

Visualizing Diffusion Tensor Images of the Mouse Spinal Cord

David H. Laidlaw* Eric T. Ahrens David Kremers
Matthew J. Avalos Russell E. Jacobs
California Institute of Technology, Pasadena, CA 91125

Carol Readhead
Cedars-Sinai Medical Center
Los Angeles, CA 90048

Abstract

Within biological systems water molecules undergo continuous stochastic Brownian motion. The rate of this diffusion can give clues to the structure of underlying tissues. In some tissues the rate is anisotropic – faster in some directions than others. Diffusion-rate images are second-order tensor fields and can be calculated from diffusion-weighted magnetic resonance images. A 2D diffusion tensor image (DTI) and an associated anatomical scalar field, created during the tensor calculation, define seven values at each spatial location. Visually representing these images is a challenge because they contain so many inter-related components. We present two new methods for visually representing DTIs. The first method displays an array of ellipsoids where the shape of each ellipsoid represents one tensor value. The novel aspect of this representation is that the ellipsoids are all normalized to approximately the same size so that they can be displayed simultaneously in context. The second method uses concepts from oil painting to represent the seven-valued data with multiple layers of varying brush strokes. Both methods successfully display most or all of the information in DTIs and provide exploratory methods for understanding them. The ellipsoid method has a simpler interpretation and explanation than the painting-motivated method; the painting-motivated method displays more of the information and is easier to read quantitatively. We demonstrate the methods on images of the mouse spinal cord. The visualizations show significant differences between spinal cords from mice suffering from Experimental Allergic Encephalomyelitis (EAE) and spinal cords from wild-type mice. The differences are consistent with differences shown histologically and suggest that our new non-invasive imaging methodology and visualization of the results could have early diagnostic value for neurodegenerative diseases.

CR Categories: H.5 [Information Systems]: Information Interfaces and Presentation; I.3.3 [Computing Methodologies]: Computer Graphics—picture/image generation; I.3.4 [Computing Methodologies]: Computer Graphics—paint systems; I.3.5 [Computing Methodologies]: Computer Graphics—computational geometry and object modeling; I.3.8 [Computing Methodologies]: Computer Graphics—applications; J.2 [Computer Applications]: Physical Sciences and Engineering; J.3 [Computer Applications]: Life and Medical Sciences; J.5 [Computer Applications]: Arts and Humanities—fine arts

Keywords: multi-valued visualization, tensor field visualization, oil painting

*D. Laidlaw is at the Dept. of Computer Science, Box 1910, Brown University, Providence RI, 02912, dhl@cs.brown.edu; E. Ahrens, and R. Jacobs are at the Biological Imaging Center, Beckman Institute 139-74, {eta,rjacobs}@gg.caltech.edu; David Kremers is in the Division of Biology, 156-29, davidkremers@caltech.edu; Matthew J. Avalos is at the Computer Graphics Laboratory 139-74, Department of Computer Science, Division of Engineering and Applied Science, avalos@gg.caltech.edu.

1 Introduction

Within biological systems water molecules undergo continuous stochastic Brownian motion. In different tissues the rate of this diffusive motion can vary by several orders of magnitude – faster in liquids like cerebro-spinal fluid, slower in tissues like muscle. In some tissues the rate is anisotropic, or faster in some directions compared to others. Magnetic resonance imaging (MRI) can acquire images with intensity values sensitive to the diffusion rate of water. A quantitative image of the diffusion rate can be calculated from a set of such MR images [4, 2, 3, 1]. From a 2D slice or 3D volume image of this directionally dependent diffusion rate we can infer underlying tissue structure and better understand the anatomy of the nervous system, neuro-degenerative diseases, and neural development [4].

