

Diffusion Tensor Visualization with Glyph Packing

Gordon Kindlmann and Carl-Fredrik Westin

Abstract—A common goal of multivariate visualization is to enable data inspection at discrete points, while also illustrating larger-scale continuous structures. In diffusion tensor visualization, glyphs are typically used to meet the first goal, and methods such as texture synthesis or fiber tractography can address the second. We adapt particle systems originally developed for surface modeling and anisotropic mesh generation to enhance the utility of glyph-based tensor visualizations. By carefully distributing glyphs throughout the field (either on a slice, or in the volume) into a dense packing, using potential energy profiles shaped by the local tensor value, we remove undue visual emphasis of the regular sampling grid of the data, and the underlying continuous features become more apparent. The method is demonstrated on a DT-MRI scan of a patient with a brain tumor.

Index Terms—Diffusion tensor, glyphs, particle systems, anisotropic sampling, fiber tractography.

1 INTRODUCTION

A tensor field may be visualized with a variety of methods. One of the simplest is tensor glyphs, which indicate at discrete locations the tensor eigenvalues and eigenvectors by the shape and orientation of glyph geometry [34]. Glyphs are only informative when unobscured, so they must be seeded in the field and scaled judiciously. When glyphs are placed on the sample points of a regular grid, the resulting visualization may inadvertently emphasize the regular sampling pattern, rather than the important patterns in the data.

Other tensor visualization methods excel at representing continuous field structures. Texture-based visualization methods, such as Line Integral Convolution (LIC) [11] are inherently continuous, because field attributes are integrated into a dense image. Lacking discrete geometry, however, texture-based methods have difficulty conveying relative scale or size information, and the textures are more challenging to visualize in three dimensions. Hyperstreamlines [12] numerically integrate tensor eigenvectors and thus represent particular continuous field features, but are only meaningful where eigenvectors are well-defined.

The goal of this work is to bridge the intrinsically discrete nature of glyph-based visualization with the continuous character of texture and hyperstreamline methods, by densely “packing” glyphs into the field, irrespective of the original data sampling grid. The packing is calculated with a particle system in which inter-particle interactions are determined by a potential energy function derived from the tensor field. Rather than emphasizing the regular sampling pattern of the data points, the resulting visualization can highlight the underlying continuous features of the tensor field. The effect is similar to that achieved by reaction-diffusion texture-based visualization [22], but with reduced computational cost.

Although theoretically applicable to other sources of symmetric second-order positive-definite tensor fields, our work so far has targeted diffusion tensor magnetic resonance imaging (DT-MRI). DT-MRI is a popular and effective means of assessing white matter structure in the central nervous system, in which the coherent organization of axons gives rise to anisotropy (directional variation in diffusivity) [5]. By fitting a second-order tensor model to the diffusion-weighted measurements, DT-MRI can quantify (at the coarse scale of MRI acquisition) the shape and orientation of white matter [4]. To

some extent, our visualization method is able to indicate the course of white pathways, even though this task is normally considered the domain of fiber tractography (based on hyperstreamlines) [29].

2 RELATED WORK

Our approach strives to mimic texture-based visualization methods, and adapts systems developed for particle-based surface modeling and anisotropic mesh generation. Tensor field interpolation generates tensor values at the arbitrary locations of particles inside the domain. This section considers these areas of previous work in turn.

Texture-based vector field visualization methods such as spot noise [45] and Line Integral Convolution (LIC) [11] have the compelling advantage of highlighting the flow pattern with a detailed, dense, and continuous pattern. LIC has been generalized to tensor visualization, either with a two-pass integration over primary and secondary eigenvectors by Hsu [19], or integration over a tensor-based polygonal probe by Zheng and Pang [50]. Sigfridsson *et al.* perform local frequency-space noise texture filtering according to tensor values [39]. Hotz *et al.* modulate the density, size, and blurring extent of spot noise to visualize the tensor field in terms of a deformation of an underlying physical texture domain [18].

Other previous work blends attributes of both texture and glyph methods, by placing or rendering glyphs in a more data-driven way. The stochastic streamline placement optimization of Turk and Banks is also demonstrated for glyphs, resulting in a dense vector visualization that maintains visual clarity in two dimensions [44]. Laidlaw *et al.* generate multi-layered diffusion tensor visualizations including glyphs stochastically placed so as to minimize overlap, which are additionally textured to indicate their out-of-plane component [24]. Kirby *et al.* use a similar placement strategy for glyphs representing vector and tensor attributes of complex flow fields [23]. Hausner describes optimized packing of rectangular patches to create decorative mosaics that respect the orientation structure of the underlying image [16]. The tensor splat method of Bengner *et al.* converts tensor values into tuned Gabor functions that are composited into 2-D and 3-D textures [7, 6].

