

A Graph Cut Approach to Image Segmentation in Tensor Space

James Malcolm Yogesh Rathi Allen Tannenbaum
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, Georgia 30332-0250
{malcolm,yogesh.rathi,tannenba}@bme.gatech.edu

Abstract

This paper proposes a novel method to apply the standard graph cut technique to segmenting multi-modal tensor valued images. The Riemannian nature of the tensor space is explicitly taken into account by first mapping the data to a Euclidean space where non-parametric kernel density estimates of the regional distributions may be calculated from user initialized regions. These distributions are then used as regional priors in calculating graph edge weights. Hence this approach utilizes the true variation of the tensor data by respecting its Riemannian structure in calculating distances when forming probability distributions. Further, the non-parametric model generalizes to arbitrary tensor distribution unlike the Gaussian assumption made in previous works. Casting the segmentation problem in a graph cut framework yields a segmentation robust with respect to initialization on the data tested.

1. Introduction

Image segmentation is often described as the process of separating an image into regions of interest—object and background—usually guided by regional statistics involving image values. However, sometimes directly computing such statistics on image values is not enough to discriminate regions. In a number of important cases, the image may be transformed into a more information rich form to yield latent discriminating features.

For example, consider the image in Figure 1. Here the intensity profile of the lizard’s scales and that of the rock are not distinct enough for intensity-based segmentation techniques to discriminate around the lizard’s throat. In such situations, texture is often a more appropriate discriminant. The structure tensor [3, 10] has been introduced for such texture analysis as a fast local computation providing a measure of the presence of edges and their orientation.




Figure 1. Lizard scales and rock are of similar intensity: initialization (top), segmentation by intensity [7] (left) and texture tensor (right)

In other cases, the image is already in tensor form. For example, diffusion tensor MRI may be represented in this manner from the direction of water diffusion at each pixel. In this case, brain structures such as nerve bundles comprise regions of similarly oriented tensors as water diffuses along the fibers. See [2, 4, 13] for more details.

Various tensor segmentation methods have been proposed including the recent work focusing on variational techniques such as active contours. The basic technique involves iteratively minimizing an energy defined over the statistics of the regions, with lower energies corresponding to better separated regions. A standard variational approach is to approximate the regions as having Gaussian distributions. The energy is defined as some measure of the similarity between two Gaussian distributions, e.g. the log-likelihood ratio, and segmentation proceeds to separate those distributions with active contours. For example, the work of [17] uses a Gaussian approximation for the texture feature channels and a non-parametric histogram for the added image intensity channel. When computing the

tensor mean, different authors employ various norms. For example, in computing the mean tensor for each region, the authors of [17, 18] use the scalar average of each element of the tensor, while [20] uses the Frobenius norm, and [19, 12] use the symmetric Kullback-Leibler divergence.

Taking into account the Riemannian structure of the space of tensors, several groups have begun to compute the statistics and define the distances on the given manifold. In particular, the authors of [13] frame the problem as that of separating Gaussian distributions yet demonstrate the utility of a Riemannian distance between tensors.

Graph cut techniques have received considerable attention as robust methods for image segmentation. Popularized by [7, 16, 5], graph cuts have found applications throughout the vision community mainly for their ability to find globally optimal solutions in many energy minimization formulations. See [6] for a history and survey of the field. Recently the technique has been demonstrated with diffusion tensor MRI [22]. Here, the object and background regional terms (7) are taken to be the average distance to the respective seed points using a divergence measure such as that used in [19]. Because distances are weighted equally, this approach is equivalent to segmenting by a mean tensor. Also, since the weights for each point are calculated as the average distance to each seed point, the region weight computations become prohibitive for large seed sets or large images since the distance metric must be calculated $O(n|\mathcal{O}|)$ and $O(n|\mathcal{B}|)$ times for the object and background weights where n is the number of pixels in the image and $|\mathcal{O}|$ and $|\mathcal{B}|$ denote the sizes of the respective seed regions. In our work, using the Fast Gauss Transform, we can reduce the computational complexity in forming the region distributions to $O(|\mathcal{O}|\log(|\mathcal{O}|))$ and $O(|\mathcal{B}|\log(|\mathcal{B}|))$, a significant decrease since the seed regions $|\mathcal{O}|$ and $|\mathcal{B}|$ may comprise a sizable fraction of the total image size n .