Diffusion-rate images calculated from MRI measurements are second-order tensor fields. A diffusion tensor image (DTI) and the associated anatomical scalar field, calculated as part of the tensor calculation, define seven related values at each spatial location. Often it is useful to simultaneously display more than a few of these values in order to visualize inter-relationships among them and to understand the data. We present two new methods for visually representing these diffusion tensor images. The first method, shown in Fig. 1, displays an array of ellipsoids where the shape of each ellipsoid represents one tensor value. The novel aspect of this representation is that the ellipsoids are all normalized to approximately the same size so that they can all be displayed in context. The second method, shown in Fig. 2, uses concepts from oil painting to represent the seven-valued data with layers of varying brush strokes. Renderings of data from a diseased mouse are shown for comparison in Figs. 8 and 10. The color plate shows both methods for both a healthy mouse and a diseased mouse. The methods are exploratory – they allow a viewer to explore all of the values comprising a tensor-valued dataset rather than distilling out some simpler derived quantity. Both are implemented as new modules within AVS[23].

In the remainder of the paper we first discuss related work in visualizing second-order tensor fields. We then describe the anatomy of the mouse spinal cord and the pathology of the mouse neurological disease that we use in our example images. Section 4 describes diffusion tensor images acquired using MRI. Section 5 describes the ellipsoid-based method and Section 6 describes the painting-motivated method. We summarize and conclude in Section 7.

2 Related work

References [7, 13, 12, 9, 8] review specific methods for tensor-field visualization and develop several tensor-field visualization methods. Portions of references [19, 10] also give informative reviews of methods. However, these methods are all designed for displaying tensor fields with a different physical structure and interpretation than the diffusion tensor. The methods do not take advantage of the fact that a diffusion tensor has orthogonal eigenvectors and non-negative eigenvalues, and thus devote critical visual bandwidth to tensor attributes not present in diffusion tensors. They also produce visual representations that intuitively represent the physical




Figure 1: Normalized diffusion ellipsoid visualization. Each ellipsoid represents the tensor value at one pixel geometrically to capture most of the information in the tensor-valued image. Fig. 2 shows the same data rendered with our oil-painting-motivated method. Figs. 8 and 10 show both methods used to render data from a diseased mouse. (See color plate.)




Figure 2: Diffusion tensor image (DTI) visualization using concepts from painting. This stroke image is composed from layers as shown in Figs. 9 and 10 and in the color plate. Component data values are represented in a blue-purple underpainting; in the direction, shape, color, and transparency of a layer of elliptical strokes over that; and in the frequency of a texture applied to the elliptical strokes. The resulting image displays all seven inter-related values that comprise a DTI.

phenomena that they are representing, but that do not represent diffusion intuitively.

Diffusion tensor imaging data have primarily been displayed by extracting scalar- or vector-valued components of the tensor at each point and visualizing the resulting scalar fields and vector fields [5, 20]. The primary limitation with this approach is that it is difficult to understand the tensor field from individual component images because relationships among the different components are spread across several images. For example, imagine trying to understand a 2D image of 3D vectors by viewing three grayscale images of the x , y , and z components of the vector. It would be difficult to understand the 3D direction of the vectors without a representation that combines the three values. Another limitation in viewing the components individually is that noise in the acquired images can bias some of the extracted values when they are viewed out of the context of the other values [20].

Arrays of ellipsoids have also been used to represent diffusion tensor data [5, 20]. They have been limited to images of small regions of interest (16×16 pixels) where each image provides an iconic representation of 256 tensor values. Fortunately, the regions displayed have diffusion rates that differ by less than a factor of 10. However, the smaller tensor values are represented by sparsely spaced small ellipsoids, and the connection between the different values is lost. For regions where the diffusion rate is more widely varying this effect would be compounded and the ellipsoids for low-rate tensor values would be all but invisible.

Reference [14] was the first to experiment with painterly effects in computer graphics. More recently reference [18] extended the approach for animation and further refined the use of layers and brush stroke characteristics for creating effective imagery. Both of these efforts were aimed toward creating art, however, and not toward visualizing scientific data. Along similar lines, references [26, 25, 22] used software to create pen and ink illustrations for artistic purposes.