Texture-based methods may also avoid noise entirely by leveraging reaction-diffusion (RD) equations [41], used previously in computer graphics [42, 48]. RD textures have been adapted to tensor visualization by Kindlmann *et al.* [22], and to flow fields and flow uncertainty by Sanderson *et al.* [35]. The effectiveness of RD textures lies in how the simulation of chemical morphogens naturally gives rise to an organized pattern of spots reflecting the local orientation, anisotropy, and scale of the tensor field. Figure 1 illustrates the natural packing of RD spots, resembling a glyph packing. Indeed, Sanderson *et al.* show how the centroids of RD spots may be used to place glyphs [35]. The utility of RD textures for visualization, however, is hampered by their sensitivity to unintuitive parameter settings, and significant computational cost. In particular, each texture spot must be large compared to the pixels or voxels composing the simulation grid, thus the quality of

- Gordon Kindlmann is with the Laboratory of Mathematics in Imaging, Department of Radiology, Brigham and Women's Hospital, Harvard Medical School, E-mail: gk@bwh.harvard.edu.
- Carl-Fredrik Westin is the Director of the Laboratory of Mathematics in Imaging, Department of Radiology, Brigham & Women's Hospital, Harvard Medical School, E-mail: westin@bwh.harvard.edu.

Manuscript received 31 March 2006; accepted 1 August 2006; posted online 6 November 2006.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org.




Fig. 1. Ellipsoid glyphs (left) and reaction-diffusion textures (right) on synthetic (top) and measured DT-MRI data (bottom). Reproduced from [22].

the visualization scales with only the square root or cube root of the memory and processing time. We avoid this cost by representing the glyph locations directly as particles, instead of indirectly as sampled intensity patterns.

Point-based modeling in computer graphics implicitly represents surfaces in a compact and flexible manner [2, 3]. One important line of work uses dynamic particle systems to sample or redistribute points within a surface [43, 40, 47, 17]. Meyer *et al.* refine the method of Heckbert [17] with scale-invariant energy functions, reducing the number of necessary particle system parameters and facilitating curvature-dependent surface sampling [26]. Glyph packing also benefits from research in unstructured anisotropic mesh generation, using a tensor formulation of the desired mesh anisotropy. The bubble packing methods of Shimada *et al.* [37] and Yamakawa and Shimada [49] use spring forces between tensor ellipsoids to equalize their spacing along the edges, faces, and interiors of the mesh domain. Bossen and Heckbert measure edge lengths according to local tensors to drive anisotropic mesh refinement and smoothing [10]. Bossen also uses particle systems to optimize the distribution of mesh vertices, based on an anisotropic potential energy function [9].

Placing glyphs at arbitrary locations in the field requires some method of interpolating tensor data, which is itself an ongoing area of work. Besides the trilinear per-coefficient interpolation standard in many hyperstreamline implementations [36], the various mathematical approaches for reconstruction of tensor fields include Wiener amalgam spaces [1], B-splines [31], NURBS [32], and PDEs [46]. More specialized approaches seek to preserve tensor invariants such as trace and anisotropy, or positive-definiteness (a physical constraint on diffusion), in the course of interpolation [14, 28].

3 METHODS

We generate glyph packings as the stable solution to the particle system described in the following sections. A basic ingredient is the glyph scaling α , which transforms from the units of the tensor value (in our case diffusivity) to units of length in the image domain. In the interests of tractability, the packing computation assumes ellipsoidal glyph geometry, described by $\alpha\mathbf{D}\mathbf{S}$ for tensor \mathbf{D} , where \mathbf{S} is the (dimensionless) unit sphere [36]. However, the glyph scaling and geometry of the packing computation need not restrict what is used in the subsequent visualization: the glyph scaling for visualization can be smaller (to better see between glyphs), and the glyph geometry can be different (for example using superquadrics [21] instead of ellipsoids).

3.1 Tensor Interpolation

The purpose of tensor interpolation is to create from the sampled data a continuous function \mathbf{D} that returns the tensor value $\mathbf{D}(\mathbf{p})$ at location \mathbf{p} . In this work, we use a cubic spline to create a C^2 field that interpolates the estimated tensors [31]. This is implemented by numerically pre-filtering the data so that subsequent convolution-based separable reconstruction with the uniform cubic B-spline can be local (four sample support along each axis) as well as interpolating [27]. The interpolation does not, however, preserve or respect any tensor invariants such as trace or anisotropy. Two points of consideration relative to our application domain and visualization purpose suggest this may not be a significant problem.

