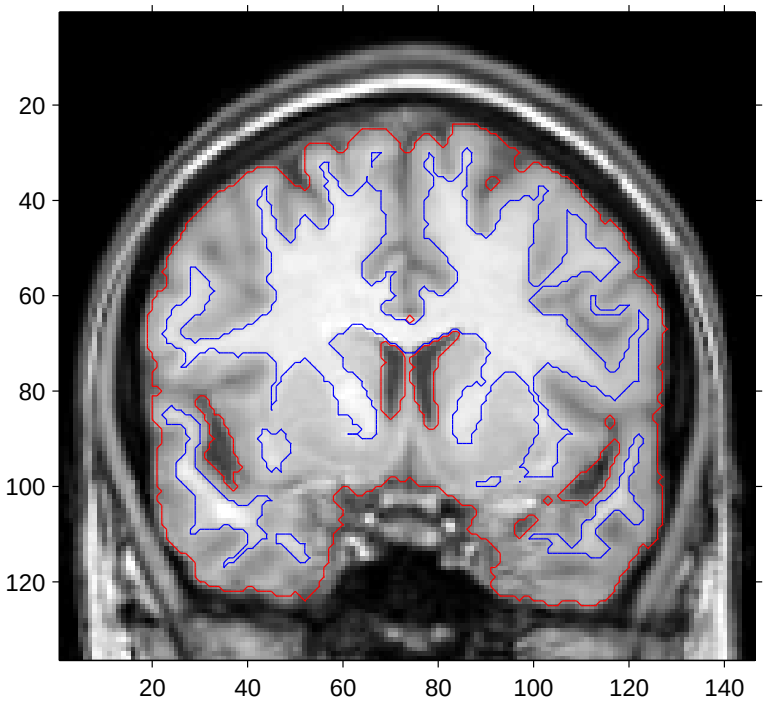

Figure 3: Overlay of the manual segmentation on the original slice

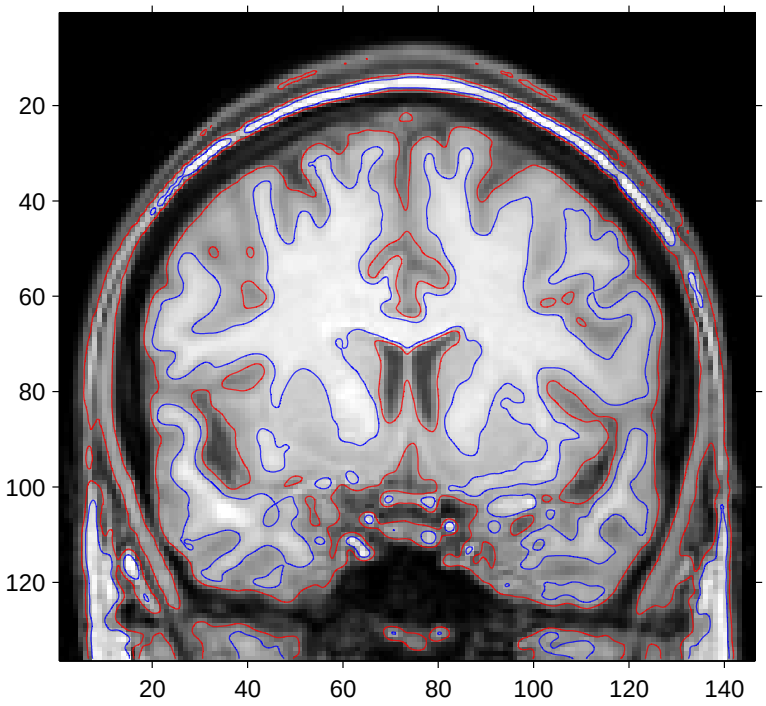

Figure 4: Overlay of the TPS segmentation on the original slice