Layering has been used in scientific visualization to show multiple items. In reference [15, 16], transparent stroked textures show surfaces without completely obscuring what is behind them. These results are related to ours, but our application is 2D, and so our layering is not as spatial as it is in the 3D case. Our layering is more in the spirit of oil painting, where layers are used more broadly, often as an organizing principle for a painting.

Much of the inspiration for our approach comes from studying the painting styles of a number of impressionist painters and from consulting books on basic artistic principles and problem solving in painting. We found Van Gogh's style particularly applicable because he used expressive, discrete strokes that are combined to represent a continuous scene. They can be read differently from different distances and can encode different information from each of those viewpoints.

The goal of our work is to visualize scientific information by building on many of these concepts. From the sources above we discovered many useful concepts and rules-of-thumb for creating effective images. Just as when an artist creates a painting, our goal is to communicate complexity in a direct, expressive, and visually pleasing manner. We used these principles as guidelines for creating our visualization methods.

There is a large body of literature on the perceptual effects of texture, color, icons, and other visual representations of scientific data [17, 24, 21]. In our painting-based work, visual choices were based not on this visualization perception literature, but instead on the artistic guidelines. The effects of many of the components of our visualizations undoubtedly are explained in the visualization perception literature. It would be interesting, but beyond the scope of this work, to better connect the literature to the world of art.




Figure 3: Anatomy of a healthy wild-type mouse spinal cord section. The three major regions will take on different appearances in the histological sections below and in our visualizations. The top of the image is dorsal and the bottom ventral.




Figure 4: Histology showing anatomy of a healthy wild-type mouse spinal cord section. The section has been stained with toluidine blue; dark regions contain myelin and correlate with regions of white matter.

3 Anatomy and pathology of examples

Our visualization examples are taken from images of the spinal cords of healthy and diseased mice. The diseased mice are transgenic and suffer from Experimental Allergic Encephalomyelitis (EAE), a disease where the mouse immune system attacks the central nervous system [11]. Many aspects of the pathology of EAE mimic multiple sclerosis in humans. The disease causes a stripping away of the insulating myelin sheath (demyelination) of axons that comprise the white matter of the spinal cord. A large amount of inflammation is also observed. EAE causes behavioral symptoms that progress in well-characterized stages: 1) limp tail, 2) uneven gait, 3) hindlimb paralysis, 4) fore and hindquarter paralysis, and 5) moribundity and death.

Fig. 3 shows the anatomy of a spinal cord section in a healthy wild-type mouse. Fig. 4 shows a histological section that has been stained with toluidine blue that labels myelin with a dark blue color. The light-blue butterfly-shaped region in the center of the spinal cord consists of gray matter containing mostly neuronal cell bodies. The darker-blue cardioid-shaped region surrounding the gray matter is white matter. The myelinated axons in this region connect the brain and other levels of the spinal cord. The cardioid shape is the boundary of the spinal cord. Outside of the cord is a layer of spinal nerves that connect the central nervous system to other parts of the body.




Figure 5: Matrix representation of a diffusion tensor image. All of the scalar elements are shown. The upper left image represents the upper left matrix element for each pixel. Understanding the tensor field from this collection of scalar-valued images is not intuitive.

4 Diffusion tensor images

The diffusion tensor images used in our visualizations were obtained using novel MRI and numerical methods described in [1]. The data are in the form of a 2D slice image with a second-order tensor value at each pixel. Each second-order tensor value can be represented by a 3×3 symmetric matrix with non-negative eigenvalues. It can also be represented geometrically by an ellipsoid. The matrix and ellipsoid representations are equivalent [6]. The principal axes of an ellipsoid are the eigenvectors (n_1, n_2, n_3) of the matrix and the principal radii are the eigenvalues $(\lambda_1, \lambda_2, \lambda_3)$ of the matrix. The shape and size of an ellipsoid encode key aspects of the diffusion tensor value we wish to visualize. It can be interpreted as the shape that a small dot of ink would take on after some fixed diffusion time. Larger ellipsoids represent faster diffusion, spheres represent isotropic diffusion, and eccentric (prolate or oblate) ellipsoids represent anisotropic diffusion.