First, the MRI scanner has likely *already* interpolated the diffusion-weighted images (DWIs) that are the basis of tensor estimation. Standard echo-planar MRI acquires DWIs on a regular grid in frequency space [25]. The grid is usually “zero-filled” by the scanner to a larger grid prior to the Fourier transform that creates the DWIs [8]. Zero-filling is effectively the same as interpolating and upsampling the DWIs in the spatial domain with the sinc function [15]. Notably, upsampling the DWIs does not preserve any scalar invariants of tensors subsequently estimated from the DWIs. In our opinion, this lessens the urgency to preserve invariants during interpolation for later computations (such as glyph packing). More specialized interpolation schemes may be important for visualizing tensor fields in other application domains.

Second, our purpose is more qualitative than quantitative: to create a visual effect that indicates large-scale structures, rather than geometrically modeling their course (as with tractography) or measuring them (with tensor attributes). We have not found the slight numeric differences between reconstruction methods to significantly affect the result, so long as they are interpolating and smooth.

3.2 Tensor-based Potential Energy

The behavior of particles in our system is determined by the forces acting upon them. The most important of these are the forces between particles, which gives rise to the glyph packing. The inter-particle forces are created by a potential energy field around each particle, shaped by the local tensor value. The energy E_{ab} at position \mathbf{p}_a due to the potential energy field around a particle at \mathbf{p}_b is the composition of functions \mathbf{g} , $|\cdot|$, and ϕ :

$$E_{ab} = \phi(|\mathbf{g}(\mathbf{y}_{ab})|) = \phi(r_{ab}) \quad (1)$$

$$\mathbf{y}_{ab} = \mathbf{p}_a - \mathbf{p}_b \quad (2)$$

$$\mathbf{x}_{ab} = \mathbf{g}(\mathbf{y}_{ab}) = \frac{\mathbf{D}_{ab}^{-1}\mathbf{y}_{ab}}{2\alpha}; \quad (3)$$

$$\mathbf{D}_{ab} = \mathbf{D}\left(\frac{\mathbf{p}_a + \mathbf{p}_b}{2}\right) \quad (4)$$

$$r_{ab} = |\mathbf{x}_{ab}| \quad (5)$$

Note that by construction, \mathbf{g} inverts the transform that creates ellipsoidal glyphs from spheres. Conceptually, \mathbf{g} maps vectors \mathbf{y}_{ab} in the field of anisotropic tensors to vectors \mathbf{x}_{ab} in an isotropic space, in which particles have a rotationally symmetric potential energy profile ϕ . O’Donnell *et al.* also use the tensor inverse as a metric for computing geodesics under diffusion tensor warping [30]. One way to characterize the local tensor value \mathbf{D}_{ab} is to sample the tensor field at the midpoint between the two particles, which ensures symmetry $E_{ab} = E_{ba}$. In Equation 3, the factor of two in the denominator of \mathbf{g} allows mutually tangent glyphs to map to mutually tangent spheres of radius $1/2$, with unit distance between centers (relevant for the later definition of ϕ).

The inter-particle force is determined by the derivative of energy E_{ab} with respect to \mathbf{y}_{ab} , which (by the chain rule) is the product of the derivatives of ϕ , $|\cdot|$, and \mathbf{g} (here omitting the ab subscripts for clarity):

$$\frac{dE}{d\mathbf{y}} = \frac{dE}{dr} \frac{dr}{d\mathbf{x}} \frac{d\mathbf{x}}{d\mathbf{y}} = \phi'(|\mathbf{x}|) \frac{\mathbf{x}^T}{|\mathbf{x}|} \frac{1}{2\alpha} \left(\mathbf{D}^{-1} + \frac{d\mathbf{D}^{-1}}{d\mathbf{y}} \mathbf{y} \right). \quad (6)$$

The second term in the last factor of Equation 6 involves a third-order tensor $d\mathbf{D}^{-1}/d\mathbf{y}$: the change in the inverse of the (second-order) tensor field with respect to \mathbf{y} . The product of this with vector \mathbf{y} is again a second-order tensor. Index notation makes the details of third-order tensor multiplication more explicit (here the ij subscripts index the vector and tensor coefficients):

$$\frac{\partial E}{\partial y_i} = \frac{\phi'(|\mathbf{x}|)}{2\alpha|\mathbf{x}|} \sum_{j=1}^3 (D_{ij}^{-1}x_j + \sum_{k=1}^3 D_{ijk}^{-1}x_j y_k), \quad (7)$$

where $D_{ijk}^{-1} = \partial D_{ij}^{-1} / \partial y_k$. Bossen mentions the theoretical role of the tensor gradient in anisotropic meshing [9]. As in that work, our computations avoid the complexity of the inverse tensor gradient by treating the tensor field as locally constant, without apparent impact on the quality of the results. The stability of the particle system, even within real-world varying tensor fields (where the tensor gradient is non-zero), is possibly related to the smoothness imposed by our spline interpolation, though this is a topic for future investigation.