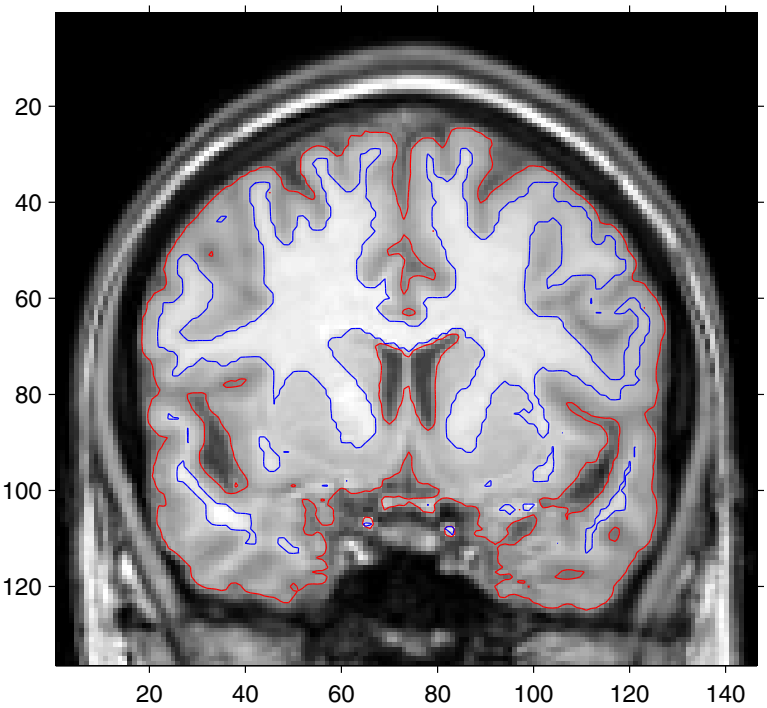

Figure 5: Overlay of the SPM segmentation on the original slice. Both the gray and the white probability outputs by SPM were thresholded at level=0.5 to get the red and the blue contour lines respectively.

