Using Statistical Pattern Recognition Techniques to Control Variable Conductance Diffusion

Terry S. Yoo
James M. Coggins

Department of Computer Science, University of North Carolina
Chapel Hill, NC 27599-3175, USA

Abstract. We present an approach for controlling relaxation parameters in variable conductance diffusion. This approach incorporates a Bayesian classifier to perform a partial labeling of an image, followed by a diffusion step. Conductance values between pixels are controlled by statistical measurements made of the partial classification. Several iterations follow, interleaving partial labeling with diffusion steps until a convergence or stopping criterion is met. The method is suitable for performing diffusion within multi-valued images. It consistently controls relaxation parameters, even in the presence of noise. The method is presented along with results on phantom and MR images.

1 Introduction

The first task in computer aided visualization of medical images is to classify organs and other solid masses (such as tumors) into tissue types. Statistical classification methods seldom incorporate either the geometry of regions within the image or the scale at which information is represented. Alternatively, structural methods often ignore the distributions of intensity values within the image when making judgments about boundary properties and measures of homogeneity within regions. These weaknesses are aggravated when these methods are individually applied to multivalued images, images with more than one value per pixel. What is needed is a segmentation method that takes into account scale, boundary information, and the probability distribution of the intensities of the multivalued image. We introduce a statistically controlled version of variable conductance diffusion (VCD) that employs user supplied control points to estimate the class conditional densities of the coherent regions within the image to control the conductance parameters of VCD.

Performing VCD on multivalued images is not a new idea. Both Whitaker and Gerig independently explored techniques in VCD on images of multiple parameters [Whitaker 91b][Gerig 91b]. Gerig uses a statistical evaluation of the feature space represented by the intensities of multiple echoes in MR imaging to evaluate the performance of his boundary preserving filter. Whitaker generates feature spaces and performs mathematical measurements in those spaces to control conductance in his diffusion process. Both have speculated on the use of statistics to generate conductance functions.

The problem is how to align the incommensurate values of a multivalued image. Our approach is to query a user for sample points identified with particular classes within an image (tissue types in medical images). From these data samples we then

Published in Barrett, H.H. and A.F. Gmitro, Information Processing in Medical Imaging, Proceedings of the 13th International Conference, IPMI '93, Flagstaff, AZ USA, June 1993, Lecture Notes in Computer Science Volume 687, (Springer-Verlag: Berlin) 1993, pp. 459-471.

determine the statistical correlation among image values through estimations of the class conditional probabilities. Using a Bayesian approach, we pursue "boundariness," discrimination between tissue classes, and control diffusion based on our ability to discriminate between one type of tissue and another.

Our method controls not only conductance but also the relaxation parameters of VCD. As the diffusion continues, we modify our estimates of the class conditional probabilities, and reduce the feature space measurement scale.

We first present an overview of the progression of VCD ideas to date. We introduce a discrete implementation of multivalued VCD, then explain the weakness that we are addressing in this work, the analysis of images with multiple incommensurate data values per pixel. We then present the user supervised statistically controlled VCD method as a mechanism for overcoming these weaknesses. We end with some results of statistically controlled VCD on a computer generated phantom as well as an MR image.

2 Some Background on VCD

There is a progression of diffusion based filtering and segmentation ideas in the research literature. Isotropic diffusion is easily expressed as solutions to the following differential equation:

$$\nabla \bullet c\nabla I = \partial I/\partial t \tag{1}$$

where the constant c is the conductance value for the image and I is the pixel intensity at a particular location. Gaussians and their derivatives are solutions to this equation. ter Haar Romeny and his associates explore the Gaussian as a filter kernel in the generation and analysis of image scale [ter Haar Romeny 1991a,b].