All of the tensor segmentation methods described above rely on the assumption that the regions approximate normal distributions. We propose to generalize these methods in order to be able to segment any image wherein the object and background can be multi-modal. This is done by taking into account the Riemannian geometry of the tensor space. To the best of our knowledge, this is the first time the graph cut technique has been applied for multi-modal tensor segmentation.

The rest of this note is organized as follows. Section 2 describes the Riemannian nature of the tensor space and the process of mapping to a Euclidean space. Section 3 outlines the standard graph cut segmenta-




Figure 2. Zebra image (top) and corresponding tensor channels \mathcal{I}_x^2 , $\mathcal{I}_x \mathcal{I}_y$, and \mathcal{I}_y^2 (right).

tion framework. Finally, Sections 4 and 5 present our algorithm and results on natural imagery and diffusion tensor MRI imagery.

2. The tensor space

Positive definite symmetric matrices (tensors) are widely used in image processing. These are defined to be $n \times n$ symmetric matrices A such that $x^T A x > 0$, for all nonzero $x \in \mathbb{R}^n$. Two typical applications are to capture structural information of an image (structure tensor) [3] and to characterize the diffusion of water molecules in diffusion tensor MRI [2, 4]. A structure tensor is used to extract important features (e.g., edges, corners, texture informational, etc.) from an image. The classic structure tensor uses the tensor product of the smoothed image gradient to form the tensor:

$$\mathbf{T} = K_\rho * [\nabla \mathcal{I} \nabla \mathcal{I}]^T = \begin{pmatrix} K_\rho * \mathcal{I}_x^2 & K_\rho * \mathcal{I}_x \mathcal{I}_y \\ K_\rho * \mathcal{I}_x \mathcal{I}_y & K_\rho * \mathcal{I}_y^2 \end{pmatrix} \quad (1)$$

where \mathcal{I} is a scalar valued image, K_ρ is a Gaussian kernel with standard deviation ρ , and the subscripts denote partial derivatives. Figure 2 gives an example of three channels of an image. For vector-valued images, e.g., color images, the structure tensor may be formed by summing along the color channels:

$$\mathbf{T} = K_\rho * \sum_{i=1}^N (\nabla \mathcal{I} \nabla \mathcal{I}^T) \quad (2)$$

where N is the number of color channels [25]. In general, augmenting the feature vector improves segmentation by increasing the information available. For example, [17] included intensity information with the image derivatives:

$$\mathbf{T} = K_\rho * [\mathcal{I} \mathcal{I}_x \mathcal{I}_y]^T [\mathcal{I} \mathcal{I}_x \mathcal{I}_y]. \quad (3)$$

This has the advantage of defaulting to intensity-based segmentation when texture information is not discrimi-




Figure 3. *Left*: 2D cone demonstrating distances on the tensor space. *Right*: Distribution of structure tensor in \mathbb{R}^3 for a typical image.

nating. To address the dislocation of edges from Gaussian smoothing, different methods of nonlinear smoothing have been investigated to retain edges [9, 21].

Comprehensive treatment of tensor manifolds can be found in [11, 14]. In this work, we briefly describe the structure of the tensor space and the required mapping to a Euclidean tangent space. For the sake of brevity, our treatment will be neither rigorous nor complete, but we want to at least outline some of the key ideas.

It well-known that the space of tensors is not a vector space, but instead forms a Riemannian manifold \mathcal{M} , more specifically, \mathcal{M} forms a convex cone. Figure 3 plots the tensors in 3D. Notice the cone-like structure of the projection. The geodesic distance between two points travels along the surface of the cone, not across the empty center.

A number of past works such as [20, 18] have used a Euclidean metric such as the Frobenius norm to compute distances between two tensor matrices. This can produce erroneous results when considering the influence of points on either side of the cone. Consider the example in Figure 3. Points A , B , and C lie on the surface of the 2D cone with C at the apex and A and B on opposite sides. Furthermore, assume that the Euclidean distance is given by $d_E(A, B) = d_E(A, C) = d_E(B, C) = d$, that is, they are equidistant in the Euclidean sense. However, the Riemannian geodesic from A to B must travel along the surface of the cone and so pass through C , yielding a distance of $d_R(A, B) = d_E(A, C) + d_E(C, B) = 2d$. Therefore, using the Euclidean metric can produce wrong estimates of the tensor statistics (e.g., mean, variance, probability distribution, etc.) which form the basis of many active contour and graph cut segmentation algorithms. In this work, we propose to account for the Riemannian geometry of the tensor manifold when computing the probability distributions used in segmentation.