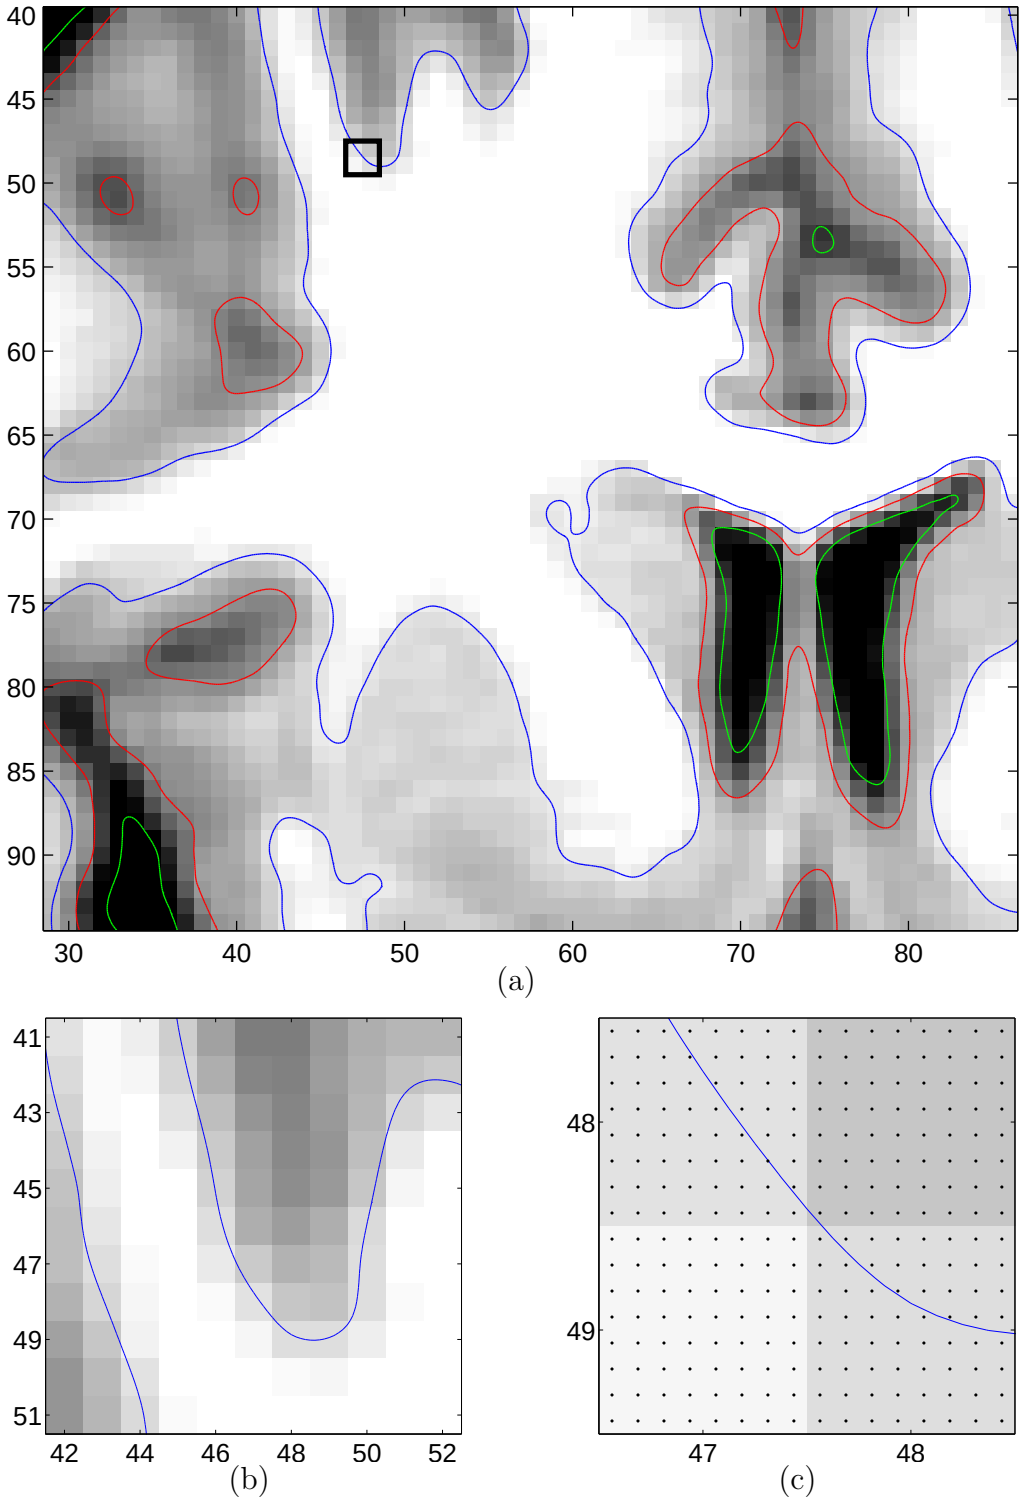

Figure 6: Zoomed plots of the TPS segmentation: (a) TPS result zoomed (b) TPS result zoomed to a smaller region (c) dot plot with the centers of the subpixels shown as small dots, from the square region outlined in heavy black in (a).