In the spinal cord, gray matter has nearly isotropic diffusion whereas white matter has very anisotropic diffusion as observed in MRI measurements [4]. In white matter water diffuses more readily along the axon than across it. The pathology of EAE leads to changes in the diffusion characteristics of the spinal cord, especially in certain regions of white matter. We intuitively show these pathologic changes with the help of our visualization methods.

Each sub-image in Fig. 5 shows one element of the matrix representation of a diffusion tensor image. The matrix elements are difficult to interpret intuitively as a tensor-valued image. One reason for this difficulty is that they are not rotationally invariant; they change if the sample is oriented differently within the experimental apparatus. These images are analogous to images of the scalar components of a vector.

Fig. 6 shows images of the eigenvalues of the same diffusion tensor image. The scalar values in these images are rotationally invariant but display no directional information. These images are analogous to an image of the magnitude of a vector. The directional information that is missing is a significant part of the tensor value that we wish to display.

In addition to the tensor-valued diffusion-rate image, the numerical calculation of the tensor image from MR images also produces a scalar-valued anatomical image as shown in Fig. 7. We name it I_0 . It shows contrast between anatomical regions arising from non-diffusion mechanisms and thus captures additional information.




Figure 6: Images of the scalar eigenvalues of a diffusion tensor image. These values extract important geometric information about the diffusion tensor but do not include the directional context.




Figure 7: An additional scalar-valued image produced as part of the data-modeling process that calculates the diffusion tensor image. This image is similar to a conventional magnetic resonance image and contains additional anatomical information beyond what is contained in the diffusion tensor image.

5 Normalized ellipsoid representation

Our first new visualization method for diffusion tensor images captures most of the tensor information into a single image. It takes advantage of the fact that a tensor value with positive eigenvalues can be represented with an ellipsoid. Our first attempts made a direct mapping from tensor values to ellipsoids, but, because the diffusion rates vary by factors of almost a thousand, so did the sizes of the ellipsoids. The ellipsoids could not all be seen in their appropriate locations with such a large variation in size.

We avoid this extreme variance in tensor values by normalizing each tensor value so that its maximum eigenvalue is 1.0. The resulting images, shown in Figs. 1 and 8 and in the color plate, contain ellipsoids with their largest radii all equal. The planar grid of ellipsoids is lit from top center with a single directional light source and viewed with perspective projection at a small tilt.

• **Observations** Each ellipsoid has a clear spatial location, and their shapes can be interpreted. Isotropic tensor values, e.g., in the gray-matter regions of Fig. 1, are represented as spheres that completely obscure the background color. Anisotropic tensor values, e.g., in white-matter regions, are skinny and smaller. This is because their largest principal radius is the same as the largest radius of the spheres but the other principal radii are much smaller. In these regions the background color shows through more.

Fig. 1 shows an image of a spinal cord from a healthy mouse.




Figure 8: Ellipsoid representation of diseased spinal cord. Note right-left asymmetries in this image. These and other differences between this image and Fig. 1 of a healthy cord show changes due to the disease. Significant changes are particularly noticeable in the ellipsoid shapes in the myelinated regions in the lower right quadrant of the image. (See color plate.)

Fig. 8 shows an image of a spinal cord from a diseased mouse. The color plate shows portions of both images. Differences due to the pathology of the disease are apparent in the ventral white matter. The presence of EAE lesions appears to have reduced the diffusion anisotropy throughout much of the white matter and caused the ellipsoids to become much more spherical than in comparable regions of the healthy cord.

This visualization method shows different information at different distances. From farther away, neighboring ellipsoids visually blend together to create a texture that gives an almost-continuous qualitative impression of the underlying anatomical regions. Close up, the ellipsoids are distinct enough to see individually, and show the relative anisotropy and directional information more quantitatively.