With the assumption of local tensor constancy, the force \mathbf{f}_{ab} on a particle at \mathbf{p}_a from a particle at \mathbf{p}_b is (written as a column vector):

$$\mathbf{f}_{ab} = - \left(\frac{dE_{ab}}{d\mathbf{y}_{ab}} \right)^T = - \frac{\phi'(|\mathbf{x}_{ab}|)}{2\alpha|\mathbf{x}_{ab}|} \mathbf{D}_{ab}^{-1} \mathbf{x}_{ab} \quad (8)$$

where $(\mathbf{D}^{-1})^T = \mathbf{D}^{-1}$ because diffusion tensors are symmetric. Note that the force between two particles is not in general aligned with the vector between them, which is an unusual property of our particle system. This is analogous to how an object's surface normals are transformed by the inverse-transpose of the object transform, as is commonly used for instancing in ray-tracing [38].

The expense of computing \mathbf{D}_{ab}^{-1} for every pair of particles is significant, due to the tensor inverse, and the convolution-based tensor interpolation (as per Section 3.1) at the mid-points of all pairs of potentially interacting particles. In the early iterations of the particle system, a fast approximation to \mathbf{D}_{ab}^{-1} suffices:

$$\mathbf{D}_{ab}^{-1} \approx \frac{\mathbf{D}^{-1}(\mathbf{p}_a) + \mathbf{D}^{-1}(\mathbf{p}_b)}{2}, \quad (9)$$

where $\mathbf{D}^{-1}(\mathbf{p})$ is the convolution-based interpolation at \mathbf{p} within a pre-computed field of tensor inverses. The approximation also assumes local linear variation in $\mathbf{D}(\mathbf{p})^{-1}$.

The remaining piece of Equation 8 is the underlying potential energy function ϕ and force function ϕ' . In typical glyph-based tensor visualizations, the density and location of glyphs is determined by their placement at dataset nodes or sampling points, and the size of the glyphs is then adjusted by the user to minimize glyph overlap and clutter while maintaining visibility. In our approach, however, the scaling parameter is primary: the user chooses α , and then the packing is calculated around that particular scaling. Thus, it makes sense to employ an energy function ϕ with an intrinsic scale, in contrast to previous particle systems wherein the scale-invariance of ϕ was shown to be beneficial [26]. Our potential energy is based on Hooke's law spring model [37, 49], though we take care to ensure that ϕ' is a C^1 function. Specifically, the linear ramp of ϕ' is extended with a Hermite spline to a tunable-width interval $[1, 1 + \gamma]$ of slight attractive force (Figure 2):

$$\phi'(r) = \begin{cases} r-1 & 0 < r < 1 \\ (r-1)(1+\gamma-r)^2/\gamma^2 & 1 \leq r \leq 1+\gamma \\ 0 & r > 1+\gamma \end{cases} \quad (10)$$

3.3 Implementation Details

In addition to the inter-particle forces that generate the packing, other forces and constraints assist the computation. For numeric stability, a drag coefficient creates an additional force opposing the particle motion:

$$\mathbf{f} = -C_{\text{drag}} \frac{d\mathbf{p}}{dt} \quad (11)$$




Fig. 2. Modified spring potential (a), and resulting force function (b), as a function of radius, with $\gamma = 0.5$.

Because diffusion tensor fields are the result of a model fitting (tensor estimation) applied to the acquired DWIs [4], there are locations with insufficient signal (in skull and air) for the tensor model to be estimated meaningfully. Our tensor estimation generates an accompanying mask image M to indicate where the DWI signal meets some confidence threshold. The gradient of the confidence mask is used to generate an additional force that pushes particles towards the mask interior:

$$\mathbf{f} = C_{\text{mask}} \nabla M(\mathbf{p}) \quad (12)$$

Finally, to generate results on two-dimensional slices through the field, particle forces are computed in three-dimensions, but particle positions are projected back onto the slice plane at every iteration.

The second-order ordinary differential equation governing the motion of the particle at position \mathbf{p}_a is an expression of $f = ma$ in terms of Equations 8 (summing over all particles), 11 and 12, with the simplification that particles have unit mass:

$$\sum_{b \neq a} \mathbf{f}_{ab} - C_{\text{drag}} \frac{d\mathbf{p}}{dt} + C_{\text{mask}} \nabla M(\mathbf{p}) = \frac{d^2\mathbf{p}}{dt^2} \quad (13)$$

We have found Euler integration to be sufficient for our purposes (with a small enough time step), because the computational goal is a minimum-energy set of particle locations, rather than the paths the particles take through the field towards their resting positions.