Analysis through isotropic Gaussian scale often discards boundary information and other image features at one scale in return for particular features at different scale. In order to retain boundary information during the diffusion, Perona and Malik proposed an edge affected diffusion process [Perona 1988]. According to these authors, an anisotropic edge affected diffusion can be controlled by a scalar conductance function, $g(\|\nabla I(x, y, t)\|)$, whose argument is the gradient magnitude of the gradually evolving image. Their edge affected diffusion equation is

$$\nabla \bullet g(\|\nabla I(x, y, t)\|)\, \nabla I = \partial I/\partial t \tag{2}$$

where the function g contains a constant relaxation parameter k, and can typically be expressed as

$$g(\alpha) = e^{-\left(\frac{\alpha}{k}\right)^2} \tag{3}$$

Equation 2 requires the instantaneous measurement of gradient magnitude at a pixel sized scale. Instabilities in these conductance measures proposed by Perona and Malik led Whitaker to explore a multiscale version of this diffusion (called variable

conductance diffusion, VCD, or geometry limited diffusion) [Whitaker 91a]. Whitaker adds a convolution with a Gaussian whose size decreases over time to control the scale of the measurement of gradient magnitude. His modified diffusion equation is

$$\nabla \bullet g(\|G(s, t) \otimes \nabla I(x, y, t)\|) \nabla I = \partial I/\partial t \tag{4}$$

where $G(\sigma, t)$ is a Gaussian convolution kernel whose standard deviation gradually decreases over time and the symbol $\otimes$ represents a convolution. This equation exhibits improved stability over the equation (2) proposed by the previous authors.

Gerig and his associates independently developed a similar approach using a different function to control conductance [Gerig 90]. These diffusion systems have been used in medicine as a preparation for Bayesian classification. Typically the statistical classifier is followed by a morphological operation and a connected component analysis to include image structure into the segmentation [Kubler 90].

Gerig and Whitaker independently explore the potential of multivalued VCD on images of two or more parameters. Whitaker shows that gradient images can be diffused to form image descriptions from higher order derivative information of single valued images [Whitaker 91b]. Moreover, he shows that dissimilarity measures other than gradient magnitude, such as gradient direction, can be used to control diffusion. He also postulates that control of VCD can be enabled through statistical measurements of image features. Gerig *et.al.* use the Euclidean norm to combine the gradient magnitude of two echoes of MR images to generate a conductance value [Gerig 91b]. Their conductance term takes on the form:

$$\frac{\delta}{\delta t}\begin{pmatrix} I_1(\bar{x},t) \\ I_2(\bar{x},t) \end{pmatrix} = \nabla c_c \begin{pmatrix} \nabla I_1 \\ \nabla I_2 \end{pmatrix} + c_c \begin{pmatrix} \nabla^2 I_1 \\ \nabla^2 I_2 \end{pmatrix} \tag{5}$$

$$c_c(\bar{x}, t) = \sqrt{c_{c1}(\bar{x}, t)^2 + c_{c2}(\bar{x}, t)^2} \tag{6}$$

where the combined conductance component $c_c(x, t)$ is used as the conductance value in both image echoes I_1 and I_2.

isotropic diffusion
edge affected diffusion
variable conductance diffusion
multivalued variable conductance diffusion
user supervised multivalued variable conductance diffusion

There has been a progression of diffusion systems throughout the development of these ideas. The multivalued work performed by Gerig and by Whitaker requires a well chosen mathematically based dissimilarity measure to be selected to control conductance. The work presented here is a step toward using measurements of the image itself to control VCD.

3 A Discrete Implementation of Multivalued VCD

We have implemented a discrete version of the continuous process described in the equations presented above. The discrete variable conductance diffusion mechanism is crude, but indicates some research areas that still require significant development. The discrete multivalued variable conductance diffusion model is driven directly from the analysis of the image through local histograms or feature space in a iterative fashion. Simply, for each pixel location, x, let $\mathbf{I}_{x,t}$ represent the multivalued intensity of the pixel during the t'th iteration. Moreover, let the value of the pixel at the t+1 iteration be determined by the following:

$$\mathbf{I}_{x,t+1} = \mathbf{I}_{x,t} + \Delta\mathbf{I}_{x,t} \tag{7}$$

where $\Delta\mathbf{I}_{x,t}$ is defined by the following equation

$$\sum_{r}^{\text{neighborhood}(x)} \omega_i\left(\|x - r\|\right)\omega_f\left(\mathbf{I}_{x,t}, \mathbf{I}_{r,t}\right)\left(\mathbf{I}_{x,t}, \mathbf{I}_{r,t}\right) = \Delta\mathbf{I}_{x,t} \tag{8}$$