Table 1: Comparison of Three Segmentation Methods with Mean and SD Summary

	subject no.	corr. coef.		kappa index	
		GM	WM	GM	WM
TPS vs Manual	1	0.660	0.827	0.835	0.872
	2	0.702	0.757	0.841	0.828
	3	0.654	0.787	0.814	0.852
	4	0.410	0.678	0.724	0.770
	5	0.612	0.791	0.776	0.840
mean (sd)		0.608(0.115)	0.768(0.056)	0.798(0.049)	0.832(0.038)
SPM vs Manual	1	0.675	0.846	0.866	0.862
	2	0.686	0.839	0.875	0.855
	3	0.637	0.810	0.831	0.838
	4	0.091	0.672	0.598	0.652
	5	0.450	0.803	0.771	0.824
mean (sd)		0.518(0.250)	0.794(0.071)	0.788(0.114)	0.806(0.087)
TPS vs SPM	1	0.806	0.883	0.875	0.869
	2	0.626	0.759	0.803	0.764
	3	0.734	0.822	0.831	0.837
	4	0.426	0.767	0.687	0.710
	5	0.645	0.800	0.778	0.805
mean (sd)		0.647(0.143)	0.806(0.050)	0.795(0.070)	0.797(0.062)

6 Conclusions

Our method is an intensity based method and it does simple thresholding. Thin plate splines are used to smooth the image. Our results show that the new method is doing a reasonably good job in terms of segmentation. The TPS method has the advantage of generating subpixel level results and smoother boundaries. The partial volume effect is addressed by the subpixel segmentation. This property further shows that the TPS method has the potential for more accurately segmenting magnetic resonance images, including curvatures of the boundaries, and potentially surfaces. Our method tackles the image non-uniformity through the local thresholding and the blending. It provides a good alternative to the other known segmentation methods.

7 Future Work

The results we have shown so far are for 2-dimensional slices only. More work is under way to study the 3-dimensional segmentation. The idea described above can be applied to the 3-d case without much change. But the 3-dimensional work will prove to be more computationally challenging and also more interesting.

To give a preview of what would happen in the 3-d TPS segmentation, we list some of the challenges. For a typical brain volume we are segmenting, the dimension is 181x217x181. To use the overlapping idea and not to stretch the computer too much on its memory and the computing time, we can only apply the 3-d TPS to a cube of about the size 20x20x20. With half of voxels in the adjacent cubes overlapped, the number of overlapping cubes could go up to $17 \times 21 \times 17 = 6069$. This poses two difficulties. First, the computing time will be enormous. But with the using of Condor (a batch processing system for clusters of machines, URL http://www.cs.wisc.edu/condor/), this problem can be greatly reduced. Secondly, because the cube size is not big enough to contain all the tissue types, the thresholding algorithm might have to deal with different numbers of classes in different cubes. A way to get around this is to combine adjacent cubes to big ones and then apply the thresholding. Based on our experiments, the combined cube of about the size 60x60x60 is big enough for thresholding.

Although the 3-d volume segmentation will certainly be much more difficult than the 2-d slice segmentation, we expect that we can get good segmentations with clever ways of modifying our algorithm.

8 Acknowledgements

The first author thanks Doug Nychka, Simon Wood, and Douglas Bates for their assistance with questions on software. Also, we thank John Carew for helpful discussions. The Condor batch processing system was used extensively in the computation, we thank the Condor group for their nice software as well. The research of Xianhong Xie and Grace Wahba was

partially supported by NSF Grant DMS0072292 and NIH Grant EY09946.

References

Ashburner, J., Friston, K. J., 1997. Multimodal image coregistration and partitioning - a unified framework. NeuroImage 6 (3), 209–217.

Ashburner, J., Friston, K. J., 2000. Voxel-Based Morphometry - The Methods. NeuroImage 11 (6), 805–821.

Ashburner, J., Neelin, P., Collins, D., Evans, A., Friston, K., 1997. Incorporating prior knowledge into image registration. NeuroImage 6, 344–352.

Bates, D. M., Lindstrom, M. J., Wahba, G., Yandell, B. S., 1987. GCVPACK-Routines for generalized cross validation. Commun. Statist.-Simula. 16 (1), 263–297.

Bezdek, J., Hall, L., Clarke, L., 1993. Review of MR image segmentation techniques using pattern recognition. Medical Physics 20 (4), 1033–1048.

Boesen, K., Rehm, K., Schaper, K., Stoltzner, S., Woods, R., Lüders, E., Rottenberga, D., 2004. Quantitative comparison of four brain extraction algorithms. NeuroImage 22, 1255–1261.