Additional information about the tensor field is used for the initialization and computation of the particle system. The dataset is traversed to find the maximal determinant \det_{max} and maximal eigenvalue λ_{max} of all tensors. Particles are then initialized at random locations, as filtered by two rejection tests. Candidate locations outside the confidence mask M are rejected. To approximate the correct density of the final packing, and knowing that glyph volume is proportional to tensor determinant, candidate location \mathbf{p} is rejected with probability $\det(\mathbf{D}(\mathbf{p}))/\det_{\text{max}}$.

Spatial binning is an important efficiency technique for speeding up the evaluation of Equation 13, by avoiding the computation of forces between particles too distant to interact [17]. Computationally, bins are dynamically updated lists of particles that occupy a particular cell within a grid overlaid on the field. Knowing the glyph scaling α and the force function support $[0, 1 + \gamma]$, the edge length of the spatial bin is set to $2\alpha\lambda_{\text{max}}(1 + \gamma)$. All points in a bin can then interact only with points in the same bin, and points in the immediately neighboring bins. Spatial binning permits a 3-D particle system to be simulated in roughly the same time as a 2-D system with the same number of particles, which is not true of reaction-diffusion textures. In our multi-threaded implementation, the spatial binning takes the additional role of quantizing work between threads, by assigning bins to worker threads on a first-come-first-served basis.

4 RESULTS

The glyph packing method can be evaluated with comparison to visualization of the same data using a regular grid of glyphs. Figure 3 demonstrates glyph packing on a simple synthetic dataset. On the regular grid (Figure 3(a)), some tensor orientations generate large areas of consistent glyph alignment and overlap, creating the false impression of a spatial field variation besides smooth orientation change. Glyph

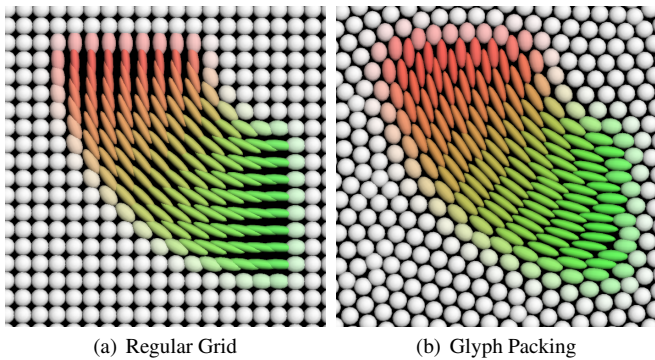


Fig. 3. Comparison of standard glyphs and glyph packing on synthetic data.

packing (Figure 3(b)) tends to consistently minimize the amount of glyph overlap, creating a more uniform field appearance. The packing can, however, also lead to a new ordered sampling grid, hexagonal instead of rectangular, which potentially could be distracting. Although this is an issue that warrants further investigation, our experience has been that in real-world datasets the hexagonal grid appearance, when it appears, is only a minor drawback.

Neurosurgeons are making increasing use of DT-MRI to plan tumor resection brain surgeries, both to assess the extent of the tumor, and to plan a minimally damaging approach through the brain [20, 13, 6]. Our neurosurgery collaborator related how all additional pieces of information about white matter structure in and around the tumor can help decide a course of action, given the amount of risk and uncertainty that accompanies invasive surgery. As a preliminary investigation of the utility of glyph packing for neurosurgical applications, we generated glyphs packings in 2D and 3D in a DT-MRI scan of a patient with a large anaplastic (malignant) oligodendroglioma near the left frontal lobe of the brain. The patient signed a statement of informed consent, and all imaging procedures followed institutional guidelines for patient safety. The diffusion tensor volume was estimated from 30 DWI at $b = 1000\text{s/mm}^2$ with 1 non-diffusion-weighted T2 image, from a 3.0 T GE scanner with ASSET reconstruction (parallel factor 2.0), at a resolution of $2 \times 2 \times 3\text{mm}$, upsampled by the scanner to $1 \times 1 \times 3\text{mm}$.

Figure 4 shows the results of glyph packing on a slice within a healthy region (left posterior) of the brain. For this and other figures, the RGB coloring on the glyphs is the standard mapping of the principal eigenvector, determined by the absolute values of the coefficients of the principal eigenvector [33]. Note that because the glyphs were constrained to lie on the slice plane, the two-dimensional forces acting on the glyphs can not always produce a clean separation of the three-dimensional glyphs viewed from above. Compared to glyphs on a regular grid (Figure 4(a)), the glyph packing (Figure 4(b)) seems to be more effective at visually indicating the curving paths of the major white matter tracts, such as the corpus callosum, cingulum, and a U-fiber, which are labeled for reference in tractography seeded in the same region (Figure 4(c)). This 2-D glyph packing (as well as that in Figure 5) ran for 2000 iterations (usable results appeared after 400), requiring about 8 minutes of computation on a single processor 1.67 GHz PowerPC G4.