The functions $\omega_i(\|x - r\|)$, and $\omega_f(\mathbf{I}_{x,t}, \mathbf{I}_{r,t})$ are scalar values and represent weighting functions based upon measurements in image space and in feature space respectively. The first weight, $\omega_i(\|x - r\|)$, assumes that the image follows a Gibbs distribution; that is, that the value of a pixel is reinforced by neighboring pixel values. The $\omega_f(\mathbf{I}_{x,t}, \mathbf{I}_{r,t})$ weight assumes that individual classes of pixel type maintain a local distribution in feature space. In the current implementation, the two weighting functions are:

$$\omega_i\left(\|x - r\|\right) = \frac{1}{\sqrt{2\pi\sigma_i^2}}\, e^{-\frac{1}{2}\frac{\|x - r\|^2}{\sigma_i^2}} \tag{9}$$

$$\omega_f\left(\mathbf{I}_{x,t}, \mathbf{I}_{r,t}\right) = \frac{1}{\sqrt{2\pi\sigma_f^2}}\, e^{-\frac{1}{2}\frac{\left\|\mathbf{I}_{x,t} - \mathbf{I}_{r,t}\right\|^2}{\sigma_f^2}} \tag{10}$$

Notice that probability distribution in both image space as well as feature space (controlled by σ_i and σ_f respectively) are isotropic and not necessarily time dependent in this implementation.

The method for multivalued diffusion presented in this paper will address the construction of multivalued diffusion covariance measures and a means for creating a value for the multicluster analog of σ_f. The value of σ_i controls the spatial scale of the diffusion and is a gradually decreasing scalar function of time, t. That is, $\sigma_i = \sigma(t)$, where the value of $\sigma(t_m) > \sigma(t_n)$ for time $t_m < t_n$.

4 Incommensurate Intra-Pixel Values

Multivalued VCD as described by the equations above do not relate the separate values of each pixel. Essentially, due to the isotropic nature of the weighting function in feature space, it is assumed that the within pixel values are "normalized" around some expected value for a pixel class. This is not always true. Consider the 2 dimensional feature space described by a two echo MR image (Proton density weighted and T2 weighted) where we select the intensity of the proton density weighted values as one feature and the T2 intensity as the other. Figure 1 shows an example of this image, and the resulting feature space.

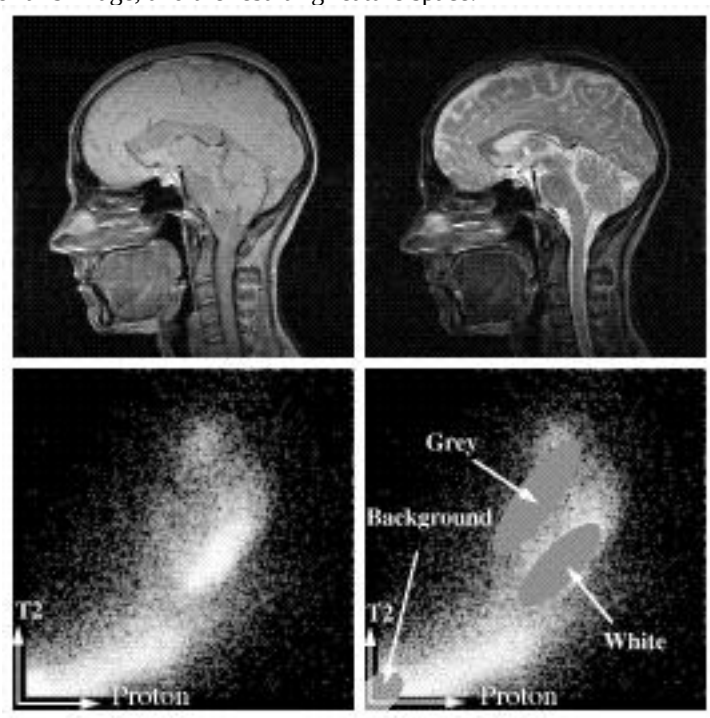

Figure 1 - Two valued MR slice of a human head. The upper images are the proton density and T2 weighted values, the lower views are the scatter plots of the image features. Lower Left: The feature space defined by the intensities of the two echoes, Lower Right: The class conditional densities associated with the control points.

If we overlay class conditional densities for white matter, gray matter, and cerebral spinal fluid (CSF), it is clear that the distributions of each of these three tissue classes are not isotropically distributed (see Figure 1, lower right). Therefore, the assumption that spatial weighting should not show directional preference in the feature space is shown to be invalid.