A Riemannian metric on a manifold \mathcal{M} smoothly assigns to each point $x \in \mathcal{M}$ an inner product on $T_x\mathcal{M}$, the tangent space of \mathcal{M} at x [11, 14]. For our case, the symmetric positive definite tensors naturally define an inner product at each point.

Two maps are defined for mapping points between the manifold \mathcal{M} and a tangent plane $T_p\mathcal{M}$ defined at point $p \in \mathcal{M}$. The first is the *exponential map* $\exp_p : T_p\mathcal{M} \rightarrow \mathcal{M}$, defined on the whole tangent space $T_p\mathcal{M}$. This map is globally one-to-one for the case of tensors, and the tangent space is the space of all symmetric matrices which is isometric to the Euclidean vector space. The second is a unique inverse map called the *log map* $\text{Log}_p : \mathcal{M} \rightarrow T_p\mathcal{M}$ that maps any point $x \in \mathcal{M}$ to the unique tangent vector at p that is the initial velocity of the unique geodesic γ from $\gamma(0) = p$ to $\gamma(1) = x$.

Note that it is this second map, the Log_p map, that is of interest in this work in order to map the Riemannian tensor manifold to a Euclidean tangent plane. Given points $p, \Lambda \in \mathcal{M}$, the log map at p can be computed as follows:

$$\text{Log}_p(\Lambda) = p^{\frac{1}{2}} \log(\Sigma) p^{\frac{1}{2}}, \quad \Sigma = p^{-\frac{1}{2}} \Lambda p^{-\frac{1}{2}}. \quad (4)$$

The geodesic distance between points $p, \Lambda \in \mathcal{M}$ can be computed using the following expression:

$$d^2(p, \Lambda) = \sum_{i=1}^N (\log(\sigma_i))^2, \quad (5)$$

where σ_i are the eigenvalues of Σ defined in (4) above.

By definition, the *intrinsic mean of a random variable* in an arbitrary metric space is the point that minimizes the expected value of the sum-of-squared distance function. As has been pointed out in [11, 14], there is no closed form expression for computing the intrinsic mean of n points lying on the Riemannian manifold \mathcal{M} ; however, due to the non-positive curvature there does exist a unique mean μ found via fixed point iteration with appropriately chosen time step dt :

$$\mu_{t+1} = \mu_t^{\frac{1}{2}} \exp \left(\frac{-dt}{n} \sum_{i=1}^n \log \left(\mu_t^{-\frac{1}{2}} p_i \mu_t^{-\frac{1}{2}} \right) \right) \mu_t^{\frac{1}{2}}.$$

See [11, 14] for a detailed treatment of the subject.

3. Graph cuts

In this section, we very briefly outline the graph cut methodology; see [7, 6, 16, 5] and the references therein. Taking advantage of efficient algorithms for global min-cut solutions, we cast the energy-based image segmentation problem in a graph structure of which the min-cut corresponds to a globally optimal segmentation.

Evaluated for a pixel object/background assignment A , such energies are designed as a data dependent term and a smoothness term. The data dependent term

evaluates the penalty for assigning a particular pixel to a given region. The smoothness term evaluates the penalty for assigning two neighboring pixels to different regions, i.e. a boundary discontinuity. These two terms are thought of as a region-based term and a boundary term, often weighted by $\lambda \geq 0$ for relative influence:

$$E(A) = \sum_{p \in \mathcal{I}} R_p(A_p) + \lambda \sum_{\substack{(p,q) \in \mathcal{N} \\ A_p \neq A_q}} B_{(p,q)} \quad (6)$$

where \mathcal{I} represents all image pixels, \mathcal{N} all unordered neighborhood pixel pairs.

To construct the graph representing this energy, each pixel is considered as a graph node in addition to two nodes representing object and background. The data dependent term is realized by connecting each pixel to both the object and background nodes with non-negative edge weights $R_p(\mathcal{O})$ and $R_p(\mathcal{B})$ representing the likelihood of object and background region presence at pixel p . Lastly, the smoothness term is realized by connecting each pairwise combination of neighboring pixels (p, q) with a non-negative edge weight determined by a penalty for boundary discontinuity, $B_{(p,q)}$. Notice that, since the min-cut sums only along the boundary, the boundary condition of $A_p \neq A_q$ in (6) may be ignored and every pair of neighboring pixels may be connected with edge weight $B_{(p,q)}$. The min-cut of the weighted graph represents the segmentation that best separates the object from its background. See [7] for more details.