Chung, M. K., Dalton, K. M., Alexander, A. L., Davidson, R. J., 2004. Less white matter concentration in autism: 2D Voxel-Based Morphometry. NeuroImage 23, 242–251.

Dale, A., Fischl, B., 1999. Cortical surface-based analysis: I. segmentation and surface reconstruction. NeuroImage 9, 179–194.

Davatzikos, C., Bryan, R., 1995. Using a deformable surface model to obtain a shape representation of the cortex. In Proceedings of the Intl. Symp. on Comp. Vision, 212–217.

Hoffmann, T. J., Chung, M. K., Dalton, K. M., Alexander, A. L., Wahba, G., Davidson, R. J., 2004. Subpixel curvature estimation of the corpus callosum via splines and its application

to autism. 10th Annual Meeting of the Organization for Human Brain Mapping.
URL http://www.stat.wisc.edu/~mchung/papers/HBM2004/HBM2004thomas.html

Kapur, T., 1995. Segmentation of brain tissue from magnetic resonance images. Tech. Rep. AITR-1566.
URL citeseer.ist.psu.edu/article/kapur95segmentation.html

Kollokian, V., 1996. Performance analysis of automatic techniques for tissue classification in MRI of the human brain. Master's thesis, Concordia University, Montreal, Canada.

Kovacevic, N., Lobaugh, N. J., Bronskill, M. J., Levine, B., Feinstein, A., Black, S. E., 2002. A robust method for extraction and automatic segmentation of brain images. NeuroImage 17, 1087–1100.

Luo, Z., Wahba, G., 1997. Hybrid adaptive splines. Journal of the American Statistical Association 92 (437), 107–116.

MacDonald, D., Kabani, N., Avis, D., Evans, A. C., 2000. Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI. NeuroImage 12, 340–356.

Morrison, M., Attikiouzel, Y., 1992. A probabilistic neural network based image segmentation network for magnetic resonance images. In Proc. Conf. Neural Networks. Vol. 3. pp. 60–65.

Mumford, D., Shah, J., 1985. Boundary detection by minimizing functionals. In IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR). pp. 22–26.

Osher, S., Paragios, N. (Eds.), 2003. Geometric level set methods in imaging, vision and graphics. Springer Verlag, New York.

Ozkan, M., Dawant, B., Maciunas, R., 1993. Neural-network-based segmentation of multi-modal medical images: a comparative and prospective study. IEEE Trans. Med. Imaging 12, 534–544.

Sethian, J. A., 1999. Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision and materials science. Cambridge University Press, Cambridge, UK.

Shan, Z. Y., Yue, G. H., Liu, J. Z., 2002. Automated histogram-based brain segmentation in T1-weighted three-dimensional magnetic resonance head images. NeuroImage 17, 1587–1598.

Terzopoulos, D., Witkin, A., Kass, M., 1988. Constraints on deformable models. Artificial Intelligence 36 (1), 91–124.

Tohka, J., Zijdenbos, A., Evans, A. C., 2004. Fast and robust parameter estimation for statistical partial volume models in brain MRI. NeuroImage 23 (1), 84–97.

Wahba, G., 1990. Spline models for observational data. SIAM, Philadelphia, PA.

Wang, Y., Adah, T., Kung, S.-Y., 1998. Quantification and segmentation of brain tissues from MR images: A probabilistic neural network approach. IEEE Transactions on Image Processing 7 (8), 1165–1181.

Wood, S., Jiang, W., Tanner, M., 2002. Bayesian mixture of splines for spatially adaptive nonparametric regression. Biometrika 89 (3), 513–528.

Zijdenbos, A. P., Dawant, B. M., Margolin, R. A., Palmer, A. C., 1994. Morphometric analysis of white matter lesions in MR images: method and validation. IEEE Transactions on Medical Imaging 13 (4), 716–724.