How then do we estimate this bias in feature space? More importantly, this directional bias will evolve during the filtering process. It is essential that we continue to track the orientation of the class conditional density and measure pixel similarity based on the changing image space.

5 Method

Probability distribution information will be supplied by the user as a set of user selected training points, providing the VCD system with semantic *a priori* knowledge of the data that would not otherwise be available. Throughout the many iterations of the discrete VCD method, new estimates of class conditional density are calculated from the control points specified in the initial user session. We employ a variation of a Bayesian classifier, slightly modified from that presented in the literature [Duda 73]. We iterate through the following steps:

1. Collect initial classification of image data
2. Calculate covariance matrices for each identified cluster using the image values from the current iteration. (use the original image for iteration 0)
3. For each pixel, calculate the diffusion vector $\Delta\mathbf{I}_{x,t}$ using a statistically based feature weighting function $\omega_f(\mathbf{I}_{x,t}, \mathbf{I}_{r,t})$
4. Adjust each pixel by the diffusion vector. (i.e. set $\mathbf{I}_{x,t+1} = \mathbf{I}_{x,t} + \Delta\mathbf{I}_{x,t}$)
5. Check for stopping criterion. If diffusion is to continue, go to step **2**, using the current image values $\mathbf{I}_{x,t+1}$ to calculate new covariance values.

Let us address each step in order. The first step is to identify the locations of tissue types of interest. The user is queried for sample points that can be identified by tissue type. A mouse based interface is provided so that the user can quickly identify many locations for each tissue class. It is important to note that the essential element acquired in this step is the location of the pixels in the training set, rather than specific intensity values. The intensity values will migrate throughout the iterative diffusion process; while intensity may vary, the spatial locations of important features remain fixed.

In step 2, let NC be the number of tissue classifications specified by the user. For each tissue class c_w (w = 1,2,...,NC) we calculate a covariance matrix $\mathbf{C}_w$. Specifically, let $\mathbf{I}_{p,t}$ represent a multivalued pixel in a particular tissue pattern and p = 1...n where n is the number of individual pixels in the user supplied tissue pattern U_w (i.e. $U_w = \{\mathbf{I}_1,...,\mathbf{I}_n\}$), and t represents the current iteration. Calculate the multivalued mean μ_w for the tissue pattern. If we represent the i-th value of μ as $\mu[i]$ and the i-th value of $\mathbf{I}_{p,t}$ as $\mathbf{I}_{p,t}[i]$, then the covariance matrix $\mathbf{C}_w$ for the tissue pattern p_w is estimated by:

$$\mathbf{C}_w(i, j) = \frac{1}{n} \sum_{p=1}^{n} \left(\mathbf{I}_{p,t}[i] - \mu[i]\right)\left(\mathbf{I}_{p,t}[j] - \mu[j]\right) \qquad (11)$$

We calculate the diffusion vector $\Delta\mathbf{I}_{x,t}$ in step 3 using a function similar to equation 4; however, we modify the feature weighting function $\omega_f(\mathbf{I}_{x,t}\ \mathbf{I}_{r,t})$ to include the available *a priori* information to normalize the similarity measures that control the diffusion. If $P(c_w)$ is the *a priori* probability of class c_w, then the new weighting function is:

$$\omega_f\left(\mathbf{I}_{x,t},\mathbf{I}_{r,t}\right)=\sum_{w=1}^{NC}P(c_w)e^{-\left((\mathbf{I}_{x,t}-\mu_w)^2/\mathbf{C}_w\right)-\left((\mathbf{I}_{r,t}-\mu_w)^2/\mathbf{C}_w\right)}e^{-\left((\mathbf{I}_{x,t}-\mathbf{I}_{r,t})^2/\mathbf{C}_w\right)} \qquad (12)$$

The first exponential term weights the importance of the pixel values of $\mathbf{I}_{x,t}$ and $\mathbf{I}_{r,t}$ relative to the tissue class c_w. The second exponential term represents a weighted measure of the similarity of the multivalued pixel intensities of $\mathbf{I}_{x,t}$ and $\mathbf{I}_{r,t}$ relative to the estimated Gaussian distribution of tissue class c_w. In the results presented in this paper, $P(c_w)$ is assumed to equal 1/NC.