Typical applications of graph cuts to image segmentation differ only in the definitions of R_p and $B_{(p,q)}$. For example, in [7] the negative log-likelihood of a pixel's fit into user-initialized intensity histograms is used in the regional term while intensity contrast is used in the boundary term:

$$R_p(\mathcal{O}) = -\ln P(\mathcal{I}_p|\mathcal{O}), \quad R_p(\mathcal{B}) = -\ln P(\mathcal{I}_p|\mathcal{B}), \quad (7)$$

$$B_{(p,q)} = \exp\left(\frac{-\|\mathcal{I}_p - \mathcal{I}_q\|^2}{2\sigma^2}\right) \frac{1}{\|p - q\|} \quad (8)$$

where σ is a user-defined parameter and $\|p - q\|$ is the standard L_2 Euclidean norm yielding pixel distance.

4. The proposed algorithm

The algorithm comprises a preprocessing step followed by probability distribution computation and finally segmentation. Gray-scale and color images are converted to tensor versions for a chosen feature vector. This paper uses the structure tensors defined in (1-3).

The preprocessing step entails mapping the data from the tensor space to a Euclidean tangent space via

the Log_p map (4). This requires choosing a suitable tangent space on which to map. In this work, the data was mapped onto the tangent space of the mean tensor μ (6), although any arbitrary tensor from the data yields suitable results in practice. A special case is given by mapping to the identity, a symmetric positive definite matrix lying on the manifold. This is equivalent to using the Log-Euclidean metric as proposed in [1] (see Figure 9).

From the user-initialized object and background seed regions, we can calculate the multi-dimensional tensor distributions using kernel density estimation [23]. Note that since symmetric matrices are isomorphic to the vector space $\mathbb{R}^{n(n+1)/2}$, the 2x2 structure tensor will only require a 3-dimensional probability space (top left, top right, bottom right). Kernel density estimation can be performed with the Improved Fast Gauss Transform [23, 15]. The basic process carried out by this algorithm is to first, for either a chosen bandwidth or a heuristically computed one, determine appropriate bins for quantization of the space. We set the bandwidth to be

$$h = \frac{\sqrt{2D}}{5} \gamma \quad (9)$$

where D is the dimension of the data (e.g. $D = 3$ for 2x2 tensors) and γ is a scaling parameter. Each sample point then contributes to bins in proportion to distance from the bin, a standard Euclidean vector norm being valid now that the data is mapped to a Euclidean space.

Graph cut segmentation proceeds much like the case of scalar intensity segmentation [7] where these tensor distributions are used as regional priors. So we construct the graph as described in Section 3, and boundary weights (8) are computed using a standard L_2 vector norm on the tensor vectors. The region term (7) evaluates the appropriate distribution for the tensor at each point in the image. Hard constraints are optionally incorporated by manually setting weights, e.g. the object link of the user-initialized object seed point is set to infinity while its background link is set to zero forcing an object labeling. The min-cut of this graph yields a globally optimal segmentation under these constraints.

5. Results

We performed various segmentations using the above methodology. To compare against using intensity alone, we show segmentations using just the mean intensity [24] and the full intensity [7]. Tensor data was formed using (1) unless specified otherwise. To demonstrate the power of the technique, several examples are shown that fail to segment with the assump-




Figure 4. Lizard showing successive integration statistical information.




Figure 5. Butterfly with multi-modal wing pattern and thin antenna: initial, Euclidean metric, Riemannian metric (*left to right*).

tion of a Euclidean space, i.e. the Log_p mapping is not performed.

For our experiments, γ in (9) ranged from $\gamma = 0.4$ to $\gamma = 0.7$. The λ controlling smoothness in (6) stayed fixed at 1 in all experiments except the diffusion tensor MRI where it was reduced to 0.1 to capture sharp features. The ρ used to form the tensors in (1-3) was held at $\rho = 1.5$ in all experiments except the zebra (Figure 10) where it was doubled to $\rho = 3$ for smoothing the included intensity values. The σ from (8) was adaptively set during segmentation to be the square root of the average squared norm used in (8):

$$\sigma = \sqrt{\frac{1}{|\mathcal{N}|} \sum_{(p,q) \in \mathcal{N}} \|T_p - T_q\|^2}$$

Neighborhoods of size 16 were used to yield smoother segmentations while neighborhoods of size 4 or 8 were seen to increase metrication artifacts [8]. Standard hard constraints were used as in [7, 22]. As there may be outlying regions of likely object, automatic post-processing is sometimes needed to take the region indicated by the object seed points.