These ellipsoid images encode most of the information from the tensor-valued diffusion rate images. They capture the spatial correlation between the tensor values and the underlying space, the orientation of the tensors, and the relative diffusion rates in each direction. They are missing both the scalar I_0 value and the absolute diffusion rate that was normalized out of the tensor values so that the ellipsoids would be a consistent size.

6 Painting-motivated representations

Our second method applies concepts from oil painting to display diffusion tensor images. We used multiple layers of brush strokes to represent the tensor image and the associated anatomical scalar image, I_0 . The brush strokes reflect the geometric nature of values derived from the tensors and of the relationships among the values. Also, the use of underpainting and saturated complementary colors evokes a sense of depth. Together, these painting concepts help create a visual representation for the data that encodes all of the data in a manner that allows us to explore the data for a more holistic understanding.

The color plate and Fig. 2 both show a completed stroke image.

diffusion data	visualization
anatomical I_0 image	underpainting lightness
voxel size	checkerboard spacing
ratio of largest to smallest eigenvalue ($\frac{\lambda_1}{\lambda_3}$)	stroke length/width ratio and transparency
principal direction (n_{1x}, n_{1y})	stroke direction
principal direction (n_{1z})	stroke red saturation
magnitude of diffusion rate	stroke texture frequency

Table 1: Mapping of data parameters to visualization parameters for diffusion tensor example.

We built the visualization in four layers that are illustrated in Figs 9 and 10. Table 1 summarizes the mapping of data values to image contents. The layers are discussed individually in the following paragraphs.

- **Underpainting** First, we wanted a layer to show the overall form or structure of the anatomy. This is analogous to the function of an underpainting in oil painting. The I_0 image provides this information and has the best signal-to-noise ratio (SNR) of any of our scalar variables. It is encoded in the lightness of the purple layer shown in the left image of Fig. 9.

- **Checkerboard layer** Second, we wanted to incorporate the underlying voxel size into the image. Without it, the image can give a false sense of resolution. The voxel-size representation needed to be subtle and not overwhelm the other information, so we encoded it as a semi-transparent checkerboard pattern layered over the underpainting. The center image in Fig. 9 shows the checkerboard layered over the underpainting.

- **Stroke layer** Third, we wanted to capture geometric information about the diffusion tensor values. Axon tracts in white matter are quite anisotropic while gray matter is close to isotropic, so the degree of anisotropy needed to be prominently displayed. We chose ellipse-shaped strokes with a length-to-width ratio equal




Figure 9: Underlayers of the visualization in Fig. 10. The images, from left to right, show an underpainting layer, a checkerboard layer composited on the underpainting, and a stroke layer composited over that. Both Fig. 10, and the color plate show a fourth and final layer composited over these that captures the diffusion rate in a texture along each stroke.




Figure 10: Diffusion tensor image (DTI) visualization using concepts from painting of a spinal cord from a mouse with EAE. This stroke image is composited from layers shown in Fig. 9. It displays all seven dependent values that comprise a diffusion tensor image. (See color plate.)

to $\frac{\lambda_1}{\lambda_3}$. Nearly circular ellipses represent isotropic regions, while skinny ones represent anisotropic regions. We compared λ_2 and λ_3 and they are always nearly equal within the cord, so we did not need to incorporate λ_2 directly. The ellipse is a natural icon to represent this ratio, since the shape echoes the structure of the underlying physical phenomenon.

We also wanted to represent the principal diffusion direction, n_1 , in this stroke layer. This is the direction of fastest diffusion and shows the direction of fiber tracts. The direction of the projection of n_1 into the image plane is encoded in the direction of the strokes. The out-of-plane component of n_1 is indicated by the saturation of the strokes; the saturated red strokes appear to point more out-of-plane because red on blue elicits a sense of depth. Note that for the more isotropic regions the direction becomes less meaningful, since the eigenvalues are all equal. Strokes in isotropic regions are made more transparent to deemphasize the regions because the data there contain less information and the changes in n_1 have less meaning.