It remains in step 4 to combine the current $\mathbf{I}_{x,t}$ with the diffusion vector $\Delta\mathbf{I}_{x,t}$ and check for the stopping criterion in step 5. In current applications, the stopping criterion is simply a number of iterations; however, termination could easily be based upon the minimum or maximum value of $|\mathbf{C}_w|$ (w = 1,...,NC) falling below some convergence threshold.

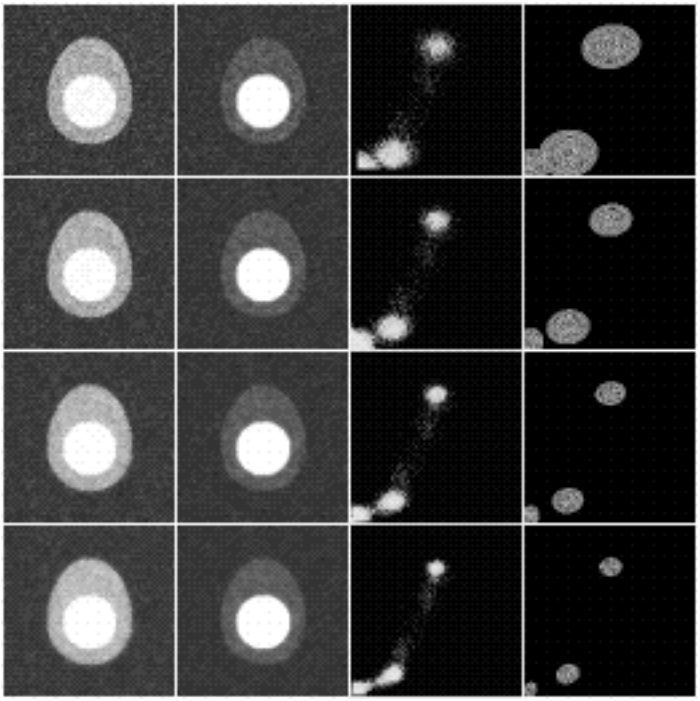

Figure 2 - simulated MR image of a computer generated egg phantom. Columns left-right: T1 image, T2 image, Feature space scatterplot, classifier distributions. Rows from top to bottom: 2 variable image through 30 iterations of statistically controlled VCD (Top original image, 2nd row - 10th iteration, 3rd row 20th iteration, 4th row 30th iteration)

6 Results

Figure 2 shows a simulated MR image of an egg shaped object as it passes through several iterations of VCD. The training classification included patterns of "egg yolk", "egg white" and background values, of approximately 100 pixels per pattern. The top of the set shows the original image, with simulated T1 and T2 weighted values followed by the scatter plot of the features. The rightmost column is a representation of the classification distributions of each of the training clusters.

Figure 3 shows a slice of an MR image taken from a scan of a human head. Only the T1 echoes are display in the figure. The left side shows the T1 weighted image with the original scatter plot of the pixel features. The right side shows the same image after several iterations using the statistically controlled VCD process. The training patterns included white matter (myelinated cerebral tissue) and gray matter (unmyelinated tissue) only. Note that although some degradation of detail or blurring is apparent, boundaries of regions in the user supplied classifications (gray matter and white matter) remain largely unaffected, while within the regions themselves, pixel values are more homogeneous.

The plots of the image features reveal that significant cluster improvement has occurred. The background pixel distribution has narrowed considerably, and delineation between feature clusters is becoming apparent. Figure 4 is a plot of the class conditional distributions based upon the covariance measures of the training pixels. After several iterations, the mean values of the user supplied partially classified pixels have remained fixed, while the variances, or the spread of the distributions have narrowed. From this view, we begin to see where the discrimination between white and gray matter lies.