Figure 5 shows visualizations of a slice through the tumor area. Our neurosurgery collaborator felt that the glyph packing results (Figure 5(b)) better helped her to develop an understanding of the “actual tensor data” than those in Figure 5(a). The accumulation of water in the areas of edema around the tumor increases diffusivity, leading to larger glyphs that need to be spaced more sparsely than in the healthy tissue, which is handled well by glyph packing (Figure 5(b)). If the scaling was reduced so the glyphs did not overlap on the regular grid (Figure 5(a)), they would be harder to see in regions of healthy tissue. The behavior of the genu of the corpus callosum around the tumor is better shown through glyph packing. The asymmetry of the corpus callosum, in particular how the fibers are partially deflected by the

tumor, is more visible in the glyph packing, and is confirmed by tractography (Figure 5(c)). As tractography can only analyze regions of linear anisotropy, it offers no insight as to why the fibers stop near the tumor; the glyph images show that while there is moderate anisotropy, it is planar rather than linear). In fact, the pattern of planar anisotropy surrounding the tumor is an interesting feature detected by the neurosurgeon in the glyph packing visualization.

Figure 4 gives one example of glyph packing in three-dimensions. Glyphs were densely packed in the tensor volume, but only those meeting an FA threshold ($FA > 0.25$) are shown. The brain region here is the corpus callosum (red and orange) and a cingulum bundles (blue and green) orthogonally oriented but adjacent to the corpus. The curving geometry of these and nearby white matter features is easier to appreciate from the glyph packings. The 3-D glyph packing (Figure 6) ran for 4000 iterations, requiring about 1.5 hours of computation on a dual processor 2.0 GHz PowerPC G5.

5 CONCLUSIONS

Glyph packing uses a particle system with anisotropic potential energy profiles to arrange glyphs into a dense pattern that displays some of the visual continuity of texture-based visualizations, while maintaining the ability to discern the full tensor information at each glyph. Our neurosurgery collaborator felt that applied to DT-MRI, the glyph packings were a more direct and informative expression of the tensor data. Although a formal user study would be a fitting way to evaluate our claims, we are still at the early stage of refining the method and iterating with the surgeon to learn about what visualization strategies are most successful. Perceptual studies could help determine how the method can be best augmented or combined with other existing methods (such as colormaps or isosurfacing) to convey as much relevant information as possible.

The computational bottleneck in our current implementation is the convolution-based sampling (Section 3.1). Under consideration are different means of accelerating this process, such as using a lower-resolution tensor volume, or selectively re-using sampled tensors between iterations. We plan to extend the glyph packing method to sample implicit surfaces in the field, either isosurfaces of important tensor attributes (such as anisotropy), or distance volumes around segmented objects of interest (such as tumors). This will likely entail analyzing the differential properties of the surface [26]. Such an analysis might complement a further exploration of the inverse tensor gradient arising in Equation 6, to determine the circumstances under which its handling is required for stable solutions.

Information about our open-source implementation is available at <http://lmi.bwh.harvard.edu/~gk/vis06>.


ACKNOWLEDGEMENTS

We gratefully acknowledge the help and feedback from our neurosurgical collaborator Alexandra Golby, M.D., and data acquisition by Stephen Whalen, in the Department of Neurosurgery, Brigham and Women’s Hospital, Harvard Medical School. This work supported by NIH grants NIBIB T32-EB002177, P41-RR13218 (NAC), and U41-RR019703-01A2, as well as the Brain Science Foundation.

REFERENCES

- [1] A. Aldroubi and P. Basser. Reconstruction of vector and tensor fields from sampled discrete data. *Contemporary Mathematics*, 247:1–15, 1999.
- [2] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. Silva. Computing and rendering point set surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 9(1):3–15, Jan 2003.
- [3] N. Amenta and Y. Kil. Defining point-set surfaces. In *Proceedings ACM SIGGRAPH 2004*, pages 264–270, 2004.
- [4] P. Basser and D. Jones. Diffusion-tensor MRI: Theory, experimental design and data analysis - A technical review. *Nuclear Magnetic Resonance in Biomedicine*, 15:456–467, 2002.
- [5] C. Beaulieu. The basis of anisotropic water diffusion in the nervous system - A technical review. *Nuclear Magnetic Resonance in Biomedicine*, 15:435–455, 2002.