Stroke placement within the image is random, with strokes arranged so that they do not overlap in the direction of their widths but do slightly in the direction of their lengths. Thus adjacent strokes create the impression of flow along their direction, similar to Van Gogh's style in some cases. The flow gives a sense of the directional fibers within the axon tracts. Strokes are not placed where the SNR is insufficient to calculate an accurate tensor. The right image of Fig. 9 shows this stroke layer composited over the two underlying layers.

- **Stroke texture layer** Finally, because the absolute magnitude of the diffusion rate is not correlated with the anisotropy ratio, $\frac{\lambda_1}{\lambda_3}$, we wanted to encode it in the fourth layer. It is less important than the other values, however, so we encoded it in a form that is interpreted less quickly. We represented the rate with the frequency of the texture on the strokes. Faster diffusion is represented with more stripes along a stroke, and slower diffusion with fewer. The higher-frequency textures create smaller features that are often used artistically to represent higher speeds. This intuitively represents the corresponding higher rate of diffusion.

- **Observations** The complete visualization is shown for a healthy cord in Fig. 2 and for a diseased cord in Fig. 10. In the healthy cord the anatomical regions are clearly distinguished. The anisotropy of the white matter and spinal nerves creates regions of narrow, opaque ellipse shapes. The regions of gray matter are shown with rounder, semi-transparent shapes. Unlike the ellipsoid representation, the stroke images distinguish the diffusion rate difference between white matter and spinal nerves. In the spinal nerves, the diffusion rate is faster, and shows in the higher-frequency texture on the strokes. This difference is consistent with the larger axon diameter of the spinal nerves as shown in histological sections. The differences in the anatomical scalar image also show through between the stroke shapes, particularly when the images are viewed from a distance, incorporating its non-diffusion information to distinguish different regions.

The pathology of the disease is apparent in the ventral white matter. The strokes near EAE lesions are much rounder and more transparent than in other white-matter regions, indicating less anisotropy. The stroke images show different information at different distances. From farther away, the underpainting, strokes, and texture all combine to give a qualitative impression of different anatomical regions. From a closer perspective, the texture becomes more apparent, and more subtle anatomical difference appear. At very close perspectives, more quantitative measures of the tensor can be seen – the anisotropy ratio (ellipse aspect ratio), or the absolute diffusion rate (texture frequency). The quantitative information is easier to read quantitatively from the stroke images than from the ellipsoid images because the stroke images are strictly 2D. There is no distortion from the perspective rendering, the projection from 3D to 2D, or the lighting calculation – all of the geometric information lies directly in the plane of the stroke image. Only the portions

of the tensor mapped to less quantitative color and transparency are less easy to quantify.

A stroke image represents all of the information in a diffusion tensor image and in the associated scalar-valued anatomical image. It captures the spatial correlation between the tensor values and the underlying space, the orientation of the tensors, the relative diffusion rates in different direction, and the absolute diffusion rate.

7 Summary and conclusions

We have presented two new methods for visualizing 2D images of second-order tensor fields representing diffusion rates. Our methods create visually rich images that can represent many values at each spatial location in an image. With this richness we have captured most or all of the information present in the data. Through applying the visualization methods we hope to better understand biological systems imaged with diffusion tensor imaging.

The first method represents the tensor values with ellipsoids. It modifies the direct tensor-to-ellipsoid mapping by normalizing the ellipsoid sizes. As a result, the ellipsoids are all approximately the same size and show the tensor structure effectively everywhere within the image.

With our second method we applied concepts borrowed from artists. Varied brush strokes and layering helped to simultaneously display many components of the data. Underpaintings showed form. We used brush strokes both individually, to encode specific values, and collectively, to show spatial connections and to generate texture and a sense of speed corresponding to the speed of diffusion. We used layering and contrast to create depth. Stroke size, texture, and contrast helped to define a focus within each image and also suggested the viewing order for different parts of an image.