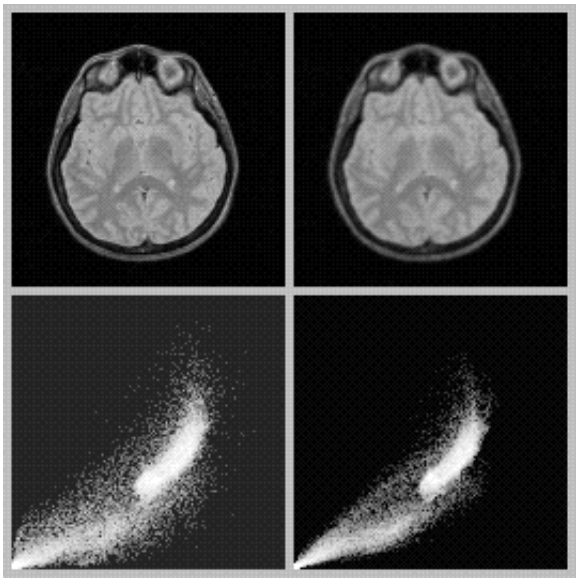

Figure 3 - VCD on an axial MR slice of a human head. The upper images are the T2 weighted values, the lower views are the scatter plots of the image features. Left: Before VCD, Right After several VCD iterations. (not shown, T1 weighted values for either before or after)

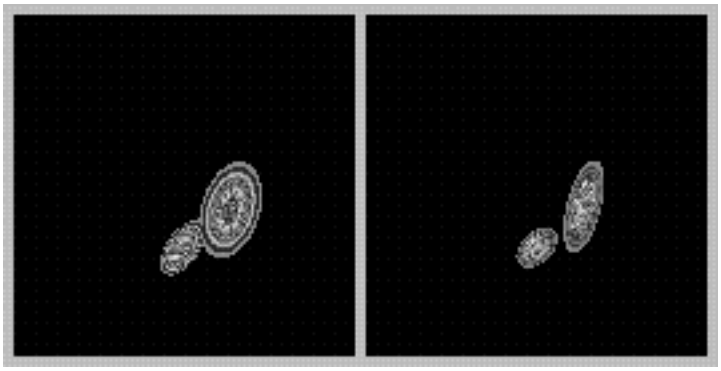

Figure 4 - Classification distributions for white matter and gray matter displayed within the feature space. Distributions are based on user supplied *a priori* locations of classified pixels. As pixel values diffuse, the means of the classification distributions remain stationary, while the variances narrow. Left: Before VCD, Right: After several VCD iterations.

7 Discussion

We have improved our control of multivalued VCD over the previous discrete implementation through the application of *a priori* information. Previously, the value for σ_f was a constant (the analog of k in a single valued VCD process), and isotropic within the feature space. Since there is no guarantee that the separate values within a pixel are commensurate, some means of understanding and compensating for the nature of the pixel features is required. Therefore, we apply statistical measures of the feature space to achieve commensurability.

We assume that there is a statistical correlation among pixel features. Measurements of pixel-to-pixel similarity are therefore made relative to estimates of class distributions from training sets supplied by the user. As pixel values diffuse, the distributions representing the classified tissue types narrow, raising the criteria for similarity, and automatically lowering conductance values.

7.1 Convergence

We can monitor the rate at which the distributions of the pre-classified pixels converge by measuring the determinants of the covariance matrices generated by each of the *a priori* pixel sequences. Figure 5 shows the covariance values through 100 iterations of the computer generated egg phantom. The large image space weight on the diffusion process early in the VCD sequence accounts for the rapid reduction in variance; during the later stages, narrow variance values slow diffusion, and the determinants subsequently do not fall as rapidly. This particular image shows exaggerated convergence due to the large contiguous regions providing significant reinforcement for the diffusion.