- [6] W. Benger, H. Bartsch, H.-C. Hege, H. Kitzler, A. Shumilina, and A. Werner. Visualizing neuronal structures in the human brain via diffusion tensor MRI. *International Journal of Neuroscience*, 116(4):461–514, April 2006.
- [7] W. Benger and H.-C. Hege. Tensor splats. In *Proceedings of SPIE Visualization and Data Analysis 2004*, pages 151–162, June 2004.
- [8] M. Bernstein, K. King, and X. Zhou. *Handbook of MRI Pulse Sequences*. Academic Press, 2004.
- [9] F. Bossen. Anisotropic mesh generation with particles. Technical Report CMU-CS-96-13, School of Computer Science, Carnegie Mellon University, May 1996.
- [10] F. J. Bossen and P. S. Heckbert. A pliant method for anisotropic mesh generation. In *The 5th International Meshing Roundtable*, pages 63–74, 1996.
- [11] B. Cabral and C. Leedom. Imaging vector fields using line integral convolution. In *Proceedings of ACM SIGGRAPH 1993*, pages 263–270. Addison Wesley, 1993.
- [12] T. Delmarcelle and L. Hesselink. Visualizing second-order tensor fields with hyper streamlines. *IEEE Computer Graphics and Applications*, 13(4):25–33, 1993.
- [13] A. Field, A. Alexander, Y.-C. Wu, K. Hasan, B. Witwer, and B. Badie. Diffusion tensor eigenvector directional color imaging patterns in the evaluation of cerebral white matter tracts altered by tumor. *Journal of Magnetic Resonance Imaging*, 20:555–562, Oct 2004.
- [14] P. Fletcher and S. Joshi. Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In *ECCV 2004 Workshop on Computer Vision Approaches to Medical Image Analysis (CVAMIA), LNCS 3117*, pages 87–98. Springer-Verlag, 2004.
- [15] R. Gonzalez and R. Woods. *Digital Image Processing*. Addison-Wesley Publishing Company, Reading, MA, 2nd edition, 2002.
- [16] A. Hausner. Simulating decorative mosaics. In *Proceedings of ACM SIGGRAPH 2001*, volume 30, pages 573–578. Addison Wesley, 2001.
- [17] P. Heckbert. Fast surface particle repulsion. In *ACM SIGGRAPH 1997 Course Notes (New Frontiers in Modeling and Texturing)*, pages 95–114, August 1997.
- [18] I. Hotz, L. Feng, H. Hagen, B. Hamann, K. Joy, and B. Jeremic. Physically based methods for tensor field visualization. In *Proceedings of IEEE Visualization 2004*, pages 123–130, 2004.
- [19] E. Hsu. Generalized line integral convolution rendering of diffusion tensor fields. In *Proceedings of the 9th Scientific Meeting and Exhibition of the International Society for Magnetic Resonance in Medicine (ISMRM)*, page 790, Glasgow, United Kingdom, 2001.
- [20] B. J. Jellison, A. S. Field, J. Medow, M. Lazar, M. S. Salamat, and A. L. Alexander. Diffusion tensor imaging of cerebral white matter: A pictorial review of physics, fiber tract anatomy, and tumor imaging patterns. *American Journal of Neuroradiology*, 25:356–369, 2004.
- [21] G. Kindlmann. Superquadric tensor glyphs. In *Proceedings of IEEE TVCG/EG Symposium on Visualization 2004*, pages 147–154, May 2004.
- [22] G. Kindlmann, D. Weinstein, and D. Hart. Strategies for direct volume rendering of diffusion tensor fields. *IEEE Transactions on Visualization and Computer Graphics*, 6(2):124–138, April-June 2000.
- [23] M. Kirby, H. Marmanis, and D. Laidlaw. Visualizing multivalued data from 2D incompressible flows using concepts from painting. In *Proceedings of IEEE Visualization 1999*, pages 333–340, 1999.
- [24] D. Laidlaw, E. Ahrens, D. Kremers, M. Avalos, R. Jacobs, and C. Readhead. Visualizing diffusion tensor images of the mouse spinal cord. In *Proceedings of IEEE Visualization 1998*, pages 127–134, 1998.
- [25] Z.-P. Liang and P. Lauterbur. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. Wiley-IEEE Press, 1999.
- [26] M. Meyer, P. Georgel, and R. Whitaker. Robust particle systems for curvature dependent sampling of implicit surfaces. In *Proceedings Shape Modeling and Applications, 2005 International Conference*, pages 124–133, June 2005.
- [27] D. Mitchell and A. Netravali. Reconstruction filters in computer graphics. In *Proceedings of ACM SIGGRAPH 1988*, pages 221–228, August 1988.
- [28] M. Moakher and P. Batchelor. Symmetric positive-definite matrices: From geometry to applications and visualization. In J. Weickert and H. Hagen, editors, *Visualization and Image Processing of Tensor Fields*, pages 285–298. Springer, 2006.