The first experiment is that of a lizard in Figure 1, the scales of which are patterned yet of intensity similar to the rock background. Segmenting using the full intensity histogram as a prior [7] fails to capture the scales around the throat and fails to capture edges correctly in many places along the border of the lizard. Using the full histogram does have the advantage of

being able to grab more of the black belly of the lizard compared to the tensor segmentation. Notice that, not only is the intensity distribution multi-modal, the tensor distributions are multi-modal. Existing methods that assume a Gaussian model will fail to segment. Figure 4(b) additionally shows segmentation using only mean intensity graph cuts [24] where the affinity for separating light and dark regions is seen.

The butterfly in Figure 5 has patterned wings and antennae distinct from the leaves in the background. Notice that if we assume the space Euclidean, the segmentation is unable to pick up the entire wing nor the antennae. This is despite the fact that a multi-modal tensor distribution was used with the Euclidean assumption. However, using the Riemannian metric gives the desired segmentation result.

In Figure 6 the feathers of the duck and the ripples in the surrounding water have very similar intensity distributions, hence tensor information is necessary. The ground-truth histogram is shown in the top right which explains the failure to segment by intensity alone [7]. Using the tensor information however, provides the features necessary to segment the duck from the water background.

The monkey in Figure 7 has a fuzzy texture compared to the surrounding leaves. Notice the white eye has the same fuzzy texture as the head and back yet is not picked up when assuming a Euclidean space while good results are shown with a Riemannian metric.

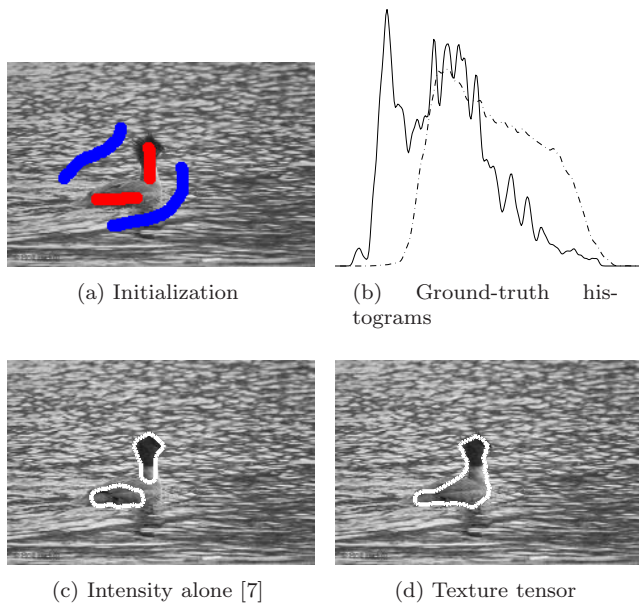


Figure 6. Duck feathers and water have overlapping intensities requiring tensor information to correctly segment.




Figure 7. Monkey and leaves are of similar intensity: initial (*top*), Euclidean metric (*left*), Riemannian metric (*right*)

The color fish in Figure 8 demonstrates incorporating additional channels into the tensor (2). Notice that, if the space is assumed Euclidean, even the yellow




Figure 8. Color fish using tensor formed across channels as in (2): initialization (*top*), Euclidean metric (*left*), Riemannian metric (*right*).




Figure 9. Color tiger using tensor formed across channels as in (2): initialization (*top*), Log-Euclidean metric (*left*), Riemannian metric (*right*). Mapping to the tangent plane of the identity matrix induces the Log-Euclidean metric and yields results similar to the Riemannian mapping to the tangent plane of the mean tensor.

scales on the underbelly are left out.

The color tiger in Figure 9 has very similar textural features to those of its background. The identity matrix is a symmetric positive definite matrix and so mapping to its tangent plane instead of the mean tensor is valid. This induces a Log-Euclidean metric space which yields results similar to the Riemannian metric space and has avoided computation of the mean tensor. The bandwidth (9) for the Log-Euclidean case was set with $\gamma = 0.2$.