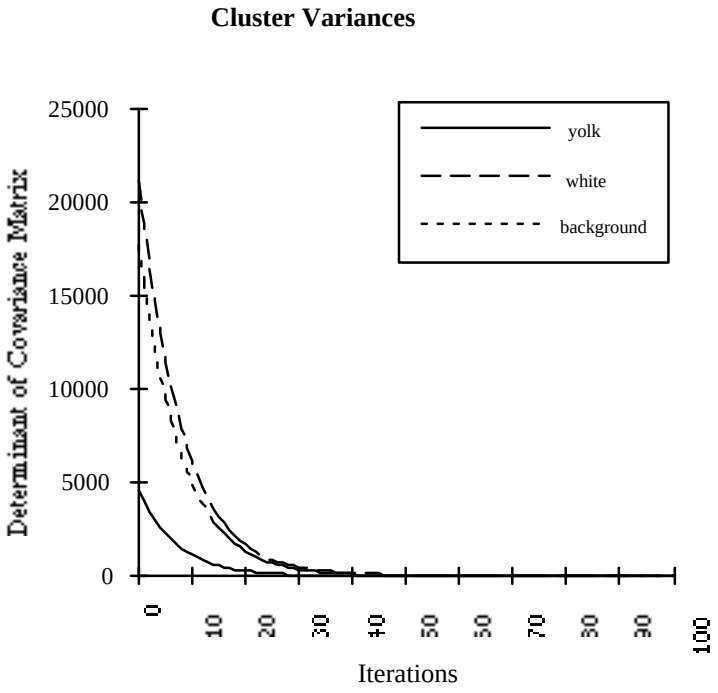

Figure 5 - Tracking the variance of the classifier distributions through 100 iterations of VCD

7.2 Watching the Means

This mechanism does not guarantee that individual pixel values remain monotonic with successive iterations; however, in order for regions to be properly classified, it is essential for the mean values of every cluster to remain relatively fixed within the feature space. Each of the weighting functions is balanced, providing for constant energy across the image throughout the VCD process.

Figure 6 shows a plot of the mean values for each of the three clusters through several iterations of VCD on the computer generated egg phantom. Note that the cluster centers do not vary significantly between iterations.

7.3 Region Boundary Pixels

Bayesian classification combined with VCD systems often have the property that the enhanced boundary gradients arising from the diffusion lead to binary classification, often ignoring partial voluming effects. The statistically based mechanism described here is sensitive enough to provide superior probability information regarding the nature of region boundary pixels.

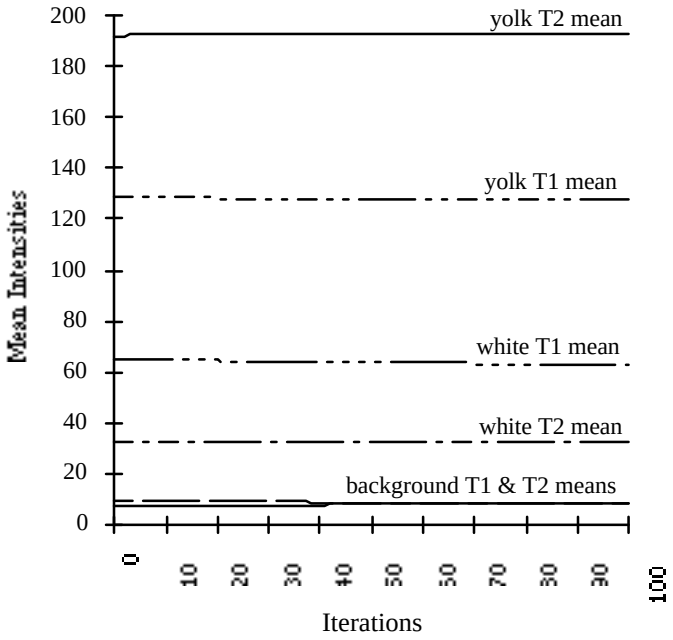

Figure 6 - Plots of the mean values for the pre-classification pixel information during 100 iterations of VCD on the computer generated egg phantom

In essence, the rate at which the process diffuses pixel values is based on the narrowness of the distributions of the pre-selected classifying pixels. If the user supplies a training pattern that has only representative values, then the variance for that cluster will begin low indicating low conductance for similar pixels. Boundary pixels appear relatively dissimilar when variance is low; subsequently, they do not migrate toward the mean value.

If a sharper classification is desired, however, it is only necessary for the user to widen the initial variance of the classification distribution by selecting pixels that are only marginally representative of the tissue class. This will widen the distributions specified by the training class initially, causing border pixels to begin diffusing toward a classification mean reinforced by the classification strength of their neighboring pixels. As VCD proceeds, however, continued narrowing of the classifications provides the appropriate reduction of interpixel conductance.

8 Future Work

Limiting the use of statistical information to local neighborhoods in controlling diffusion may make possible an adaptive filter that will perform segmentation even in the presence of non-linear distortions of signal intensity. "Cold" spots or sensitivity depressions are the subject of much discussion over MR data, and make traditional Bayesian classification difficult. In order to best adapt VCD systems to MR data, a study of the statistical nature of multivalued MR data, or more precisely, a study of the correlation between MR echoes of a single data set is intended.