The images created with both methods are effective because they display many data values simultaneously. We know of no other methods that simultaneously display as many components of second-order tensor-valued data. Our methods also qualitatively represent the underlying phenomena intuitively and geometrically. From different perspectives, they show the data at different levels of abstraction – more qualitatively from a distance, more quantitatively up close. Finally, the images emphasize different data values to different degrees leading a viewer through the temporal process of understanding the relationship among them.

Both of our methods visually represent multi-valued data with images at a much higher spatial resolution than the data. Unfortunately, this can lead to cluttered images that are difficult and time-consuming to interpret. We believe that incorporating all of the information is important for exploring the relationships among values in multi-valued data. We have found that by attempting to design the temporal interpretation of our images we can create images that can be interpreted at different levels of detail in different amounts of time. For datasets with many more spatial samples, the challenge will become greater and our methods as they stand may not be as effective, but the design methodology should still apply.

We are in the process of performing usability studies comparing these and other methods. Our subjects are spinal-cord researchers and we plan to report results in a future publication.

We have made some very specific visual choices in generating our examples. Hopefully, those choices will stimulate others to make different choices as a further exploration of this fascinating domain. We have only begun to explore the space of visualization methods that borrow conceptually from painting. We believe that methods that incorporate concepts from art and other fields will continue to lead to many new, effective, and powerful visualization algorithms for multi-valued scientific data from many disciplines.

8 Acknowledgements

This work was supported by NSF (CCR-96-19649). Additional support was provided by NSF (ASC-89-20219) as part of the NSF STC for Computer Graphics and Scientific Visualization, by the DOE (DE-FG03-92ER25134) as part of the Center for Research in Computational Biology, and by the National Institute on Drug Abuse, the National Institute of Mental Health, and the National Science Foundation as part of the Human Brain Project. All opinions, findings, conclusions, or recommendations expressed in this document are those of the author(s) and do not necessarily reflect the views of the sponsoring agencies. Thanks also to Scott Fraser, and Al Barr for access to great facilities, including the MRI machine and the graphics lab, and for a stimulating environment.