To compare with the previous work of [17] we show an example of segmenting a zebra in Figure 10. It was shown in [17] that, even assuming the space to be Euclidean, segmentation proceeds correctly with active




Figure 10. Zebra using the tensor defined in (3): initialization (*left*), Euclidean metric (*middle*), Riemannian metric (*right*). As expected, segmentation works with or without mapping to the tangent plane [17].




Figure 11. Zebra with initialization failing to capture top of back. Figure 10 shows user-added seed points that resulted in correct segmentation..

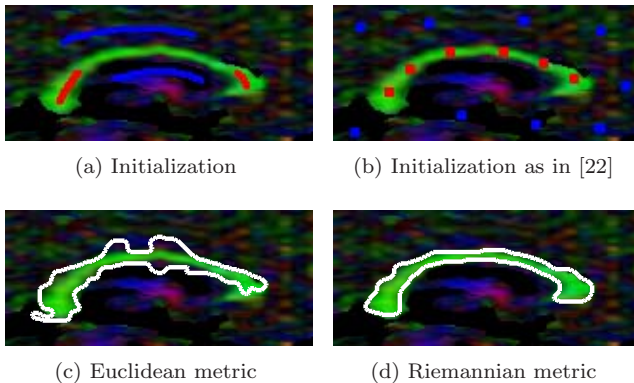


Figure 12. Diffusion tensor MRI of the brain with corpus callosum segmented. Both initializations gave comparable results. The color-coded fractional anisotropy image is shown for visualization while segmentation was performed on the underlying tensor data.

contours. Figure 11 shows a naive initialization that failed to capture the top of the zebra’s back. Figure 10 shows that additional seed points along the back led to correct segmentation.

In Figure 12, the proposed method was also tested against diffusion tensor MRI of the brain with two different initializations—a few brush strokes or a set of points as in [22]—both yielding comparable segmentations. In order to view the underlying tensor data, the color coded fractional anisotropy image was gen-

erated, yet segmentation was performed on the underlying tensor data. Despite the corpus callosum comprising a unimodal tensor distribution, Figure 12(c) shows a poor segmentation under the Euclidean assumption. Further, under the Euclidean metric, parameters were highly sensitive and the segmentation shown is the best we could obtain. Segmenting after mapping (Figure 12(d)) was robust to parameter selection. In this experiment, $\lambda = 0.1$ to capture the thin structure and sharp corners.

We should note finally that we employed the publicly available implementations of the Improved Fast Gauss Transform¹ and Max-Flow graph cut², in our experiments.

6. Conclusion

In this paper, we proposed a new method to perform segmentation of tensor valued images taking into account the natural Riemannian structure of the tensor manifold. Segmentation was performed via standard interactive graph cuts. Examples show that without taking into account the Riemannian structure, tensor distributions fail to capture the true variation of object and background.

The algorithm is summarized as follows. After forming the tensor data, we map the data to a Euclidean space where distances can be taken via a vector norm. From user-initialized object and background seed regions, multidimensional distributions are computed. Similar to the work of [7], these distributions are then used as priors in calculating edge weights for subsequent segmentation via graph cut energy minimization.

Future work will explore scale invariant features; see [17]. Indeed, the method may fail when two textures differ only in scale. Future work on diffusion tensor MRI will require quantitative comparison with other techniques and hand-segmented ground truth. In such a study, care must be taken to define the ini-

¹<http://www.umiacs.umd.edu/vikas/>

²<http://www.adastral.ucl.ac.uk/vladkolm/>

tial prior distributions equivalently across techniques. Also, since the graph cut technique yields the global optima under the constraints of user initialization, it may be difficult to compare against methods sensitive to local minima such as active contours.

Acknowledgements

This work was supported in part by grants from NSF, AFOSR, ARO, MURI, MRI-HEL as well as by a grant from NIH (NAC P41 RR-13218) through Brigham and Women's Hospital. This work is part of the National Alliance for Medical Image Computing (NAMIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149. Information on the National Centers for Biomedical Computing can be obtained from <http://nihroadmap.nih.gov/bioinformatics>.