Using this information, we propose to construct diffusion systems that depend on local statistical information, rather than traditional global cluster information. The problems of small sample size are expected to be the most significant obstacle in this proposed direction. The goal is to enable VCD filtering in regions of low sensitivity, intensity, or contrast, even as they cross sensitivity distortion boundaries.

9 Conclusions

We have demonstrated a means of setting relaxation control parameters for variable conductance diffusion by allowing the user to specify tissue types of interest in medical images. The resulting images show significant stability with regions of interest converging on a limited range of intensities. Border pixels of homogeneous regions show pronounced stability. They do not leak into surrounding regions.

A weakness of this method is that it continues to depend upon a relaxation/annealing schedule that is arbitrarily set by the user. We continue to explore methods for determining better means of controlling of relaxation parameters to work toward VCD in the presence of non-linear distortions of image intensity.

Acknowledgments

We'd like to thank Ross T. Whitaker, Stephen M. Pizer, Bryan Morse, and Penny Rheingans for their inspiration, contributions, and support for this research.

Funding for this research has been provided in part by the Medical Image Presentation Project (**NIH** Grant No. PO1 CA 47982) and the Science and Technology Center for Computer Graphics and Scientific Visualization (**NSF** cooperative agreement no. ASC-8920219).

References

[Duda 73] Duda, Richard O. and Peter E. Hart, *Pattern Classification and Scene Analysis*, John Wiley & Sons, New York, ©1973.

[Gerig 91a] Gerig, Guido, John Martin, Olaf Kubler, Ron Kikinis, Martha Shenton and Ferenc A. Jolesz, “Automating Segmentation of Dual-Echo MR Head Data,” *Information Processing in Medical Imaging* (Lecture Notes in Computer Science **511**, A.C.F Colchester and D. J. Hawkes, *eds.*): 175-187, Springer-Verlag, Berlin, ©1991.

[Gerig 91b] Gerig, Guido, Olaf Kubler, Ron Kikinis and Ferenc A. Jolesz, "Nonlinear Anisotropic Filtering of MRI Data," *Department of Radiology, Brigham and Women's Hospital, Technical Report, BIWI-TR-124* , (Harvard Medical School, June 1991) *submitted to IEEE TMI*.

[Gerig 90] Gerig, Guido, Ron Kikinis and Ferenc A. Jolesz, "Image Processing of Routine Spin-Echo MR Images to Enhance Vascular Structures: Comparison with MR Angiography" 3D Imaging in Medicine (NATO ASI Series F: Vol. 60): 121-132, Springer-Verlag, Berlin, ©1990.

[Kubler 90] Kubler, Olaf, and Guido Gerig, "Segmentation and Analysis of MultiDimensional Data-Sets in Medicine" 3D Imaging in Medicine (NATO ASI Series F: Vol. 60): 63-82, Springer-Verlag, Berlin, ©1990.

[Perona 88] Perona, Pietro and Jitendra Malik, "Scale-Space and Edge Detection Using Anisotropic Diffusion," *Computer Science Division (EECS), UC-Berkeley, Report Number: UCB/CSD 88/483*, (UC-Berkeley, December 1988).

[ter Haar Romeny 91a] ter Haar Romeny, Bart M., and Luc M. Florack, "A Multiscale Geometric Approach to Human Vision," *Perception of Visual Information* (B. Hendee and P.N.T. Wells, *eds.*): Springer-Verlag, Berlin, ©1991.

[ter Haar Romeny 91b] ter Haar Romeny, Bart M., Luc M. Florack, Jan J. Koenderink, and Max A. Viergever, "Scale Space: Its Natural Operators and Differential Invariants," *Information Processing in Medical Imaging* (Lecture Notes in Computer Science **511**, A.C.F Colchester and D. J. Hawkes, *eds.*): 239-255, Springer-Verlag, Berlin, ©1991.

[Whitaker 91a] Whitaker, Ross T. and Stephen M. Pizer, "A Multi-scale Approach to Nonuniform Diffusion," *Department of Computer Science, UNC-Chapel Hill, Technical Report TR91-040* , (UNC-CH, September 1991).

[Whitaker 91b] Whitaker, Ross T., "Geometry Limited Diffusion in the Characterization of Geometric Patches in Images," *Department of Computer Science, UNC-Chapel Hill, Technical Report, TR91-039* , (UNC-CH, September 1991).