- [29] S. Mori and P. V. Zijl. Fiber tracking: Principles and strategies - A technical review. *Nuclear Magnetic Resonance in Biomedicine*, 15:468–480, 2002.
- [30] L. O'Donnell, S. Haker, and C.-F. Westin. New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. In *Proceedings of the 5th Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 459–466, September 2002.
- [31] S. Pajevic, A. Aldroubi, and P. Basser. A continuous tensor field approximation of discrete DT-MRI data for extracting microstructural and architectural features of tissue. *Journal of Magnetic Resonance*, 154:85–100, 2002.
- [32] S. Pajevic, A. Aldroubi, and P. Basser. Continuous tensor field approximation of diffusion tensor MRI data. In J. Weickert and H. Hagen, editors, *Visualization and Image Processing of Tensor Fields*, pages 299–314. Springer, 2006.
- [33] S. Pajevic and C. Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, 42(3):526–540, 1999.
- [34] F. Post, F. Post, T. van Walsum, and D. Silver. Iconic techniques for feature visualization. In *Proceedings of IEEE Visualization 1995*, pages 288–295, 1995.
- [35] A. Sanderson, M. Kirby, and C. Johnson. Display of vector fields using a reaction diffusion model. In *Proceedings of IEEE Visualization 2004*, pages 115–122, October 2004.
- [36] W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit: An Object Oriented Approach to Graphics*. Kitware, Inc., Clifton Park, New York, 3rd edition, 2004.
- [37] K. Shimada, A. Yamada, and T. Itoh. Anisotropic triangulation of parametric surfaces via close packing of ellipsoids. *International Journal of Computational Geometry and Applications*, 10(4):400–424, 2000.
- [38] P. Shirley and R. Morley. *Realistic Ray Tracing*. AK Peters, Wellesley, Massachusetts, 2003.
- [39] A. Sigfridsson, T. Ebbens, Heiberg, and L. Wigström. Tensor field visualization using adaptive filtering of noise fields combined with glyph rendering. In *Proceedings IEEE Visualization 2002*, pages 371–378, Boston, Massachusetts, 2002.
- [40] R. Szeliski, D. Tonnesen, and D. Terzopoulos. Modeling surfaces of arbitrary topology with dynamic particles. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '93)*, pages 140–152, New York, June 1993.
- [41] A. Turing. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London*, 237(B):37–72, 1952.
- [42] G. Turk. Generating textures on arbitrary surfaces using reaction-diffusion. In *Proceedings of ACM SIGGRAPH 1991*, volume 25, pages 289–298. Addison Wesley, 1991.
- [43] G. Turk. Re-tiling polygonal surfaces. In *Proceedings ACM SIGGRAPH 1992*, pages 55–64, July 1992.
- [44] G. Turk and D. Banks. Image-guided streamline placement. In *Proceedings of ACM SIGGRAPH 1996*, volume 30, pages 453–460. Addison Wesley, 1996.
- [45] J. van Wijk. Spot noise: Texture synthesis for data visualization. In *Proceedings of ACM SIGGRAPH 1991*, volume 25, pages 309–318. Addison Wesley, 1991.
- [46] J. Weickert and M. Welk. Tensor field interpolation with PDEs. In J. Weickert and H. Hagen, editors, *Visualization and Image Processing of Tensor Fields*, pages 315–325. Springer, 2006.
- [47] A. Witkin and P. Heckbert. Using particles to sample and control implicit surfaces. In *Proceedings ACM SIGGRAPH '94*, volume 28, pages 269–277, 1994.
- [48] A. Witkin and M. Kass. Reaction-diffusion textures. In *Proceedings of ACM SIGGRAPH 1991*, volume 25, pages 299–308, 1991.
- [49] S. Yamakawa and K. Shimada. High quality anisotropic tetrahedral mesh generation via packing ellipsoidal bubbles. In *The 9th International Meshing Roundtable*, pages 263–273, 2000.
- [50] X. Zheng and A. Pang. Hyperlic. In *Proceedings of IEEE Visualization 2003*, pages 249–256, Seattle, Washington, 2003.




(a) Regular Grid




(a) Regular Grid




(b) Glyph Packing



(b) Glyph Packing




(c) Fiber Tractography




(c) Fiber Tractography

Fig. 4. Standard glyphs (a), glyph packing (b), and fiber tractography (c) on a slice through a healthy region of the brain. CC = corpus callosum; CB = cingulum, UF = U-fiber

Fig. 5. Standard glyphs (a), glyph packing (b), and fiber tractography (c) on a slice through the brain tumor. The region of edema surrounds the tumor.



(a) Regular Grid



(b) Particle System

Fig. 6. Comparison of glyphs on 3D grid and from Particle System