References

- [1] Eric T. Ahrens, David H. Laidlaw, Carol Readhead, Celia F. Brosnan, Scott E. Fraser, and Russell E. Jacobs. MR microscopy of transgenic mice that spontaneously acquire experimental allergic encephalomyelitis. *Magnetic Resonance in Medicine*, 40(1):119–132, 1998.
- [2] E.T. Ahrens, D.H. Laidlaw, C. Readhead, R.E. Jacobs, P.T. Narasimhan, and S.E. Fraser. Quantitative diffusion tensor imaging of the transgenic experimental allergic encephalomyelitis mouse spinal cord. In *Proc. ISMRM, 5th Scientific Meeting*, number 506, Vancouver, 1997.
- [3] E.T. Ahrens, C. Readhead, D. H. Laidlaw, C. Brosnan, R. E. Jacobs, and S. E. Fraser. Magnetic resonance microscopy of transgenic mice expressing a myelin basic protein-specific t cell receptor. *Journal of Neurochemistry*, 69:S124, 1997.
- [4] P. J. Basser, J. Mattiello, and D. Lebihan. Estimation of the effective self-diffusion tensor from the NMR spin-echo. *J. Magnetic Resonance B*, 103(3):247–254, march 1994.
- [5] Peter J. Basser and Carlo Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion tensor MRI. *J. Magnetic Resonance B*, 111(3):209–219, june 1996.
- [6] J. Crank. *The Mathematics of Diffusion*. Oxford Univesity Press, Oxford, England, 1975.
- [7] Thierry Delmarcelle. *The Visualization of Second-Order Tensor Fields*. PhD thesis, Stanford University, 1994.
- [8] Thierry Delmarcelle and Lambertus Hesselink. Visualizing second-order tensor fields with hyperstream lines. *IEEE Computer Graphics and Applications*, 13(4):25–33, July 1993.
- [9] R. R. Dickinson. A unified approach to the design of visualization software for the analysis of field problems. In *Three-dimensional Visualization and Display Technologies (Proceedings of SPIE)*, volume 1083, pages 173–180, 1989.
- [10] Richard S. Gallagher, editor. *Computer Visualization: Graphics Techniques for Scientific and Engineering Analysis*. CRC Press, Boca Raton, 1995.
- [11] Joan Goverman, Andrea Woods, Lisa Larson, Leslie P. Weiner, Leroy Hood, and Dennis M. Zaller. Transgenic mice that express a myelin basic protein-specific T cell receptor develop spontaneous autoimmunity. *Cell*, 72:551–560, 1994.
- [12] R. B. Haber. Visualization techniques for engineering mechanics. *Computing Systems in Engineering*, 1(1):37–50, 1990.
- [13] Robert B. Haber and David A. McNabb. Visualization idioms: A conceptual model for scientific visualization systems. *Visualization in scientific computing*, pages 74–93, 1990.
- [14] Paul E. Haeberli. Paint by numbers: Abstract image representations. In Forest Baskett, editor, *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 207–214, August 1990.
- [15] Victoria Interrante, Henry Fuchs, and Stephen M. Pizer. Conveying the 3D shape of smoothly curving transparent surfaces via texture. *IEEE Transactions on Visualization and Computer Graphics*, 3(2), April–June 1997. ISSN 1077-2626.
- [16] Victoria L. Interrante. Illustrating surface shape in volume data via principal direction-driven 3D line integral convolution. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 109–116. ACM SIGGRAPH, Addison Wesley, August 1997. ISBN 0-89791-896-7.
- [17] Jock Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5(2):110–141, 1986.
- [18] Barbara J. Meier. Painterly rendering for animation. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 477–484. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.
- [19] Clifford A. Pickover and Stuart K. Tewksbury, editors. *Frontiers of Scientific Visualization*. Wiley, New York, 1994.
- [20] Carlo Pierpaoli and Peter. J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36:893–906, 1996.
- [21] Penny Rheingans. Color, change, and control for quantitative data display. In *Visualization '92*, pages 252–259, 1992.
- [22] Michael P. Salisbury, Sean E. Anderson, Ronen Barzel, and David H. Salesin. Interactive pen-and-ink illustration. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 101–108. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [23] Craig Upson, Thomas Faulhaber, Jr., David Kamins, David Laidlaw, David Schlegel, Jeffrey Vroom, Robert Gurwitz, and Andries van Dam. The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics and Applications*, 9(4):30–42, July 1989.
- [24] Colin Ware and William Knight. Using visual texture for information display. *ACM Transactions on Graphics*, 14(1):3–20, January 1995.
- [25] Georges Winkenbach and David H. Salesin. Computer-generated pen-and-ink illustration. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 91–100. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [26] Georges Winkenbach and David H. Salesin. Rendering parametric surfaces in pen and ink. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 469–476. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.




Figure 1: Two methods for visualizing a slice from a 3D second-order tensor field. The data measure the molecular diffusion rate of water for each voxel in a mouse spinal cord. Each image shows half of a spinal cord section. The top images use an ellipsoid shape to represent the diffusion rate for each voxel and capture most of the information from the tensor-valued data. The bottom images use concepts from oil painting to represent the data. The images are composited from layers as shown in Fig. 9. The data values are decomposed into different components. The components are represented visually in a blue-purple underpainting; in the direction, shape, color, and transparency of a layer of elliptical strokes over that; and in the frequency of a texture applied to the elliptical strokes. The resulting image displays all seven related values from the data. Both methods show significant differences between the healthy spinal cord in the left images and the diseased spinal cord in the right images.