References

- [1] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Log-euclidean metrics for fast and simple calculus on diffusion tensors. *J. of Magnetic Resonance in Medicine*, 56(2):411–421, 2006.
- [2] P. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.
- [3] J. Bigun, G. Granlund, and J. Wiklund. Multidimensional orientation estimation with applications to texture analysis and optical flow. *PAMI*, 13(8):775–790, 1991.
- [4] D. L. Bihan, J. Mangin, C. Poupon, C. Clark, S. Pappata, N. Molko, and H. Chabriat. Diffusion tensor imaging: Concepts and applications. *J. of Magnetic Resonance Imaging*, 13:534–546, 2001.
- [5] A. Blake, C. Rother, M. Brown, and P. Torr. Interactive image segmentation using an adaptive GMMRF model. In *ECCV*, volume 3021, pages 428–441, 2004.
- [6] Y. Boykov and G. Funka-Lea. Graph cuts and efficient N-D image segmentation. *IJCV*, 70:109–131, 2006.
- [7] Y. Boykov and M. Jolly. *Interactive graph cuts* for optimal boundary and region segmentation of objects in N-D images. In *ICIP*, volume 1, pages 105–112, 2001.
- [8] Y. Boykov and V. Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. In *ICCV*, pages 26–33, 2003.
- [9] T. Brox and J. Weickert. Nonlinear matrix diffusion for optic flow estimation. *Pattern Recognition*, 2449:446–453, 2002.
- [10] R. de Luis-Garcia, M. Rousson, R. Deriche, and C. Alberola-Lopez. Tensor processing for texture and color segmentation. In *Scandinavian Conference on Image Analysis*, pages 1117–1127, 2005.
- [11] P. Fletcher and S. Joshi. Principle geodesic analysis on symmetric spaces. In *ECCV, Workshops CVAMIA-MMBIA*, pages 87–98, 2004.
- [12] C. Lenglet, M. Rousson, and R. Deriche. Segmentation of 3D probability density fields by surface evolution: Application to diffusion MRI. In *MICCAI*, pages 18–25, 2004.
- [13] C. Lenglet, M. Rousson, R. Deriche, and O. Faugeras. Statistics on the manifold of multivariate normal distributions: Theory and application to diffusion tensor MRI processing. *J. of Mathematical Imaging and Vision*, 2006.
- [14] X. Pennec, P. Fillard, and N. Ayache. A riemannian framework for tensor computing. *IJCV*, 66:41–66, 2006.
- [15] V. Raykar and R. Duraiswami. The improved fast gauss transform with applications to machine learning. In L. Bottou, O. Chapelle, D. Decoste, and J. Weston, editors, *Large Scale Kernel Machines*. MIT Press, 2006. to appear.
- [16] C. Rother, V. Kolmogorov, and A. Blake. GrabCut: Interactive foreground extraction using iterated graph cuts. In *ACM Trans. on Graphics (SIGGRAPH)*, 2004.
- [17] M. Rousson, T. Brox, and R. Deriche. Active unsupervised texture segmentation on a diffusion based feature space. In *CVPR*, pages 699–704, 2003.
- [18] M. Rousson, C. Lenglet, and R. Deriche. Level set and region based propagation for diffusion tensor MRI segmentation. In *ECCV, Workshops CVAMIA-MMBIA*, pages 123–134, 2004.
- [19] Z. Wang and B. Vemuri. An affine invariant tensor dissimilarity measure and its application to tensor-valued image segmentation. In *CVPR*, 2004.
- [20] Z. Wang and B. Vemuri. Tensor field segmentation using region based active contour model. In *ECCV*, volume 3024, 2004.
- [21] J. Weickert and T. Brox. Diffusion and regularization of vector and matrix-valued images. In M. Nashed and O. Scherzer, editors, *Contemporary Mathematics*, volume 313, pages 251–268, 2002.
- [22] Y. Wedeselassie and G. Harmarneh. DT-MRI segmentation using graph cuts. In *Medical Imaging*, volume 6512 of *SPIE Proceedings*, pages 1–9. SPIE, 2007.
- [23] C. Yang, R. Duraiswami, N. Gumerov, and L. Davis. Improved fast gauss transform and efficient kernel density estimation. In *ICCV*, pages 464–471, 2003.
- [24] X. Zeng, W. Chen, and Q. Peng. Efficiently solving the piecewise constant mumford-shah model using graph cuts. Technical report, Dept. of Computer Science, Zhejiang University, P.R. China, 2006.
- [25] S. D. Zenzo. A note on the gradient of a multi-image. *Computer Vision, Graphics, and Image Processing*, 33:116–125, 1986.